MODEL THEORY

STEFAN GESCHKE

Contents

1. First order logic

1.1. **Structures.** A **vocabulary** τ is a set consisting of relation symbols, function symbols and constant symbols. We will usually use relation symbols such as P, Q, R, $\leq$, ..., function symbols such as f, g, h, $\cdot$, $+$, ..., and constant symbols such as c, d, 0, 1, ... To every relation symbol and every function symbol there is a natural number ≥ 1 attached to it, the **arity** of the symbol.

Now fix a vocabulary τ. A **structure** $\mathcal{A}$ for τ (a τ**-structure**) is a nonempty set A together with

(S1) relations $R^{\mathcal{A}} \subseteq A^n$ for every n-ary relation symbol $R \in \tau$,
(S2) functions $f^{\mathcal{A}} : A^m \to A$ for every m-ary function symbol $f \in \tau$ and
(S3) constants $c^{\mathcal{A}} \in A$ for every constant symbol $c \in \tau$.

A structure $\mathcal{A}$ is often identified with its underlying set A. Sometimes we will denote the interpretation $R^{\mathcal{A}}$, $f^{\mathcal{A}}$ or $c^{\mathcal{A}}$ of a symbol R, f or c by R, f, respectively c as well.

A τ-structure $\mathcal{B}$ is a **substructure** of $\mathcal{A}$ if

(U1) the underlying set B of $\mathcal{B}$ is a subset of A,
(U2) every relation $R^{\mathcal{B}}$ is the restriction of $R^{\mathcal{A}}$ to B,
(U3) every function $f^{\mathcal{B}}$ is the restriction of $f^{\mathcal{A}}$ to B and
(U4) every constant $c^{\mathcal{B}}$ agrees with $c^{\mathcal{A}}$.

Observe that (U3) implies that B is closed under all the functions $f^{\mathcal{A}}$. Similarly, (U4) implies that all constants $c^{\mathcal{A}}$ are already elements of B.

Two τ-structures $\mathcal{A}$ and $\mathcal{B}$ are **isomorphic** if there is a bijection $i : A \to B$ such that

(I1) for every n-ary relation symbol $R \in \tau$ and all $a_1, \dots, a_n \in A$,
$$(a_1, \dots, a_n) \in R^{\mathcal{A}} \Leftrightarrow (i(a_1), \dots, i(a_n)) \in R^{\mathcal{B}},$$
(I2) for every n-ary function symbol $f \in \tau$ and all $a_1, \dots, a_n \in A$,
$$i(f^{\mathcal{A}}(a_1, \dots, a_n)) = f^{\mathcal{B}}(i(a_1), \dots, i(a_n))$$
and
(I3) for all constant symbols $c \in \tau$,
$$i(c^{\mathcal{A}}) = c^{\mathcal{B}}.$$

1.2. **Formulas.** A **first order formula** over τ is a finite sequence over the alphabet
$$\{\exists, \vee, \neg, =, (,)\} \cup \tau \cup \mathrm{Var},$$
where Var is a countably infinite set of variables. We tacitly assume that the sets $\{\exists, \vee, \neg, =, (,)\}$, τ and Var are pairwise disjoint. The variables are usually denoted by $x, y, z, \dots$. Before defining formulas, let us define **terms**.

A term over τ is a finite sequence of characters that can be obtained by finitely many applications of the following rules:

(T1) All constant symbols in τ and all variables are terms.
(T2) If $t_1, \dots, t_n$ are terms and $f \in \tau$ is an n-ary function symbol, then $f(t_1, \dots, t_n)$ is a term.

A first order formula over τ is a finite sequence of characters that can be obtained by finitely many applications of the following rules:

(F1) If t_1 and t_2 are terms over τ, then $(t_1 = t_2)$ is a formula.
(F2) If $R \in \tau$ is an n-ary relation symbol and if $t_1, \dots, t_n$ are terms over τ, then $R(t_1, \dots, t_n)$ is a formula.
(F3) If φ is a formula, then so is $\neg\varphi$.
(F4) If φ and ψ are formulas, then $(\varphi \vee \psi)$ is a formula.
(F5) If φ is a formula and x a variable, then $\exists x\varphi$ is a formula.

If φ and ψ are formulas, we use $(\varphi \wedge \psi)$, $(\varphi \to \psi)$, $(\varphi \leftrightarrow \psi)$, and $\forall x\varphi$ as abreviations for $\neg(\neg\varphi \vee \neg\psi)$, $(\neg\varphi \vee \psi)$, $\neg(\neg(\neg\varphi \vee \psi) \vee \neg(\varphi \vee \neg\psi))$, and $\neg\exists x\neg\varphi$, respectively. Also, we omit parentheses as long the readability does not suffer.

A variable x occurs **freely** in φ if x occurs outside the scope of a quantifier $\exists x$ or $\forall x$. A formula without free variables is a **sentence**.

We use the notation $\varphi(x_1, \dots, x_n)$ to indicate that $x_1, \dots, x_n$ are pairwise distinct variables and that the free variables of φ are among the x_i, $1 \leq i \leq n$.

1.3. **Semantics.** We fix a τ-structure $\mathcal{A} = (A, \dots)$. Given a τ-term $t(x_1, \dots, x_n)$ and $a_1, \dots, a_n \in A$ we define $t(a_1, \dots, a_n)$ (or, in more accurat notation, $t^{\mathcal{A}}(a_1, \dots, a_n)$) as follows:

(TA1) $x_i(a_1, \dots, a_n) = a_i$
(TA2) If $c \in \tau$ is a constant symbol, let $c(a_1, \dots, a_n) = c^{\mathcal{A}}$.
(TA3) If $f \in \tau$ is an m-ary function symbol and $t_1(x_1, \dots, x_n), \dots, t_m(x_1, \dots, x_n)$ are terms, let

$$f(t_1, \dots, t_m)(a_1, \dots, a_n) = f^{\mathcal{A}}(t_1(a_1, \dots, a_n), \dots, t_m(a_1, \dots, a_n)).$$

Now, for every formula $\varphi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in A$ we define the **validity** of $\varphi(a_1, \dots, a_n)$ in $\mathcal{A}$:

(FG1) If $t_1(x_1, \dots, x_n)$ and $t_2(x_1, \dots, x_n)$ are terms, then $(t_1 = t_2)(a_1, \dots, a_n)$ holds in $\mathcal{A}$ iff

$$t_1(a_1, \dots, a_n) = t_2(a_1, \dots, a_n).$$

(FG2) If $t_1(x_1, \dots, x_n), \dots, t_m(x_1, \dots, x_n)$ are terms and R is an m-ary relation symbol, then $R(t_1, \dots, t_m)(a_1, \dots, a_n)$ holds in $\mathcal{A}$ iff

$$(t_1(a_1, \dots, a_n), \dots, t_m(a_1, \dots, a_n)) \in R^{\mathcal{A}}.$$

(FG3) If $\varphi(x_1, \dots, x_n)$ and $\psi(x_1, \dots, x_n)$ are formulas, then $(\varphi \vee \psi)(a_1, \dots, a_n)$ holds in $\mathcal{A}$ iff at least one of $\varphi(a_1, \dots, a_n)$ and $\psi(a_1, \dots, a_n)$ holds in $\mathcal{A}$.
(FG4) If $\varphi(x_1, \dots, x_n)$ is a formula, then $\neg\varphi(a_1, \dots, a_n)$ holds in $\mathcal{A}$ iff $\varphi(a_1, \dots, a_n)$ does not hold in $\mathcal{A}$.

(FG5) If $\varphi(x, x_1, \ldots, x_n)$ is a formula, then $\exists x \varphi(a_1, \ldots, a_n)$ holds in $\mathcal{A}$ iff there is $a \in A$ such that $\varphi(a, a_1, \ldots, a_n)$ holds in $\mathcal{A}$.

If $\varphi(a_1, \ldots, a_n)$ holds in $\mathcal{A}$, we write $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$.

We extend the **model relation** $\models$ to (possibly infinite) sets of formulas. Let $a : \mathrm{Var} \to A$ be any function, an **assignment** of elements of A to each of the variables. Also, let Φ be a set of formulas over τ. Then Φ holds in $\mathcal{A}$ under the assignment a (or with respect to a) iff for every formula $\varphi(x_1, \ldots, x_n) \in \Phi$ we have

$$\mathcal{A} \models \varphi(a(x_1), \ldots, a(x_n)).$$

In this case we write $\mathcal{A} \models \Phi[a]$ and say that $(\mathcal{A}, a)$ is a **model** of Φ. If Φ holds in $\mathcal{A}$ with respect to every assignment, then we write $\mathcal{A} \models \Phi$ and say that $\mathcal{A}$ is a model of Φ.

Observe that the validity of a set of sentences in $\mathcal{A}$ is independent of the particular assignment. If $i : A \to B$ is an isomorphism between the τ-structures $\mathcal{A}$ and $\mathcal{B}$, every assignment $a : \mathrm{Var} \to A$ induces an assignment $b = i \circ a : \mathrm{Var} \to B$ such that for every set Φ of formulas the following holds:

$$\mathcal{A} \models \Phi[a] \Leftrightarrow \mathcal{B} \models \Phi[b]$$

In particular, isomorphic structures satisfy exactly the same sentences.

The symbol $\models$ is also used for the (semantic) implication between set of formulas. Let Φ and Ψ be sets of formulas over τ. Then Φ **implies** Ψ if for all τ-structures and all assignments a the following holds:

$$\mathcal{A} \models \Phi[a] \Rightarrow \mathcal{A} \models \Psi[a]$$

In this case we write $\Phi \models \Psi$.

We also allow single formulas on either side of $\models$, with the obvious meaning.

1.4. **Completeness.** One of the most important results of Mathematical Logic is the **Completeness Theorem** which states that the relation $\models$ between sets of formulas can be defined in a purely syntactical way.

In order to do this, given a vocabulary τ, one fixes a set of rules and a set of axioms. The axioms are specific formulas over τ such as $(x = y) \to (y = x)$ for all variables x and y and $\varphi \vee \neg\varphi$ for every formula φ over τ. The rules typically are pairs of sets of formulas such as $(\{\varphi, \psi\}, \{(\varphi \wedge \psi)\})$. The rule $(\{\varphi, \psi\}, \{(\varphi \wedge \psi)\})$ allows it to **deduce** the formula $(\varphi \wedge \psi)$ from φ and ψ.

Given a set Φ of formulas over τ and a formula ψ, a **proof** of ψ from Φ is a finite sequence of formulas that ends with ψ and in which every formula is an axiom or an element of Φ or follows from previous formulas in the sequence by application of one of the rules. The formula ψ can be **deduced** from Φ if there is a proof of ψ from Φ. In this case we write $\Phi \vdash \psi$. If Ψ is a set of formulas and for all $\psi \in \Psi$, $\Phi \vdash \psi$, then we write $\Phi \vdash \Psi$.

The rules and axioms are chosen in such a way that the following theorem holds true:

Theorem 1.1 (Gödel's Completeness Theorem)**.** *For every vocabulary* τ*, every set* Φ *of formulas over* τ *and every formula* ψ *over* τ*,*

$$\Phi \vdash \psi \quad \Leftrightarrow \quad \Phi \models \psi.$$

A set Φ of formulas is **consistent** if there is a formula ψ which cannot be deduced from Φ. It turns out Φ is consistent iff no contradiction such as $\psi \wedge \neg\psi$ can be deduced from Φ.

The Completeness Theorem is usually proved in a slightly different form:

Theorem 1.2. *A set* Φ *of formulas over a vocabulary* τ *is consistent iff it has a model.*

Exercise 1.3. Show that Theorem 1.1 implies Theorem 1.2.

Hint: This implication can be proved without showing that Theorem 1.1 or Theorem 1.2 are actually true. Moreover, you have to be a little bit careful since we are considering sets of formulas, not just sets of sentences. A model of a set of formulas consists of both a structure and an assignment.

Exercise 1.4. Show that Theorem 1.2 implies Theorem 1.1.

Hint: The same as in Exercise 1.3.

1.5. **Examples of structures and first order theories.** A **theory** over a vocabulary τ is a set of sentences over τ. Given a structure $\mathcal{A}$, the theory of $\mathcal{A}$ is the set $\mathrm{Th}(\mathcal{A})$ of all sentences φ over τ such that $\mathcal{A} \models \varphi$. If $\mathcal{C}$ is a class of τ-structures, then

$$\mathrm{Th}(\mathcal{C}) = \bigcap_{\mathcal{A} \in \mathcal{C}} \mathrm{Th}(\mathcal{A})$$

is the theory of $\mathcal{C}$. If Φ is a theory, then $\mathrm{Mod}(\Phi)$ is the class of all structures that are models of Φ.

We briefly discuss a few concrete examples of vocabularies, structures and theories.

1.5.1. *Groups.* There are several choices for vocabularies of group theory, namely a single binary function symbol $\cdot$ for the multiplication, a binary function symbol together with a constant symbol e for the identity element or $\cdot$, e and a unary function symbol ${}^{-1}$ for the inversion of group elements.

If $\tau = \{\cdot\}$, we say a **group** is a τ-structure $(G, \cdot)$ satisfying the following sentences:

(G1) $\forall x \forall y \forall z((x \cdot y) \cdot z = x \cdot (y \cdot z))$
(G2) $\exists x \forall y(x \cdot y = y \wedge y \cdot x = y)$
(G3) $\forall x \exists y \forall z(\forall z'(z \cdot z' = z' \wedge z' \cdot z = z') \to (x \cdot y = z \wedge y \cdot x = z))$

Group theory is the theory of the class of all groups, which by the Completeness Theorem is the **deductive closure** of the set of axioms of group theory stated above.

If we use the vocabulary $\{\cdot, e, {}^{-1}\}$, we can state the axioms as follows:

(G1) $\forall x \forall y \forall z((x \cdot y) \cdot z = x \cdot (y \cdot z))$
(G2a) $\forall x(e \cdot x = x \wedge x \cdot e = x)$
(G3a) $\forall x(x \cdot x^{-1} = e \wedge x^{-1} \cdot x = e)$

Observe that these axioms are all of the same form: universal quantifiers followed by an equation (or a conjunction of equations). A formula is **atomic** if it contains no quantifiers or logical connectives ($\neg$, $\vee$).

Exercise 1.5. Let τ be a vocabulary and let Φ be a set of sentences that start with some universal quantifiers followed by an atomic formula. If $\mathcal{A}$ is a model of Φ and $\mathcal{B}$ is a substructure of $\mathcal{A}$, then $\mathcal{B}$ is a model of Φ.

Notice that in a group $(G, \cdot)$ the identity element and the inversion of elements are definable in the following sense: there are formulas $\varphi(x)$ and $\psi(x, y)$ such that for all $a \in G$, $(G, \cdot) \models \varphi(a)$ iff a is the identity element of G and for all $b, c \in G$, $(G, \cdot) \models \psi(b, c)$ iff b is the inverse of c. Namely, let $\varphi(x)$ be the formula $\forall y(x \cdot y = y \wedge y \cdot x = y)$ and let $\psi(x, y)$ be the formula

$$\forall z(\varphi(z) \to (x \cdot y = z \wedge y \cdot x = z)).$$

1.5.2. *The natural numbers.* Now consider the structure $(\mathbb{N}, S, 0)$, where S denotes the function that maps every natural number to its immediate successor. Consider the following **Peano Axioms**:

(PA1) $\forall x \neg(S(x) = 0)$
(PA2) $\forall x \forall y (S(x) = S(y) \to x = y)$
(PA3) Whenever $A \subseteq \mathbb{N}$ contains 0 and is closed under the function S, then $A = \mathbb{N}$.

It is well known that these three axioms determine $(\mathbb{N}, S, 0)$ up to isomorphism. However, there is no way to express (PA3) in first order logic. A reasonable approximation is the following **scheme**: for every formula $\varphi(x, y_1, \ldots, y_n)$ consider

$$\text{(PA3}(\varphi)) \ \forall y_1 \ldots \forall y_n ((\varphi(0, y_1, \ldots, y_n) \wedge \forall x (\varphi(x, y_1, \ldots, y_n) \\ \to \varphi(S(x), y_1, \ldots, y_n))) \to \forall x \varphi(x, y_1, \ldots, y_n)).$$

Here $\varphi(0, y_1, \ldots, y_n)$ and $\varphi(S(x), y_1, \ldots, y_n)$ are used to denote the formulas obtained from φ by replacing every free occurrence of x by 0 or $S(x)$, respectively.

PA is the theory consisting of the axioms (PA1), (PA2) and (PA3(φ)) for all formulas $\varphi(x, y_1, \ldots, y_n)$ over the vocabulary $\{0, S\}$. As we will see later, PA does not determine $(\mathbb{N}, S, 0)$ up to isomorphism.

1.5.3. *Fields.* If for fields we choose the vocabulary $\{+, -, \cdot, ^{-1}, 0, 1\}$, where we consider $-$ as a unary function symbol, we have to be careful with the axioms. Namely, for a field $(F, +, -, \cdot, ^{-1}, 0, 1)$, the function $x \mapsto x^{-1}$ is only a partial function since it is not defined on 0. However, our concept of a structure does not allow for partial functions. So, if we insist on having $^{-1}$ in our vocabulary, we should add a new axiom to the usual field axioms, such as $0^{-1} = 0$, understanding that we always have to exclude 0 when we talk about multiplicative inverses of field elements.

1.5.4. *Vector spaces.* In linear algebra one typically considers at the same time various vector spaces over the same field, rather than vector spaces over different fields. One way to treat vector spaces as structures is to fix a field F and to introduce for each $a \in F$ a unary function symbol f_a that represents the multiplication of elements of the vector space with a. So, let $\tau = \{+, -, 0\} \cup \{f_a : a \in F\}$. In the case of an uncountable field F such as $F = \mathbb{R}$ or $F = \mathbb{C}$, this gives a natural example of an uncountable vocabulary.

Observe that in the case that F is infinite, we need infinitely many axioms to axiomatize the interaction between field elements and elements of the vector space. For example, for all $a, b \in F$ we need the axiom $\forall x (f_a(f_b(x)) = f_{ab}(x))$. This scheme of axioms corresponds to the single axiom

"for all $a, b \in F$ and all $v \in V$, $a \cdot (b \cdot v) = (a \cdot b) \cdot v$"

that is usually used in linear algebra. Here V stands for the vector space under consideration.

Exercise 1.6. Give a complete axiomatization of vector spaces as structures over τ.

1.6. **Compactness.** One of the important properties of first order logic is its **compactness**. If there is a proof of a formula ψ from a set of formulas Φ, then this proof uses only finitely many formulas from Φ. Therefore we have the following theorem:

Theorem 1.7 (Finiteness Theorem)**.** *If Φ is a set of formulas over a vocabulary τ and ψ is a formula over τ, then $\Phi \vdash \tau$ iff there is a finite set $\Phi_0 \subseteq \Phi$ such that $\Phi_0 \vdash \psi$.*

Using the Completeness Theorem, we obtain that $\Phi \models \psi$ iff there is a finite set $\Phi_0 \subseteq \Phi$ such that $\Phi_0 \models \psi$. However, we will give a purely semantic proof of the Finiteness Theorem for the relation $\models$. The Finiteness Theorem for $\models$ follows from

Theorem 1.8 (Compactness Theorem)**.** *Let Φ be a set of formulas over τ. Then Φ has a model iff every finite subset of Φ does.*

In order to prove the Compactness Theorem, we will use **ultraproducts**, which are formed using **ultrafilters**.

Definition 1.9. Let I be a nonempty set. A nonempty collection $\mathcal{F} \subseteq \mathcal{P}(I)$ is a **filter** on I if

(F1) $\emptyset \notin \mathcal{F}$
(F2) If $A \in \mathcal{F}$ and $A \subseteq B \subseteq I$, then $B \in \mathcal{F}$.
(F3) If $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$.

A filter $\mathcal{F}$ is an **ultrafilter** if it is a maximal filter (with respect to set-theoretic inclusion).

A collection $\mathcal{S} \subseteq \mathcal{P}(I)$ has the **finite intersection property** if for all $A_1, \dots, A_n \in \mathcal{S}$, $A_1 \cap \dots \cap A_n \neq \emptyset$.

Lemma 1.10. *Every family $\mathcal{S}$ of subsets of I with the finite intersection property can be extended to an ultrafilter.*

Proof. Let

$$\mathcal{F} = \left\{ A \subseteq I : \text{there is a finite } \mathcal{T} \subseteq \mathcal{S} \text{ such that } \bigcap \mathcal{T} \subseteq A \right\}.$$

$\mathcal{F}$ is the smallest filter that includes $\mathcal{S}$. Consider the partial order of all filters on I that extend $\mathcal{F}$, ordered by set-theoretic inclusion. It is easily checked that the union of every chain of filters is again a filter on I. Hence, by Zorn's Lemma, the partial order has a maximal element, which is an ultrafilter that extends $\mathcal{S}$. □

Lemma 1.11. *Let $\mathcal{F}$ be a filter on a set I. Then the following are equivalent:*

(1) *$\mathcal{F}$ is an ultrafilter.*
(2) *For all $A \subseteq I$, $A \in \mathcal{F}$ iff $I \setminus A \notin \mathcal{F}$.*
(3) *For all $A, B \subseteq I$, if $A \cup B \in \mathcal{F}$, then $A \in \mathcal{F}$ or $B \in \mathcal{F}$.*

Proof. (1)⇒(2): Since $\mathcal{F}$ is a filter, it contains at most one of the sets A and $I \setminus A$. Suppose $I \setminus A \notin \mathcal{F}$. Since $\mathcal{F}$ is closed under taking supersets, this implies that $\mathcal{F}$ contains no subset of $I \setminus A$. In other words, every element of $\mathcal{F}$ intersects A. Since $\mathcal{F}$ is closed under finite intersections, it follows that $\mathcal{F} \cup \{A\}$ has the finite intersection property. Hence there is an ultrafilter $\mathcal{G}$ on I such that $\mathcal{F} \cup \{A\} \subseteq G$. Since $\mathcal{F}$ is a maximal filter, $\mathcal{F} = \mathcal{G}$ and thus $A \in \mathcal{F}$.

(2)⇒(3): Suppose neither A nor B are elements of $\mathcal{F}$. By (2), $I \setminus A, I \setminus B \in \mathcal{F}$ and therefore $I \setminus A \cap I \setminus B = I \setminus (A \cup B) \in \mathcal{F}$. It follows that $A \cup B \notin \mathcal{F}$.

(3)⇒(1): Let $\mathcal{F}$ be a filter satisfying (3). We show that $\mathcal{F}$ is maximal. Let $A \subseteq I$ be such that $\mathcal{F} \cup \{A\}$ is contained in a filter, i.e., has the finite intersection property. Then $I \setminus A \notin \mathcal{F}$. However, $(I \setminus A) \cup A = I \in \mathcal{F}$ since $\mathcal{F}$ is nonempty and closed under taking supersets. By (3), $A \in \mathcal{F}$. This shows the maximality of $\mathcal{F}$. □

Exercise 1.12. Let $\mathcal{F}$ be an ultrafilter on I. Show that if $A_1, \dots, A_n \subseteq I$ and $A_1 \cup \dots \cup A_n \in \mathcal{F}$, then at least one of the sets A_i, $i \in \{1, \dots, n\}$, is an element of $\mathcal{F}$.

We now fix a vocabulary τ.

Definition 1.13. Let I be a set and for each $i \in I$ let $\mathcal{A}_i$ be a τ-structure. Let $\prod_{i \in I} A_i$ be the usual product of the sets A_i, i.e.,

$$\prod_{i \in I} A_i = \{a : a \text{ is a function defined on } I \text{ such that for all } i \in I,\ a(i) \in A_i\}.$$

For every τ-formula $\varphi(x_1, \dots, x_n)$ and for all $a_1, \dots, a_n \in \prod_{i \in I} A_i$ let

$$[\![\varphi(a_1, \dots, a_n)]\!] = \{i \in I : \mathcal{A}_i \models \varphi(a_1(i), \dots, a_n(i))\}.$$

If $\mathcal{U}$ is an ultrafilter on I and $a, b \in \prod_{i \in I} A_i$, we let $a \sim_{\mathcal{U}} b$ iff

$$[\![a = b]\!] = \{i \in I : a(i) = b(i)\} \in \mathcal{U}.$$

Lemma 1.14. *Let $(\mathcal{A}_i)_{i \in I}$ be a family of τ-structures and let $\mathcal{U}$ be an ultrafilter on I.*

a) $\sim_{\mathcal{U}}$ is an equivalence relation on $\prod_{i \in I} A_i$.

b) If $f \in \tau$ is an n-ary function symbol and $a_j \sim_{\mathcal{U}} b_j$ for every $j \in \{1, \dots, n\}$, then $[\![f(a_1, \dots, a_n) = f(b_1, \dots, b_n)]\!] \in \mathcal{U}$.

c) For every n-ary relation symbol $R \in \tau$, if $a_j \sim_{\mathcal{U}} b_j$ for every $j \in \{1, \dots, n\}$, then

$$[\![R(a_1, \dots, a_n)]\!] \in \mathcal{U} \quad \Leftrightarrow \quad [\![R(b_1, \dots, b_n)]\!] \in \mathcal{U}.$$

Proof. a) It is clear that $\sim_{\mathcal{U}}$ is symmetric and reflexive. In order to show transitivity, let $a \sim_{\mathcal{U}} b$ and $b \sim_{\mathcal{U}} c$. Now $[\![a = b]\!], [\![b = c]\!] \in \mathcal{U}$. Since $\mathcal{U}$ is closed under finite intersections, $[\![a = b]\!] \cap [\![b = c]\!] \in \mathcal{U}$. Now $[\![a = b]\!] \cap [\![b = c]\!] \subseteq [\![a = c]\!]$ and thus $[\![a = c]\!] \in \mathcal{U}$. It follows that $a \sim_{\mathcal{U}} c$.

b) This is similar to the proof of a). If $a_j \sim_{\mathcal{U}} b_j$ for every $j \in \{1, \dots, n\}$, then $[\![a_1 = b_1]\!] \cap \dots \cap [\![a_n = b_n]\!] \in \mathcal{U}$. Now for all

$$i \in [\![a_1 = b_1]\!] \cap \dots \cap [\![a_n = b_n]\!]$$

we have

$$f(a_1(i), \dots, a_n(i)) = f(b_1(i), \dots, b_n(i)).$$

It follows that $[\![f(a_1, \dots, a_n) = f(b_1, \dots, b_n)]\!] \in \mathcal{U}$ and thus

$$f(a_1, \dots, a_n) \sim_{\mathcal{U}} f(b_1, \dots, b_n).$$

c) This is similar to the proof of b). If $a_j \sim_{\mathcal{U}} b_j$ for every $j \in \{1, \dots, n\}$, then $[\![a_1 = b_1]\!] \cap \dots \cap [\![a_n = b_n]\!] \in \mathcal{U}$. If $i \in [\![a_1 = b_1]\!] \cap \dots \cap [\![a_n = b_n]\!]$, then $(a_1(i), \dots, a_n(i)) \in R^{\mathcal{A}_i}$ iff $(b_1(i), \dots, b_n(i)) \in R^{\mathcal{A}_i}$. Now, if $[\![R(a_1, \dots, a_n)]\!] \in \mathcal{U}$, then

$$J = [\![R(a_1, \dots, a_n)]\!] \cap [\![a_1 = b_1]\!] \cap \dots \cap [\![a_n = b_n]\!] \in \mathcal{U}.$$

For every $i \in J$ we have $i \in [\![R(b_1, \dots, b_n)]\!]$ and thus $[\![R(b_1, \dots, b_n)]\!] \in \mathcal{U}$. The other direction of c) is symmetric. □

Definition 1.15. Let $(\mathcal{A}_i)_{i \in I}$ and $\mathcal{U}$ be as in Lemma 1.14. The underlying set of the strucure $\mathcal{A} = \prod_{i \in I} \mathcal{A}_i / \mathcal{U}$ is the set $A = \prod_{i \in I} A_i / \sim_{\mathcal{U}}$ of all $\sim_{\mathcal{U}}$-equivalence classes. For every constant symbol $c \in \tau$ we let $c^{\mathcal{A}}$ be the $\sim_{\mathcal{U}}$-equivalence class of the function that assigns to every $i \in I$ the value $c^{\mathcal{A}_i} \in A_i$.

Given $a \in \prod_{i \in I} A_i$, by $[a]_{\mathcal{U}}$ we denote the $\sim_{\mathcal{U}}$-equivalence class of a. For every n-ary function symbol $f \in \tau$ and $a_1, \dots, a_n \in \prod_{i \in I} A_i$ let $f^{\mathcal{A}}([a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}})$ be the $\sim_{\mathcal{U}}$-equivalence class of the function that assigns to each $i \in I$ the value $f^{\mathcal{A}_i}(a_1(i), \dots, a_n(i))$.

For every n-ary relation symbol $R \in \tau$ and all $a_1, \dots, a_n \in \prod_{i \in I} A_i$ let $([a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}}) \in R^{\mathcal{A}}$ iff $[\![R(a_1, \dots, a_n)]\!] \in \mathcal{U}$. Observe that the structure $\mathcal{A}$ is well-defined by Lemma 1.14. $\mathcal{A}$ is the **ultraproduct** of the structures $\mathcal{A}_i$, $i \in I$, with respect to the ultrafilter $\mathcal{U}$.

Exercise 1.16. Let $(\mathcal{A}_i)_{i \in I}$ be a family of τ-structures. Let $i_0 \in I$ and let $\mathcal{U}$ be the ultrafilter on I that consists of all subsets of I that contain i_0. Show that $\prod_{i \in I} \mathcal{A}_i / \mathcal{U}$ is isomorphic to $\mathcal{A}_{i_0}$.

Theorem 1.17 (Łoś). *Let $\mathcal{A}$ be the ultraproduct of the structures $\mathcal{A}_i$, $i \in I$, with respect to the ultrafilter $\mathcal{U}$ on I. Then for every formula $\varphi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in \prod_{i \in I} A_i$,*

$$\mathcal{A} \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}}) \quad \Leftrightarrow \quad [\![\varphi(a_1, \dots, a_n)]\!] \in \mathcal{U}.$$

Proof. We prove the theorem by induction on the complexity of formulas. For atomic formulas, Łoś's Theorem follows immediately from the definition of the ultraproduct.

Now let $\varphi(x_1,\dots,x_n)$ and $\psi(x_1,\dots,x_n)$ be formulas and $a_1,\dots,a_n \in \prod_{i\in I} A_i$. Assume that for $\chi=\varphi$ and $\chi=\psi$ it holds that

$$[\![\chi(a_1,\dots,a_n)]\!] \in \mathcal{U} \text{ iff } \mathcal{A} \models \chi([a_1]_{\mathcal{U}},\dots,[a_n]_{\mathcal{U}}).$$

We have

$$[\![(\varphi\vee\psi)(a_1,\dots,a_n)]\!] = [\![\varphi(a_1,\dots,a_n)]\!] \cup [\![\psi(a_1,\dots,a_n]\!].$$

Now

$$[\![\varphi(a_1,\dots,a_n)]\!] \cup [\![\psi(a_1,\dots,a_n]\!] \in \mathcal{U}$$

iff $[\![\varphi(a_1,\dots,a_n)]\!] \in \mathcal{U}$ or $[\![\psi(a_1,\dots,a_n)]\!] \in \mathcal{U}$. Hence, by our assumption on φ and ψ,

$$[\![(\varphi\vee\psi)(a_1,\dots,a_n)]\!] \in \mathcal{U} \text{ iff } \mathcal{A} \models (\varphi\vee\psi)([a_1]_{\mathcal{U}},\dots,[a_n]_{\mathcal{U}}).$$

Also,

$$[\![\neg\varphi(a_1,\dots,a_n)]\!] = I \setminus [\![\varphi(a_1,\dots,a_n)]\!].$$

Since $\mathcal{U}$ is an ultrafilter,

$$I \setminus [\![\varphi(a_1,\dots,a_n)]\!] \in \mathcal{U} \text{ iff } [\![\varphi(a_1,\dots,a_n)]\!] \notin \mathcal{U}.$$

This shows that

$$\mathcal{A} \models \neg\varphi([a_1]_{\mathcal{U}},\dots,[a_n]_{\mathcal{U}}) \text{ iff } [\![\neg\varphi(a_1,\dots,a_n)]\!] \in \mathcal{U}.$$

Finally, consider the formula $\exists x\varphi(x,y_1,\dots,y_n)$, let $b_1,\dots,b_n \in \prod_{i\in I} A_i$ and assume that for all $a \in \prod_{i\in I} A_i$,

$$[\![\varphi(a,b_1,\dots,b_n)]\!] \in \mathcal{U} \text{ iff } \mathcal{A} \models \varphi([a]_{\mathcal{U}},[b_1]_{\mathcal{U}},\dots,[b_n]_{\mathcal{U}}).$$

If $[\![(\exists x\varphi)(b_1,\dots,b_n)]\!] \in \mathcal{U}$, for each $i \in [\![(\exists x\varphi)(b_1,\dots,b_n)]\!]$ we choose $a(i) \in A_i$ such that $\mathcal{A}_i \models \varphi(a(i),b_1(i),\dots,b_n(i))$. For all $i \in I \setminus [\![(\exists x\varphi)(b_1,\dots,b_n)]\!]$ choose $a(i) \in A_i$ arbitrarily. Now

$$[\![(\exists x\varphi)(b_1,\dots,b_n)]\!] \subseteq [\![\varphi(a,b_1,\dots,b_n)]\!]$$

and therefore $[\![\varphi(a,b_1,\dots,b_n)]\!] \in \mathcal{U}$. By our assumption on φ,

$$\mathcal{A} \models \varphi([a]_{\mathcal{U}},[b_1]_{\mathcal{U}},\dots,[b_n]_{\mathcal{U}}).$$

It follows that

$$\mathcal{A} \models (\exists x\varphi)([b_1]_{\mathcal{U}},\dots,[b_n]_{\mathcal{U}}).$$

If

$$\mathcal{A} \models (\exists x\varphi)([b_1]_{\mathcal{U}},\dots,[b_n]_{\mathcal{U}}),$$

there is some $a \in \prod_{i\in I} A_i$ such that

$$\mathcal{A} \models \varphi([a]_{\mathcal{U}},[b_1]_{\mathcal{U}},\dots,[b_n]_{\mathcal{U}}).$$

By our assumption on φ, $[\![\varphi(a,b_1,\dots,b_n)]\!] \in \mathcal{U}$. It follows that

$$[\![(\exists x\varphi)(b_1,\dots,b_n)]\!] \in \mathcal{U},$$

finishing the proof of Łoś's theorem. □

Definition 1.18. An ultrafilter $\mathcal{U}$ on a set I is **free** if it does not contain any finite sets.

Exercise 1.19. Let $\mathcal{U}$ be an ultrafilter on an infinite set I. Show that $\mathcal{U}$ contains a finite set iff $\mathcal{U}$ is of the form $\{S \subseteq I : i \in S\}$ for some $i \in I$.

We have now gathered the necessary tools in order prove the Compactness Theorem.

Proof of Theorem 1.8. Clearly, if Φ has a model, then so does every finite subset of Φ.

Now assume that every finite subset of Φ has a model. Let I denote the collection of all finite subsets of Φ. A set $S \subseteq I$ is **upward closed** if for every $\Phi_0 \in S$ and every $\Phi_1 \in I$ with $\Phi_0 \subseteq \Phi_1$ we have $\Phi_1 \in S$. If $S, T \subseteq I$ are nonempty and upward closed, then so is $S \cap T$. Namely, choose $\Phi_S \in S$ and $\Phi_T \in T$. Then $\Phi_S \cup \Phi_T$ is finite and, by the upward closedness of S and T, an element of both S and T. Hence $S \cap T \neq \emptyset$. Also, if $\Phi_0 \in S \cap T$ and $\Phi_0 \subseteq \Phi_1 \in I$, then Φ_1 is both in S and in T and thus in $S \cap T$.

This implies that the collection of all nonempty upward closed $S \subseteq I$ has the finite intersection property. By Lemma 1.10, the family of all nonempty upward closed subsets of I extends to an ultrafilter $\mathcal{U}$ on I. For every $\Phi_0 \in I$ choose a τ-structure $\mathcal{A}_{\Phi_0}$ that is a model of Φ_0. Let

$$\mathcal{A} = \prod_{\Phi_0 \in I} \mathcal{A}_{\Phi_0} / \mathcal{U}.$$

Claim 1.20. $\mathcal{A} \models \Phi$

Let $\varphi \in \Phi$. Let $S = \{\Phi_0 \in I : \varphi \in \Phi_0\}$. Clearly, S is upward closed and nonempty. It follows that $S \in \mathcal{U}$. For every $\Phi_0 \in S$, $\mathcal{A}_{\Phi_0} \models \varphi$. Hence, by Łoś's theorem, $\mathcal{A} \models \varphi$. This shows that $\mathcal{A}$ is a model of all of Φ. $\square$

Corollary 1.21 (Finiteness Theorem for $\models$)**.** *For every set Φ of sentences over τ and every sentence ψ over τ, $\Phi \models \psi$ iff there is a finite set $\Phi_0 \subseteq \Phi$ such that $\Phi_0 \models \psi$.*

Proof. Clearly, if $\Phi_0 \models \psi$ for some subset of Φ, then $\Phi \models \psi$. Now assume that for every finite $\Phi_0 \subseteq \Phi$, $\Phi_0 \not\models \psi$. Then for every finite $\Phi_0 \subseteq \Phi$, $\Psi_0 \cup \{\neg\psi\}$ has a model. It follows that every finite subset of $\Phi \cup \{\neg\psi\}$ has a model. By Theorem 1.8, $\Phi \cup \{\neg\psi\}$ has a model. This shows that $\Phi \not\models \psi$. $\square$

A special case of ultraproducts are **ultrapowers** where each factor is the same.

Exercise 1.22. Let I be a set, $\mathcal{U}$ an ultrafilter on I, and $\mathcal{A}$ a τ-structure. Let $\mathcal{B} = \prod_{i \in I} \mathcal{A}/\mathcal{U}$. Let B be the underlying set of $\mathcal{B}$. For each $a \in A$ let $\overline{a} : I \to A$ be the map that assigns to all $i \in I$ the value a. Consider the map

$$j : A \to B; a \mapsto [\overline{a}]_{\mathcal{U}}.$$

Show that $j[A]$ carries a substructure of $\mathcal{B}$ and that j is a isomorphism between $\mathcal{A}$ and the substructure of $\mathcal{B}$ that lives on $j[A]$.

We will now see our first example of a model of the first order Peano Axioms that is not isomorphic to $\mathbb{N}$.

Exercise 1.23. Let $\mathcal{A} = (\mathbb{N}, 0, S)$. Let $I = \mathbb{N}$ and let $\mathcal{U}$ be a free ultrafilter on I. Consider the ultrapower $\mathcal{B} = \prod_{i \in I} \mathcal{A}/\mathcal{U}$ and define $j : A \to B$ as in Exercise 1.22. Show that j is not onto.

We give another application of ultraproducts.

Theorem 1.24. *Let Φ be a theory over a vocabulary τ. If Φ has arbitrarily large finite models, then it has an infinite model.*

Proof. For every $n \in \mathbb{N}$ let $\mathcal{A}_n$ be a model of Φ of size at least n. Let $I = \mathbb{N}$ and consider the **Fréchet filter**

$$\mathcal{F} = \{S \subseteq \mathbb{N} : \mathbb{N} \setminus S \text{ is finite }\}.$$

Let $\mathcal{U}$ be an ultrafilter that extends $\mathcal{F}$. It is easily checked that an ultrafilter $\mathcal{U}$ is free iff it extends $\mathcal{F}$. Let $\mathcal{A} = \prod_{n \in \mathbb{N}} \mathcal{A}_n$.

For all $n \in \mathbb{N}$ and all $m \geq n$, $\mathcal{A}_m$ satisfies the sentence

$$\varphi_n = \exists x_1 \ldots \exists x_n (x_1 \neq x_2 \wedge \cdots \wedge x_1 \neq x_n \wedge x_2 \neq x_3 \wedge \cdots \wedge x_{n-1} \neq x_n).$$

By the choice of $\mathcal{U}$, for all $n \in \mathbb{N}$ the set $\{m \in \mathbb{N} : m > n\}$ is an element of $\mathcal{U}$. Hence, by Łoś's theorem, for all $n \in \mathbb{N}$, $\mathcal{A} \models \varphi_n$. It follows that $\mathcal{A}$ is an infinite structure. □

This theorem can also easily be proved using the Compactness Theorem. Namely, if Φ has arbitrarily large finite models, then for all $n \in \mathbb{N}$, $\Phi \cup \{\varphi_n\}$ has a model, where φ_n is defined as in the proof of the previous theorem. Since φ_m implies φ_n if $m \geq n$, we have that every finite subset of $\Phi \cup \{\varphi_n : n \in \mathbb{N}\}$ has a model. Hence $\Phi \cup \{\varphi_n : n \in \mathbb{N}\}$ has a model. But a model of $\Phi \cup \{\varphi_n : n \in \mathbb{N}\}$ is a model of Φ and has to be infinite.

Another application of ultraproducts is to prove the finite Ramsey Theorem from the infinite Ramsey Theorem. We only deal with the Ramsey Theorem for graphs. Let V be a set and let $G = (V, E)$ be a graph on V, i.e., let E be a binary relation on V that is both irreflexive ($\forall x \neg E(x, x)$) and symmetric ($\forall x \forall y (E(x, y) \leftrightarrow E(y, x))$). A set $H \subseteq V$ is **homogeneous** if either any two elements of H are related or no two elements of H are related.

Theorem 1.25 (Infinite Ramsey Theorem for graphs)**.** *Every infinite graph has an infinite homogeneous subset.*

Proof. Let $G = (V, E)$ be an infinite graph. Without loss of generality we may assume that $V = \mathbb{N}$. We define a strictly increasing sequence $(a_n)_{n\in\mathbb{N}}$ of natural numbers, a sequence $(i_n)_{n\in\mathbb{N}}$ of 0's and 1's, and a sequence $(A_n)_{n\in\mathbb{N}}$ of subsets of $\mathbb{N}$ such that $A_0 \supseteq A_1 \supseteq \dots$ and such that for all $n \in \mathbb{N}$ and all $m \geq n$, $a_m \in A_n$. Let $A_0 = \mathbb{N}$ and let $a_0 = 0$. Suppose a_n and A_n have been defined. One of the two sets $A_n^1 = \{a \in A_n : a > a_n \wedge E(a_n, a)\}$ and $A_n^0 = \{a \in A_n : a > a_n \wedge \neg E(a_n, a)\}$ is infinite.

Let $i_n \in 2$ be such that $A_n^{i_n}$ is infinite. Let $A_{n+1} = A_n^{i_n}$ and let $a_{n+1} = \min A_{n+1}$. This finishes the recursive definition of the three sequences. Now for all $n, m \in \mathbb{N}$ with $n < m$, whether or not $E(a_n, a_m)$ holds only depends on n. Namely, $E(a_n, a_m)$ holds iff $i_n = 1$.

Let $i \in 2$ be such that the set $\{n \in \mathbb{N} : i_n = i\}$ is infinite. Now $H = \{a_n : i_n = i\}$ is an infinite homogeneous subset of V. □

The finite Ramsey Theorem says that for every $n \in \mathbb{N}$ there is $m \in \mathbb{N}$ such that every graph of size at least m has a homogeneous subset of size at least n.

Theorem 1.26 (Finite Ramsey Theorem for graphs)**.** *For every $n \in \mathbb{N}$ there is $m \in \mathbb{N}$ such that every graph of size at least m has a homogeneous subset of size at least n.*

Proof. Suppose the finite Ramsey theorem fails. Then there is $n \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ there is a finite graph G_m of size at least m without a homogeneous set of size n. Let $\mathcal{U}$ be a free ultrafilter on $\mathbb{N}$. Let $G = \prod_{m\in\mathbb{N}} G_m/\mathcal{U}$. Now G is an infinite graph that satisfies the statement

$$\neg\exists x_1, \dots, x_n(\text{“the } x_i \text{ are pairwise distinct and form a homogeneous set”}),$$

which can be easily expressed in first order logic. This shows that G is a counter example to the infinite Ramsey Theorem. □

Exercise 1.27. Use the Compactness Theorem to deduce the finite Ramsey Theorem from the infinite.

We conclude this section on compactness with an observation about fields of positive characteristic. Recall that a field F is of characteristic 0 iff for all $n > 0$ the sum $1 + \dots + 1$ with n summands is different from 0. F is of characteristic $n > 0$ if n is the minimal integer > 0 such that the sum $1 + \dots + 1$ with n summands is equal to 0. It turns out that the characteristic of a field is either 0 or a prime number.

Theorem 1.28. *Let $\tau = \{+, \cdot, 0, 1\}$. If Φ is a theory over τ that is satisfied by fields of arbitrarily large characteristic, then Φ is satisfied by a field of characteristic 0.*

Proof. For every $n > 0$ let ψ_n be the sentence over τ that says that the sum $1 + \cdots + 1$ with n summands is different from 0. We may assume that Φ contains the axioms of field theory.

Consider the theory $\Phi \cup \{\psi_n : n \in \omega\}$. If $\Phi_0 \subseteq \Phi \cup \{\psi_n : n \in \omega\}$ is finite, then there is $m > 0$ such that $\Phi_0 \subseteq \Phi \cup \{\psi_n : n < m\}$. But since Φ is satisfied by fields of arbitrarily large characteristic, the theory $\Phi \cup \{\psi_n : n < m\}$ has a model. It follows that Φ_0 has a model.

By the Compactness Theorem, $\Phi \cup \{\psi_n : n \in \mathbb{N}\}$ has a model, which is a field of characteristic 0. □

This theorem in particular shows that fields of positive characteristic cannot be axiomatized in first order logic.

Exercise 1.29. Show that if a sentence φ holds for every field of characteristic 0, then there is $p > 0$ such that φ holds in every field of characteristic at least p.

1.7. **Elementary substructures.**

Definition 1.30. Let $\mathcal{B}$ be a τ-structure and let $\mathcal{A}$ be a substructure of $\mathcal{B}$. A τ-formula $\varphi(x_1, \ldots, x_n)$ is **absolute** between $\mathcal{A}$ and $\mathcal{B}$ iff for all $a_1, \ldots, a_n \in A$,

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

$\mathcal{A}$ is an **elementary substructure** (or an **elementary submodel**) of $\mathcal{B}$ ($\mathcal{A} \preccurlyeq \mathcal{B}$) if every τ-formula is absolute between $\mathcal{A}$ and $\mathcal{B}$.

Two τ-structures are **elementary equivalent** if they satisfy the same sentences over τ. We write $\mathcal{A} \equiv \mathcal{B}$ if two structures $\mathcal{A}$ and $\mathcal{B}$ are elementary equivalent.

Lemma 1.31 (Tarski-Vaught Criterion)**.** *Let $\mathcal{A}$ be a substructure of $\mathcal{B}$. Then $\mathcal{A} \preccurlyeq \mathcal{B}$ iff for every τ-formula $\varphi(x, y_1, \ldots, y_n)$ and all $a_1, \ldots, a_n \in A$,*

(1) *if there is $a \in B$ such that $\mathcal{B} \models \varphi(a, a_1, \ldots, a_n)$, then*
there is $a \in A$ such that $\mathcal{B} \models \varphi(a, a_1, \ldots, a_n)$.

Proof. Clearly, if $\mathcal{A} \preccurlyeq \mathcal{B}$, then for all formulas $\varphi(x, y_1, \ldots, y_n)$ and all $a_1, \ldots, a_n \in A$ we have (1).

Now suppose for all formulas $\varphi(x, y_1, \ldots, y_n)$ and all $a_1, \ldots, a_n \in A$ we have (1). Since $\mathcal{A}$ is a substructure of $\mathcal{B}$, every atomic formula is absolute between $\mathcal{A}$ and $\mathcal{B}$. Also, it is easily checked that the class of absolute formulas is closed under negation and disjunction.

Now assume that $\varphi(x, y_1, \ldots, y_n)$ is absolute between $\mathcal{A}$ and $\mathcal{B}$ and let $a_1, \ldots, a_n \in A$. Suppose that $\mathcal{A} \models (\exists x\varphi)(a_1, \ldots, a_n)$. Let $a \in A$ be such that $\mathcal{A} \models \varphi(a, a_1, \ldots, a_n)$. By the absoluteness of φ, $\mathcal{B} \models \varphi(a, a_1, \ldots, a_n)$ and thus $\mathcal{B} \models (\exists x\varphi)(a_1, \ldots, a_n)$.

On the other hand, if $\mathcal{B} \models (\exists x\varphi)(a_1, \ldots, a_n)$, then by (1), there is $a \in A$ such that $\mathcal{B} \models \varphi(a, a_1, \ldots, a_n)$. By the absoluteness of φ,

$\mathcal{A} \models \varphi(a, a_1, \ldots, a_n)$. It follows that $\mathcal{A} \models (\exists x\varphi)(a_1, \ldots, a_n)$, showing that $\exists x\varphi$ is absolute between $\mathcal{A}$ and $\mathcal{B}$. □

Theorem 1.32 (Downward Löwenheim-Skolem Theorem)**.** *Let κ be an infinite cardinal and let τ be a vocabulary of size at most κ. Then for every structure $\mathcal{B}$ and every set $X \subseteq B$ of size at most κ there is an elementary substructure $\mathcal{A}$ of $\mathcal{B}$ such that $X \subseteq A$ and $|A| \leq \kappa$.*

Proof. Let $X \subseteq B$ be a set of size at most κ. For each formula $\varphi(x, y_1, \ldots, y_n)$ we define a function $f_\varphi : B^n \to B$ as follows: for all $b_1, \ldots, b_n \in B$ with $\mathcal{B} \models (\exists x\varphi)(b_1, \ldots, b_n)$ choose an element $a \in B$ such that $\mathcal{B} \models \varphi(a, b_1, \ldots, b_n)$ and let $f_\varphi(b_1, \ldots, b_n) = a$. If $\mathcal{B} \not\models (\exists x\varphi)(b_1, \ldots, b_n)$, then let $f_\varphi(b_1, \ldots, b_n)$ be an arbitrary element of B. The functions f_φ are **Skolem functions**.

We now define a sequence $(X_n)_{n\in\mathbb{N}}$ of subsets of B. Let X_0 be a nonempty subset of B of size at most κ such that $X \subseteq X_0$. If X_m has been defined, let

$$X_{m+1} = X_m \cup \{f_\varphi(b_1, \ldots, b_n) : b_1, \ldots, b_n \in X_n \text{ and } \varphi(x, y_1, \ldots, y_n) \text{ is a } \tau\text{-formula}\}.$$

Let $A = \bigcup_{m\in\mathbb{N}} X_m$.

Now, whenever $\varphi(x, y_1, \ldots, y_n)$ is a τ-formula and $a_1, \ldots, a_n \in A$, then for some $m \in \mathbb{N}$, $a_1, \ldots, a_n \in X_m$ and therefore

$$f_\varphi(a_1, \ldots, a_n) \in X_{m+1} \subseteq A.$$

It follows that A is closed under all the Skolem functions.

Clearly, for every n-ary function symbol $f \in \tau$ and all $b_1, \ldots, b_n \in B$, $\mathcal{B} \models \exists x(x = f(b_1, \ldots, b_n))$. Let $\varphi(x, y_1, \ldots, y_n)$ be the formula $x = f(y_1, \ldots, y_n)$. Then $f_\varphi = f^{\mathcal{B}}$. Since A is closed under all the Skolem functions, for every n-ary function symbol $f \in \tau$, A is closed under $f^{\mathcal{B}}$. Also, for every constant symbol $c \in \tau$ and every $a \in A$, $f_{c=x}(a) = c^{\mathcal{B}}$ and thus $c^{\mathcal{B}} \in A$. This shows that A is the underlying set of a substructure $\mathcal{A}$ of $\mathcal{B}$.

We now apply the Tarski-Vaught Criterion to show that $\mathcal{A} \preccurlyeq \mathcal{B}$. Let $\varphi(x, y_1, \ldots, y_n)$ be a τ-formula and let $a_1, \ldots, a_n \in A$ be such that $\mathcal{B} \models (\exists x\varphi)(a_1, \ldots, a_n)$. Let $a = f_\varphi(a_1, \ldots, a_n)$. By the choice of A, $a \in A$. By the choice of f_φ, $\mathcal{B} \models \varphi(a, a_1, \ldots, a_n)$. Hence by the Tarski-Vaught Criterion, $\mathcal{A} \preccurlyeq \mathcal{B}$. □

Exercise 1.33. $\mathcal{Q} = (\mathbb{Q}, +, \cdot, 0, 1)$ is a substructure of $\mathcal{R} = (\mathbb{R}, +, \cdot, 0, 1)$. Show that $\mathcal{Q}$ is not an elementary substructure of $\mathcal{R}$.

Observe that by the Löwenheim-Skolem Theorem, $\mathcal{R}$ has a countable elementary submodel that contains all rational numbers. We will see later that $(\mathbb{Q}, <)$ is an elementary submodel of $(\mathbb{R}, <)$

Theorem 1.34 (Upward Löwenheim-Skolem Theorem, weak version)**.** *Let Φ be a theory that has an infinite model. Then Φ has arbitrarily* *large models. In particular, for every infinite τ-structure $\mathcal{A}$ there are arbitrarily large τ-structures that are elementary equivalent to $\mathcal{A}$.*

Proof. Let Φ be a theory that has an infinite model. Let I be any set. For each $i \in I$ we introduce a new constant symbol c_i. Let σ be the vocabulary obtained by adding the constant symbol c_i, $i \in I$, to τ. Let

$$\Phi_I = \Phi \cup \{c_i \neq c_j : i, j \in I \text{ and } i \neq j\}.$$

For every finite subset Φ_0 of Φ_I there is a finite set $I_0 \subseteq I$ such that $\Phi_0 \subseteq \Phi \cup \{c_i \neq c_j : i, j \in I_0 \text{ and } i \neq j\}$. Let $\mathcal{A}$ be an infinite model of Φ. If $I_0 \subseteq I$ is finite, choose pairwise distinct interpretations $c_i^{\mathcal{A}'} \in A$ of the c_i, $i \in I_0$. This is possible since $\mathcal{A}$ is infinite. Now let $\mathcal{A}'$ be the σ-structure with underlying set is A, in which all the constant symbols, function symbols and relation symbols of τ are interpreted as in $\mathcal{A}$, in which the constant symbols c_i, $i \in I_0$, are interpreted by the respective $c_i^{\mathcal{A}'}$, and in which the constant symbols c_i, $i \in I \setminus I_0$, are interpreted arbitrarily.

By the choice of the $c_i^{\mathcal{A}'}$, $i \in I_0$,

$$\mathcal{A}' \models \Phi \cup \{c_i \neq c_j : i, j \in I_0 \text{ and } i \neq j\}.$$

It follows that every finite subset of Φ_i has a model. Hence, by the Compactness Theorem, Φ_I has a model $\mathcal{B}$. Since the c_i, $i \in I$, have pairwise distinct interpretations in $\mathcal{B}$, the underlying set B is at least of size $|I|$. Since I was arbitrary, this shows that Φ has arbitrarily large models. □

Notice that for the proof of the Upward Löwenheim-Skolem Theorem it is actually enough to assume that Φ has an infinite model or arbitrarily large finite models. However, we have already seen that a theory with arbitrarily large finite models has an infinite model.

It follows from the Upward Löwenheim-Skolem Theorem that there are arbitrarily large models of the theory of $(\mathbb{N}, 0, S)$ and of the theory of $(\mathbb{N}, +, \cdot, 0, 1)$. This shows that the natural numbers cannot be axiomatized in first order logic up to isomorphism.

On the other hand, if the usual axioms for set theory (ZFC) are consistent, and we have no reason to believe otherwise, then, by the Completeness Theorem together with the Löwenheim-Skolem Theorems, there is a countable model of set theory. In this model there are only countably elements that the model considers to be real numbers, even though the model itself believes that there are uncountably many real numbers. But this just means that the notion "countable" of the model is different from the corresponding notion in the real world, for instance because the model does not know the bijection between the countably many real numbers of the model and the natural numbers of the model that exists in the real world.

Definition 1.35. Let $\mathcal{A}$ and $\mathcal{B}$ be τ-structures. Then $j : A \to B$ is an **elementary embedding** of $\mathcal{A}$ into $\mathcal{B}$ if j is an isomorphism from $\mathcal{A}$ onto an elementary substructure of $\mathcal{B}$. In other words, j is an elementary embedding iff for all formulas $\varphi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in A$ we have

$$\mathcal{A} \models \varphi(a_1, \dots, a_n) \quad \Leftrightarrow \quad \mathcal{B} \models \varphi(j(a_1), \dots, j(a_n)).$$

Elementary embeddings play a major role in some areas of set theory and in nonstandard analysis.

Exercise 1.36. Let $\mathcal{A}$ be a τ-structure, I a set and $\mathcal{U}$ an ultrafilter on I. Show that the embedding j of $\mathcal{A}$ into the ultrapower $\mathcal{A}^I/\mathcal{U}$ as defined in Exercise 1.22 is elementary.

Definition 1.37. Let σ and τ be vocabularies such that $\sigma \subseteq \tau$. If $\mathcal{A}$ is a σ-structure and $\mathcal{A}'$ is a τ-structure on the same underlying set A such that every symbol in σ has the same interpretation in $\mathcal{A}$ as in $\mathcal{A}'$, then $\mathcal{A}$ is the **reduct** of $\mathcal{A}'$ to σ and $\mathcal{A}'$ is an **expansion** of $\mathcal{A}$ to τ. We write $\mathcal{A} = \mathcal{A}' \restriction \sigma$.

Now let $\mathcal{A}$ be a τ-structure. For each $a \in A$ let c_a be a new constant symbol. Let $\tau_A = \tau \cup \{c_a : a \in A\}$. The **diagram** of $\mathcal{A}$ is the theory

$$\begin{aligned}\operatorname{diag}(\mathcal{A}) = \{\varphi(c_{a_1}, \dots, c_{a_n}) &: a_1, \dots, a_n \in A, \\ &\varphi(x_1, \dots, x_n) \text{ is an atomic formula or the negation} \\ &\qquad \text{of an atomic formula and } \mathcal{A} \models \varphi(a_1, \dots, a_n)\}.\end{aligned}$$

The **elementary diagram** of $\mathcal{A}$ is the theory

$$\begin{aligned}\operatorname{eldiag}(\mathcal{A}) = \{\varphi(c_{a_1}, \dots, c_{a_n}) &: a_1, \dots, a_n \in A, \\ &\varphi(x_1, \dots, x_n) \text{ is a } \tau\text{-formula and } \mathcal{A} \models \varphi(a_1, \dots, a_n)\}.\end{aligned}$$

Observe that $\operatorname{eldiag}(\mathcal{A})$ is a **complete theory** in the sense that for every τ_A-sentence φ either $\varphi \in \operatorname{eldiag}(\mathcal{A})$ or $\neg\varphi \in \operatorname{eldiag}(\mathcal{A})$.

Exercise 1.38. Let $\mathcal{A}$ be a τ-structure. Show that for every τ_A-structure $\mathcal{B}$ with $\mathcal{B} \models \operatorname{diag}(\mathcal{A})$, the reduct $\mathcal{B} \restriction \tau$ has a substructure that is isomorphic to $\mathcal{A}$.

Lemma 1.39. *Let $\mathcal{A}$ be a τ-structure. If $\mathcal{B}$ is a model of* $\operatorname{eldiag}(\mathcal{A})$*, then the map $e : A \to B; a \mapsto c_a^{\mathcal{B}}$ is an elementary embedding of $\mathcal{A}$ into $\mathcal{B} \restriction \tau$.*

Proof. Let $\varphi(x_1, \dots, x_n)$ be a τ-formula and $a_1, \dots, a_n \in A$. Then $\mathcal{A} \models \varphi(a_1, \dots, a_n)$ iff $\varphi(c_{a_1}, \dots, c_{a_n}) \in \operatorname{eldiag}(\mathcal{A})$ iff $\mathcal{B} \models \varphi(c_{a_1}, \dots, c_{a_n})$ iff

$$\mathcal{B} \restriction \tau \models \varphi(e(a_1), \dots, e(a_n)).$$

This shows that e is an elementary embedding. $\square$

Corollary 1.40 (Upward Löwenheim-Skolem Theorem, strong version)**.** *Let $\mathcal{A}$ be an infinite τ-structure. Then there are arbitrarily large τ-structures $\mathcal{B}$ such that $\mathcal{A}$ elementarily embeds into $\mathcal{B}$.*

Proof. Interpreting every constant c_a, $a \in A$, by a itself, we obtain an expansion of $\mathcal{A}$ that is a model of eldiag($\mathcal{A}$). Hence eldiag($\mathcal{A}$) has an infinite model. By the weak version of the Upward Löwenheim-Skolem Theorem, eldiag($\mathcal{A}$) has arbitrarily large models and by Lemma 1.39, $\mathcal{A}$ elementarily embeds into the reduct of any model of eldiag($\mathcal{A}$) to τ. □

Corollary 1.41. *Let Φ be a theory over a vocabulary τ. Let κ be an infinite cardinal $\geq |\tau|$. If Φ has an infinite model, then Φ has a model of size κ.*

Definition 1.42. Let $(I, \leq)$ be a partial order, i.e., let $\leq$ be reflexive $(\forall i(i \leq i))$, transitive $(\forall i, j, k((i \leq j \land j \leq k) \to i \leq k))$ and anti-symmetric $(\forall i, j((i \leq j \land j \leq i) \to i = j))$. Then $(I, \leq)$ is **directed** if for all $i, j \in I$ there is $k \in I$ such that $i \leq k$ and $j \leq k$.

A family $(\mathcal{A}_i)_{i \in I}$ together with a family $(e_{ij})_{i,j \in I \land i \leq j}$ is an **elementary directed system** if for all $i, j \in I$ with $i \leq j$, e_{ij} is an elementary embedding of $\mathcal{A}_i$ into $\mathcal{A}_j$ and for all $i, j, k \in I$ with $i \leq j \leq k$, $e_{jk} \circ e_{ij} = e_{ik}$.

If $\mathcal{E} = ((\mathcal{A}_i)_{i \in I}, (e_{ij})_{i,j \in I \land i \leq j})$ is an elementary directed system, then a τ-structure $\mathcal{A}$ together with a family $(e_i)_{i \in I}$ is the **limit** of the system $\mathcal{E}$ if for all $i \in I$, e_i is an elementary embedding of $\mathcal{A}_i$ into $\mathcal{A}$ such that for all $j \in I$ with $i \leq j$ we have $e_i = e_j \circ e_{ij}$ and moreover, $A = \{e_i(a) : i \in I \text{ and } a \in A_i\}$.

The definition of an elementary directed system sounds more complicated than it actually is. The next theorem and its proof show that if $\mathcal{E} = ((\mathcal{A}_i)_{i \in I}, (e_{ij})_{i,j \in I \land i \leq j})$ is an elementary directed system, we may assume, without loss of generality, that for all $i, j \in I$ with $i \leq j$, $\mathcal{A}_i \preccurlyeq \mathcal{A}_j$ and e_{ij} is the identity function from A_i to A_j. In this situation, the underlying set of the limit of $\mathcal{E}$ is the union of the A_i, the e_i are again identity functions, and the $\mathcal{A}_i$ are elementary submodels of the limit $\mathcal{A}$.

Theorem 1.43. *Every elementary directed system*

$$\mathcal{E} = ((\mathcal{A}_i)_{i \in I}, (e_{ij})_{i,j \in I \land i \leq j})$$

over a vocabulary τ has a limit $\mathcal{A}$.

Proof. We may assume that the $\mathcal{A}_i$ are pairwise disjoint. Let $B = \bigcup_{i \in I} A_i$. We define an equivalence relation $\sim$ on B as follows: for $a, b \in B$ let $a \sim b$ if there are $i, j, k \in I$ such that $i, j \leq k$ $a \in A_i$, $b \in A_j$ and $e_{ik}(a) = e_{jk}(b)$, The relation $\sim$ is symmetric by definition and reflexive since $\leq$ is. Showing the transitivity of $\sim$ requires slightly more work.

Let $a \sim b$ and $b \sim c$. Let $i_0, j_0, k_0, i_1, j_1, k_1 \in I$ be such that $i_0, j_0 \leq k_0$ and $i_1, j_1 \leq k_1$ and $a \in A_{i_0}$ and $b \in A_{j_0}$ and $b \in A_{i_1}$ and $c \in A_{j_1}$ and $e_{i_0k_0}(a) = e_{j_0k_0}(b)$ and $e_{i_1k_1}(b) = e_{j_1k_1}(c)$. Since the A_i were assumed to be pairwise disjoint, j_0 and i_1 are the same. Since I is directed, there is $k \in I$ such that $k_0, k_1 \leq k$. Now

$$\begin{aligned} e_{i_0k}(a) = e_{k_0k}(e_{i_0k_0}(a)) = e_{k_0k}(e_{j_0k_0}(b)) \\ = e_{k_1k}(e_{i_1k_1}(b)) = e_{k_1k}(e_{j_1k_1}(c)) = e_{j_1k}(c) \end{aligned}$$

and therefore $a \sim c$.

Now let A be the set $B/\sim$ of $\sim$-equivalence classes. For each $i \in I$ and each $a \in A_i$ let $e_i(a)$ be the $\sim$-equivalence class of a. By the definition of $\sim$, for all $i, j \in I$ with $i \leq j$ and for all $a \in A_i$ and $b \in A_j$ with $e_{ij}(a) = b$ we have $e_i(a) = e_j(b)$ and thus $e_i = e_j \circ e_{ij}$. Clearly,

$$A = \{e_i(a) : i \in I \text{ and } a \in A_i\}.$$

We now define the interpretations of the symbol in τ in $\mathcal{A}$.

If $c \in \tau$ is a constant symbol, for every $i \in I$ we have $c^{\mathcal{A}_i} \in A_i$. Given $i, j \in I$, there is $k \in I$ such that $i, j \leq k$. Since e_{ik} and e_{jk} are elementary embeddings, $e_{ik}(c^{\mathcal{A}_i}) = c^{\mathcal{A}_k} = e_{jk}(c^{\mathcal{A}_j})$ and therefore $c^{\mathcal{A}_i} \sim c^{\mathcal{A}_j}$. Let $c^{\mathcal{A}}$ be the $\sim$-equivalence class of $c^{\mathcal{A}_i}$ for an arbitrary $i \in I$. By the previous argument, $c^{\mathcal{A}}$ is well defined.

Now let $R \in \tau$ be an n-ary relation symbol. Let $a_1, \dots, a_n \in B$. The a_m are representatives of equivalence classes in A. Each a_m is an element of some A_{i_m}. Since I is directed, there is $j \in I$ such that $i_1, \dots, i_n \leq j$. Now let

$$(e_{i_1}(a_1), \dots, e_{i_n}(a_n)) \in R^{\mathcal{A}} \quad \Leftrightarrow \quad (e_{i_1j}(a_1), \dots, e_{i_nj}(a_n)) \in R^{\mathcal{A}_j}.$$

$R^{\mathcal{A}}$ is well defined since I is directed and the e_{ij} are elementary embeddings.

Finally, let $f \in \tau$ be an n-ary function symbol. Let $a_1, \dots, a_n \in A$ and choose $i_1, \dots, i_n, j \in I$ as in the definition of the interpretation of a relation symbol. Now let

$$f^{\mathcal{A}}(e_{i_1}(a_1), \dots, e_{i_n}(a_n)) = e_j(f^{\mathcal{A}_j}(e_{i_1j}(a_1), \dots, e_{i_nj}(a_n))).$$

Again, $f^{\mathcal{A}}$ is well defined by the directedness of I and since the e_{ij} are elementary embeddings.

This finishes the definition of the structure $\mathcal{A}$. We have to show that the e_i, $i \in I$, are elementary embeddings. We use the same induction on the complexity of formulas as in the proof of the Tarski-Vaught Criterion. Again, the only nontrivial case is the case of an existential formula.

So let $\varphi(x, y_1, \dots, y_n)$ be a formula. Assume that for all $i \in I$ and $a, b_1, \dots, b_n \in A_i$ we have

$$\mathcal{A}_i \models \varphi(a, b_1, \dots, b_n) \quad \Leftrightarrow \quad \mathcal{A} \models \varphi(e_i(a), e_i(b_1), \dots, e_i(b_n)).$$

Let $i \in I$, $b_1, \ldots, b_n \in A_i$ and suppose that $\mathcal{A}_i \models (\exists x\varphi)(b_1, \ldots, b_n)$. Choose $a \in A_i$ such that $\mathcal{A} \models \varphi(a, b_1, \ldots, b_n)$. By our assumption,

$$\mathcal{A} \models \varphi(e_i(a), e_i(b_1), \ldots, e_i(b_n))$$

and hence

$$\mathcal{A} \models (\exists x\varphi)(e_i(b_1), \ldots, e_i(b_n)).$$

Now suppose that

$$\mathcal{A} \models (\exists x\varphi)(e_i(b_1), \ldots, e_i(b_n)).$$

We show that $\mathcal{A}_i \models (\exists x\varphi)(b_1, \ldots, b_n)$. Let $a \in A$ be such that

$$\mathcal{A} \models \varphi(a, e_i(b_1), \ldots, e_i(b_n)).$$

Choose $j \in I$ and $b \in A_j$ such that $a = e_j(b)$. Let $k \in I$ be such that $i, j \leq k$. By our assumption,

$$\mathcal{A}_k \models \varphi(e_{jk}(b), e_{ik}(b_1), \ldots, e_{ik}(b_n))$$

and thus

$$\mathcal{A}_k \models (\exists x\varphi)(e_{ik}(b_1), \ldots, e_{ik}(b_n)).$$

Since e_{ik} is an elementary embedding, $\mathcal{A}_i \models (\exists x\varphi)(b_1, \ldots, b_n)$. This finishes the proof of the theorem. $\square$

Corollary 1.44 (Tarksi's Elementary Chain Theorem)**.** *Let $(I, \leq)$ be a linear order. For each $i \in I$ let $\mathcal{A}_i$ be a τ-structure and assume that for all $i, j \in I$ with $i \leq j$ we have $\mathcal{A}_i \preccurlyeq \mathcal{A}_j$. Let $A = \bigcup_{i \in I} A_i$ and define a τ-structure $\mathcal{A}$ on A in the natural way. Then for every $i \in I$, $\mathcal{A}_i \preccurlyeq \mathcal{A}$.*

Exercise 1.45 (For people with some background in set theory)**.** Let κ be an uncountable cardinal. Recall that κ is regular if it is not the supremum of a set $A \subseteq \kappa$ of size $< \kappa$. A set $C \subseteq \kappa$ is **unbounded** in κ if for all $\alpha < \kappa$ there is $\beta \in C$ such that $\alpha \leq \beta$. $C \subseteq \kappa$ is **closed** if for all $S \subseteq C$, $\sup S = \kappa$ or $\sup S \in C$. C is **club** if C is both closed and unbounded in κ.

Let τ be a countable vocabulary and let κ be an uncountable regular cardinal. Show that for every τ-structure $\mathcal{A}$ on κ, the set of $\alpha < \kappa$ such that α carries an elementary substructure of $\mathcal{A}$ is club in κ.

Exercise 1.46. Let τ be a countable vocabulary and let $\mathcal{A}$ be a τ-structure. Suppose the underlying set A of $\mathcal{A}$ is of some regular size κ. If $(A_\alpha)_{\alpha<\kappa}$ is an increasing sequence of subsets of A such that all A_α are of size $< \kappa$ and for all limit ordinals $\beta < \kappa$, $A_\beta = \bigcup_{\alpha<\beta} A_\alpha$. Show that there is $\alpha < \kappa$ such that A_α carries an elementary substructure of $\mathcal{A}$.

2. Properties of first order theories

2.1. Categoricity and completeness.

Definition 2.1. Let Φ be a first order theory over a vocabulary τ. The **deductive closure** of Φ is the theory

$$\mathrm{Ded}(\Phi) = \{\varphi : \Phi \models \varphi\}.$$

Φ is **deductively closed** if $\Phi = \mathrm{Ded}(\Phi)$. Φ is **complete** if it is consistent and for all sentences φ, either $\varphi \in \Phi$ or $\neg\varphi \in \Phi$.

Now let κ be a cardinal. A theory Φ is **κ-categorical** if up to isomorphism, it has exactly one model of size κ.

Observe that every complete theory is deductively closed and for every τ-structure $\mathcal{A}$, $\mathrm{Th}(\mathcal{A})$ is complete.

Lemma 2.2. *Let Φ be a deductively closed theory having only infinite models. If Φ is κ-categorical for some $\kappa \geq |\tau|$, then Φ is complete.*

Proof. Since Φ has a model of size κ, it is consistent. Assume that Φ is not complete. Then there is a sentence φ such that $\varphi, \neg\varphi \notin \Phi$. Note that since $\varphi \notin \Phi$ and Φ is deductively closed, $\Phi \not\models \varphi$. It follows that $\Phi \cup \{\neg\varphi\}$ is has a model. By the Löwenheim-Skolem Theorem, $\Phi \cup \{\neg\varphi\}$ has a model $\mathcal{A}$ of size κ. Similarly, $\Phi \cup \{\varphi\}$ has a model $\mathcal{B}$ of size κ. Since $\mathcal{A}$ and $\mathcal{B}$ are not elementarily equivalent, they cannot be isomorphic, contradicting the κ-categoricity of Φ. □

Example 2.3. Let τ be the empty vocabulary. The only relation that we can talk about using this vocabulary is equality. For each $n \in \mathbb{N}$ let φ_n be the sentence over τ that says that there are at least n dinstinct elements. Now $\Phi = \mathrm{Ded}(\{\varphi_n : n \in \mathbb{N}\})$ is the **theory of infinite sets**. Φ is κ-categorical for every infinite κ. By Lemma 2.2, Φ is complete.

Example 2.4. Let τ be the vocabulary of vector spaces over $\mathbb{Q}$. Let Φ be the deductive closure of the set of axioms for vector spaces over $\mathbb{Q}$ together with an axiom that says that the vector space has at least two elements. A $\mathbb{Q}$-vector space V is countable if and only if its dimension is at least 1 and at most $\aleph_0$. If V is uncountable, then the dimension of V is equal to the cardinality. It follows that Φ is not $\aleph_0$-categorical but κ-categorical for all uncountable κ. Again by Lemma 2.2, Φ is complete.

Our next example of a complete theory is the **theory of dense linear orders without endpoints**. The vocabulary consists of a single binary relation symbol $\leq$. We freely use $x < y$ as an abbreviation for $x \leq y \wedge x \neq y$. The theory is the deductive closure of the set consisting of the following axioms:

(DO1) $\forall x(x \leq x)$ (Reflexivity)
(DO2) $\forall x \forall y \forall z((x \leq y \wedge y \leq z) \to x \leq z)$ (Transitivity)

(DO3) $\forall x \forall y((x \leq y \land y \leq x) \to x = y)$ (Anti-symmetry)
(DO4) $\forall x \forall y(x \leq y \lor y \leq x)$ (Linearity)
(DO5) $\forall x \forall z(x < z \to \exists y(x < y \land y < z))$ (Density)
(DO6) $\forall x \exists y \exists z(y < x \land x < z)$ (No endpoints)

Clearly, both $(\mathbb{Q}, \leq)$ and $(\mathbb{R}, \leq)$ are dense linear orders without endpoints.

Theorem 2.5. *The theory of dense linear orders without endpoints is $\aleph_0$-categorical.*

Proof. Let $(A, \leq)$ and $(B, \leq)$ be countable linear orders without endpoints. Let $(a_n)_{n\in\mathbb{N}}$ and $(b_n)_{n\in\mathbb{N}}$ be 1-1 enumerations of A and B, respectively. We will recursively define 1-1 sequences $(p_n)_{n\in\mathbb{N}}$ and $(q_n)_{n\in\mathbb{N}}$ in A, respectively B such that the map that assigns q_n to p_n is an isomorphism between $(A, \leq)$ and $(B, \leq)$.

Let $p_0 = a_0$ and $q_0 = b_0$. Suppose for some $n \in \mathbb{N}$ we have defined p_m and q_m for all $m \leq n$. If n even, let p_{n+1} be the $a_\ell \in A \setminus \{p_0, \dots, p_n\}$ with the smallest index. We distinguish three cases:

(1) For all $m \leq n$, $p_{n+1} < p_m$.
(2) For all $m \leq n$, $p_{n+1} > p_m$.
(3) Not (1) or (2).

In case (1) choose $q_{n+1} \in B$ such that $q_{n+1} < q_m$ for all $m \leq n$. This is possible since $(B, \leq)$ has no endpoints. Similarly, in case (2) we can choose $q_{n+1} \in B$ such that $q_{n+1} > q_m$ for all $m \leq n$. In case (3) there are $m_0, m_1 \leq m$ such that p_{m_0} is the $\leq$-maximal element of $\{p_m : m \leq n, p_m < p_{n+1}\}$ and p_{m_1} is the $\leq$-minimal element of $\{p_m : m \leq n, p_m > p_{n+1}\}$. Choose $q_{n+1} \in B$ such that $q_{m_0} < q_{n+1} < q_{m_1}$. This is possible since $(B, \leq)$ is a dense linear order.

If n is odd, we choose q_{n+1} to be the $b_\ell \in B \setminus \{q_0, \dots, q_n\}$ with the smallest index. We choose p_{n+1} exactly as we chose q_{n+1} in the case of even n, with the roles of A and B, a and b, and p and q switched. The finishes the definition of the sequences $(p_n)_{n\in\mathbb{N}}$ and $(q_n)_{n\in\mathbb{N}}$. It is easily verified that the map $f : A \to B; p_n \mapsto q_n$ is an isomorphism between the structures $(A, \leq)$ and $(B, \leq)$. □

This way of constructing an isomorphism between two structures is known as **back-and-forth argument**.

Corollary 2.6. *The theory of dense linear orders without endpoints is complete.*

Note that the theory of dense linear orders without endpoints is not $2^{\aleph_0}$-categorical, i.e., not $|\mathbb{R}|$-categorical. $(\mathbb{R}, \leq)$ and $(\mathbb{R} \setminus \{0\}, \leq)$ are both dense linear orders without endpoints, but they are not isomorphic since one is Dedekind complete (every bounded subset has a least upper bound) and the other is not. In fact, the theory is not κ-categorical for any uncountable κ.

Exercise 2.7. Show that every countable linear order embeds into $(\mathbb{Q}, \leq)$.

Exercise 2.8. Let $(L, \leq)$ be a linear order. Then $D \subseteq L$ is **dense in** L if every nonempty open interval of L contains an element of D. $(L, \leq)$ is **separable** if it has a countable dense subset. Show that a linear order is separable iff in embeds into lexicographically ordered product $\mathbb{R} \times \{0, 1\}$.

Note that the theory of dense linear orders without endpoints is not $2^{\aleph_0}$-categorical, i.e., not $|\mathbb{R}|$-categorical. $(\mathbb{R}, \leq)$ and $(\mathbb{R} \setminus \{0\}, \leq)$ are both dense linear orders without endpoints, but they are not isomorphic since one is Dedekind complete (every bounded subset has a least upper bound) and the other is not. In fact, we have the following theorem:

Theorem 2.9. *The theory of dense linear orders without endpoints is not κ-categorical for any uncountable κ. Moreover, for each uncountable κ, there are 2^κ pairwise non-isomorphic dense linear orders without endpoints of size κ.*

Proof. The main building blocks are the following two dense linear orders without endpoints: One is the familiar linear order $\mathcal{Q} = (\mathbb{Q}, \leq)$. In order to define the second linear order, let $X \subseteq (0, 1)$ be a set of size $\aleph_1$. By the Downward Löwenheim-Skolem Theorem there is a set $P \subseteq \mathbb{R}$ of size $\aleph_1$ such that $X \subseteq P$ and

$$(P, 0, 1, +, \cdot) \preccurlyeq (\mathbb{R}, 0, 1, +, \cdot).$$

$\mathcal{P} = (P, \leq)$ is the second linear order we will use.

$\mathcal{P}$ has the following property: For all $p, q \in P$ with $p < q$, the set $\{r \in P : p < r < q\}$ is of size $\aleph_1$. We say that P is **$\aleph_1$-dense**.

This can be seen as follows: Since $(P, 0, 1, +, \cdot)$ is an elementary substructure of $(\mathbb{R}, 0, 1, +, \cdot)$, $1 \in P$ and P is closed under addition, multiplication and division. It follows that $\mathbb{Q} \subseteq P$. Now, if $p, q \in P$ are such that $p < q$, then there is a affine linear map f over $\mathbb{Q}$ that is 1-1 and maps $(0, 1)$ into (p, q). Since $\mathbb{Q} \subseteq P$ and since P is closed under addition and multiplikation, $f[X] \subseteq P$. Now $f[X]$ is a subset of $\{r \in P : p < r < q\}$ of size $\aleph_1$. This shows the $\aleph_1$-density of P.

From $\mathcal{P}$ and $\mathcal{Q}$ we construct two new linear orders, $\mathcal{A}_0$ and $\mathcal{A}_1$. $\mathcal{A}_0$ simply consists of one copy of $\mathcal{Q}$ followed by a copy of $\mathcal{P}$. $\mathcal{A}_1$ consists of a copy of $\mathcal{Q}$ followed by ω_1 copies of $\mathcal{P}$, where all elements of the α-th copy of $\mathcal{P}$ are smaller than all elements of the β-th copy if $\alpha < \beta < \omega_1$.

Now, for each function $f : \kappa \to \{0, 1\}$ we define a structure $\mathcal{B}_f$ as follows: B_f is the union of κ disjoint linear orders. For $\alpha < \kappa$, the α-th linear order is a copy of $\mathcal{A}_{f(\alpha)}$. Again, if $\alpha < \beta < \kappa$, then all elements of the α-th linear order are below all elements of the β-th linear order. It is clear that each $\mathcal{B}_f$ is a dense linear order without endpoints of size κ.

We finish the proof of the theorem by showing that for $f, g : \kappa \to \{0,1\}$ with $f \neq g$ we have $\mathcal{B}_f \not\cong \mathcal{B}_g$. A subset of a linear order is **convex** if with any two points it contains all the points between the two. The copies of $\mathcal{Q}$ that have been used in the construction of $\mathcal{B}_f$ and $\mathcal{B}_g$ are maximal countably infinite convex subsets of B_f, respectively B_g. Now suppose that $f \neq g$ and $\mathcal{B}_f \cong \mathcal{B}_g$. Let $i : B_f \to B_g$ be an isomorphism. An easy transfinite induction shows that for each $\alpha < \kappa$, i maps the α-th maximal countably infinite subset of B_f to the α-th maximal countably infinite subset of B_g. Now let $\alpha < \kappa$ be minimal with $f(\alpha) \neq g(\alpha)$. Without loss of generality we may assume that $f(\alpha) = 0$ and $g(\alpha) = 1$.

Then i has maps the copy of $\mathcal{A}_0$ in $\mathcal{B}_f$ that contains the α-th maximal countably infinite subset of B_f to the copy of $\mathcal{A}_1$ in $\mathcal{B}_g$ that contains the α-th maximal countably infinite subset of B_g. It follows that i restricts to an isomorphism between a copy of $\mathcal{A}_0$ and a copy of $\mathcal{A}_1$. But $\mathcal{A}_0 \not\cong \mathcal{A}_1$, a contradiction. □

We now produce an example of a complete theory that is not κ-categorical for any infinite κ. The easiest way to come up with a complete theory is to consider the theory of a stucture.

Let $\mathcal{A}$ be the structure whose underlying set is

$$A = (\mathbb{N} \times \{0\}) \cup (\mathbb{N} \times \{1\})$$

and that has a single binary relation $\leq$ defined as follows: for all $(a,i),(b,j) \in A$ let $(a,i) \leq (b,j)$ if $i = j = 0$ and $a \leq b$ in $\mathbb{N}$. In other words, $\mathcal{A}$ is the disjoint union of a copy of $(\mathbb{N}, \leq)$ and a countably infinite set with no relation between the two or within the countably infinite set.

Theorem 2.10. *Let $\mathcal{A}$ be the structure defined above. Then* $\mathrm{Th}(\mathcal{A})$ *is a complete theory that is not κ-categorical for any infinite κ.*

Proof. First observe that the two parts of the structure are definable. Let $\varphi(x)$ be the formula $\exists y(x < y)$. Now for all $a \in A$ we have $\mathcal{A} \models \varphi(a)$ iff a is of the form $(n,0)$ for some $n \in \mathbb{N}$. Given a model $\mathcal{B}$ of $\mathrm{Th}(\mathcal{A})$, we call $\{b \in B : \mathcal{B} \models \varphi(b)\}$ the **$\mathbb{N}$-part** of the structure $\mathcal{B}$.

We first show that $\mathrm{Th}(\mathcal{A})$ is not $\aleph_0$-categorical. We introduce a new constant symbol c. For each $n \in \mathbb{N}$ let ψ_n be the sentence

$$\varphi(c) \wedge \exists x_0 \ldots \exists x_n (x_0 < x_1 \wedge \cdots \wedge x_{n-1} < x_n \wedge x_n < c).$$

Intuitively, ψ_n says that c is in the $\mathbb{N}$-part of the structure and is at least $n+1$.

Interpreting c in $\mathcal{A}$ by a pair $(n,0)$ with sufficiently large n, we see that every finite subset of $\mathrm{Th}(\mathcal{A}) \cup \{\psi_n : n \in \mathbb{N}\}$ has a model. Hence, by the Compactness Theorem, $\mathrm{Th}(\mathcal{A}) \cup \{\psi_n : n \in \mathbb{N}\}$ has a model $\mathcal{B}$. By the Downward Löwenheim-Skolem Theorem, we may assume that $\mathcal{B}$ is countable. Since $\mathcal{B} \models \varphi(c)$ and in $\mathcal{B}$ the interpretation of c has infinitely many elements below it, $\mathcal{A} \not\cong \mathcal{B}$. Note that it is unnecessary to use $\varphi(c)$ in the formulas ψ_n since once $c^{\mathcal{B}}$ is comparable with any other element of B, it has to be in the $\mathbb{N}$-part of $\mathcal{B}$.

To see that $\mathrm{Th}(\mathcal{A})$ is not κ-categorical for any uncountable κ, we argue as follows. Let κ be an uncountable cardinal. We add pairwise distinct constant symbols c_α and d_α, $\alpha < \kappa$, to our vocabulary. Now consider the theory

$$\Phi = \mathrm{Th}(\mathcal{A}) \cup \{\varphi(c_\alpha) \wedge \neg\varphi(d_\alpha) : \alpha < \kappa\} \cup \{c_\alpha \neq c_\beta : \alpha < \beta < \kappa\} \cup \{d_\alpha \neq d_\beta : \alpha < \beta < \kappa\}.$$

As in the proof of the Upward Löwenheim-Skolem Theorem, it is easily checked that every subset of Φ has a model. Hence Φ has a model by the Compactness Theorem. By the Downward Löwenheim-Skolem Theorem, Φ has a model $\mathcal{B}$ of size κ.

We use the same notation $\mathcal{B}$ for the structure $\mathcal{B}$ and for its reduct to the original vocabulary $\{\leq\}$. Clearly, both the $\mathbb{N}$-part of $\mathcal{B}$ and its complement are of size κ. Now consider the structure $\mathcal{B}'$ that consists of the $\mathbb{N}$-part of $\mathcal{B}$ and a countably infinite subset of the complement of the $\mathbb{N}$-part. A straight forward application of the Tarski-Vaught-Criterion shows that $\mathcal{B}'$ is an elementary substructure of $\mathcal{B}$.

Since the $\mathbb{N}$-part of $\mathcal{B}'$ is the same as of $\mathcal{B}$, $\mathcal{B}'$ is of size κ. But since the complement of the $\mathbb{N}$-part of $\mathcal{B}'$ is only countable, $\mathcal{B} \not\cong \mathcal{B}'$. Hence $\mathrm{Th}(\mathcal{A})$ is not κ-categorical. □

We provide another example of a theory that is $\aleph_0$-categorical and not κ-categorical for any uncountable κ.

Definition 2.11. Recall that a **graph** G is a set V of **vertices** with a binary relation E that is irreflexive and symmetric. E is the **edge relation** of G and two vertices $v, w \in V$ are connected by an **edge** if they are related by E. The unordered pairs $\{x, y\}$, $(x, y) \in E$, are the **edges** of G

A graph G on a countably infinite set V of vertices is **random** if for all finite disjoint sets $A, B \subseteq V$ there is a vertex $v \in V \setminus (A \cup B)$ such that all $w \in A$ are connected to v by an edge and no $u \in B$ is connected to v. For all $n, m \in \mathbb{N}$ with $n, m > 0$ let $\varphi_{n,m}$ be the sentence

$$\forall x_1 \dots \forall x_n \forall y_1 \dots \forall y_n (\text{“}\{x_1, \dots, x_n\} \text{ and } \{y_1, \dots, y_m\} \text{ are disjoint”} \to \exists z (z \neq x_1 \wedge \dots \wedge z \neq x_n \wedge z \neq y_1 \wedge \dots \wedge z \neq y_m \wedge E(x_1, z) \wedge \dots \wedge E(x_n, z) \wedge \dots \wedge \neg E(y_1, z) \wedge \dots \wedge \neg E(y_m, z)).$$

If $n = 0$ and $m > 0$ let $\varphi_{n,m}$ be the sentence

$$\forall y_1 \dots \forall y_m \exists z (\neg E(y_1, z) \wedge \dots \wedge \neg E_(y_m, z)).$$

If $n > 0$ and $m = 0$ let $\varphi_{n,m}$ be the sentence

$$\forall x_1 \dots \forall x_n \exists z (E(x_1, z) \wedge \dots \wedge E_(x_n, z)).$$

Let $\varphi_{0,0}$ be the sentence $\exists z(z = z)$. The $\varphi_{n,m}$ are **extension axioms**.

Clearly, a countable structure (V, E), where E is a binary relation, is a random graph iff it satisfies the theory

$$\Phi = \{\forall x(\neg E(x,x)), \forall x \forall y(E(x,y) \leftrightarrow E(y,x))\} \cup \{\varphi_{n,m} : n, m \in \mathbb{N}\}.$$

Theorem 2.12. *Up to isomorphism, there is exactly one random graph. In other words, Φ is $\aleph_0$-categorical.*

Proof. We first prove the existence of the random graph. Let $V = \mathbb{N}$. Let $(A_n)_{n\in\mathbb{N}}$ and $(B_n)_{n\in\mathbb{N}}$ be two sequences of finite subsets of $\mathbb{N}$ such that for all finite sets $A, B \subseteq \mathbb{N}$ there are infinitely many $n \in \mathbb{N}$ such that $A_n = A$ and $B_n = B$. Let

$$F = \{(n,m) : n < m,\ A_m \text{ and } B_m \text{ are disjoint subsets of } \{0, \dots, m-1\}, \text{ and } n \in A_m\}$$

Let $E = \{(n,m) : (n,m) \in F \text{ or } (m,n) \in F\}$. Clearly, (V, E) is a countably infinite graph. To show that it is random, let $A, B \subseteq V$ be disjoint and finite. By the choice of the sequences $(A_n)_{n\in\mathbb{N}}$ and $(B_n)_{n\in\mathbb{N}}$ there is $m > \max(A \cup B)$ such that $A_m = A$ and $B_m = B$. Now the vertex m is connected to all vertices in A and not connected to the vertices in B.

Next we show that any two random graphs are isomorphic. We use a back-and-forth argument. Let $G_i = (V_i, E_i)$ be a random graph for $i \in \{0, 1\}$. Let $(a_n)_{n\in\mathbb{N}}$ be a 1-1 enumeration of V_0 and let $(b_n)_{n\in\mathbb{N}}$ be a 1-1 enumeration of V_1. We construct two sequences $(p_n)_{n\in\mathbb{N}}$ and $(q_n)_{n\in\mathbb{N}}$ as follows:

Let $p_0 = a_0$ and $q_0 = b_0$. If we have chosen p_m and q_m for all $m < n$, we distinguish two cases. If n is even, let p_n be the $a_\ell \in V_0 \setminus \{p_m : m < n\}$ of the smallest index. Let $A = \{m : m < n \text{ and } (p_m, p_n) \in E_0\}$ and $B = \{m : m < n \text{ and } (p_m, p_n) \notin E_0\}$. Choose $q_n \in B \setminus \{q_m : m < n\}$ such that q_n is connected to all vertices in $\{q_m : m \in A\}$ and not connected to any vertex in $\{q_m : m \in B\}$.

If n is odd, find p_n and q_n as above, but after switching the roles of p and q, (V_0, E_0) and (V_1, E_1) and of a and b, respectively. Now $f : V_0 \to V_1; p_n \mapsto q_n$ is an isomorphism between (V_0, E_0) and (V_1, E_1). $\square$

Because of this theorem, we refer to **the** random graph rather than **a** random graph.

Corollary 2.13. $\mathrm{Ded}(\Phi_{\mathrm{random}})$ *is complete.*

Exercise 2.14. Show that every countable graph embeds into the random graph.

Theorem 2.15. *The theory of the random graph is not $\aleph_1$-categorical.*

Proof. We construct two non-isomorphic models of Φ_{random} of size $\aleph_1$. We construct two strictly increasing sequences $(G_\alpha)_{\alpha<\omega_1}$ and $(H_\alpha)_{\alpha<\omega_1}$ of countable graphs. Recall that for a graph G, by $V(G)$ we denote the set of vertices of G and by $E(G)$ we denote the edge relation of G.

We start by letting G_0 be the countably infinite graph with no edges and H_0 the countably infinite complete graph, i.e., a countably infinite graph in which any two vertices are connected by an edge. If $\delta < \omega_1$ is a limit ordinal, let

$$G_\delta = \left(\bigcup_{\alpha<\delta} V(G_\alpha), \bigcup_{\alpha<\delta} E(G_\alpha) \right)$$

and

$$H_\delta = \left(\bigcup_{\alpha<\delta} V(H_\alpha), \bigcup_{\alpha<\delta} E(H_\alpha) \right).$$

Now suppose that G_α and H_α have been constructed for some $\alpha < \omega_1$. Let $(A_n)_{n\in\mathbb{N}}$ be an enumeration of all finite subsets of $V(G_\alpha)$. Let $(B_n)_{n\in\mathbb{N}}$ be an enumeration of all finite subsets of $V(H_\alpha)$. Choose pairwise distinct a_n and b_n, $n \in \mathbb{N}$, outside $V(G_\alpha)$, respectively $V(H_\alpha)$. Let $V(G_{\alpha+1}) = V(G_\alpha) \cup \{a_n : n \in \mathbb{N}\}$ and $V(H_{\alpha+1}) = V(H_\alpha) \cup \{b_n : n \in \mathbb{N}\}$. Choose $E(G_{\alpha+1})$ such that $E(G_{\alpha+1}) \restriction V(G_\alpha) = E(G_\alpha)$ and for every $n \in \mathbb{N}$, a_n is connected to every element of A_n and not connected to any element of $V(G_\alpha) \setminus A_n$. Choose $E(H_{\alpha+1})$ such that $E(H_{\alpha+1}) \restriction V(H_\alpha) = E(H_\alpha)$ and for every $n \in \mathbb{N}$, b_n is connected to every element of $V(H_\alpha) \setminus B_n$ and not connected to any element of B_n. Finally, let

$$G = \left(\bigcup_{\alpha<\omega_1} V(G_\alpha), \bigcup_{\alpha<\omega_1} E(G_\alpha) \right)$$

and

$$H = \left(\bigcup_{\alpha<\omega_1} V(H_\alpha), \bigcup_{\alpha<\omega_1} E(H_\alpha) \right).$$

We claim that G and H are both models of Φ_{random}, but not isomorphic. It is clear that G and H are graphs. If A and B are disjoint finite subsets of $V(G)$, then there is $\alpha < \omega_1$ with $A \cup B \subseteq V(G_\alpha)$. For the construction of $G_{\alpha+1}$ we chose an enumeration $(A_n)_{n\in\mathbb{N}}$ of all finite subsets of $V(G_\alpha)$. Let $n \in \mathbb{N}$ be such that $A = A_n$. We chose some a_n such that $a_n \in V(G_{\alpha+1}) \setminus V(G_\alpha)$ and a_n is connected to all elements of A_n but not connected to any element of $V(G_\alpha) \setminus A_n$. Since B is disjoint from A, a_n is not connected to any element of B. This shows that G satisfies all the $\psi_{n,m}$. A symmetric argument shows that H satisfies all the $\psi_{n,m}$.

We now show that G and H are not isomorphic. Let $f : V(G) \to V(H)$ be a bijection. We show that f cannot be an isomorphism.

We choose a strictly increasing sequence $(\alpha_n)_{n\in\mathbb{N}}$ of ordinals $< \omega_1$ as follows:

Let $\alpha_0 = 0$. If n is even, let α_{n+1} be such that $f[V(G_{\alpha_n}] \subseteq V(H_{\alpha_{n+1}})$. This is possible since $V(G_{\alpha_n})$ is countable and f is onto. If n is odd, let α_{n+1} be such that $f^{-1}[V(H_{\alpha_n})] \subseteq V(G_{\alpha_{n+1}})$. This is possible since $V(H_{\alpha_n})$ is countable and f is 1-1. Finally let $\alpha = \sup\{\alpha_n : n \in \mathbb{N}\}$.

Now we have $f[V(G_\alpha)] = V(H_\alpha)$. By the construction of G, every vertex in $V(G) \setminus V(G_\alpha)$ is connected to only finitely many vertices in $V(G_\alpha)$. By the construction of H, every vertex $V(H) \setminus V(H_\alpha)$ is connected to infinitely many vertices in $V(H_\alpha)$. It follows that f is not an isomorphism. □

Exercise 2.16. Show that the graph H in the proof of the previous theorem has an uncountable complete subgraph while G does not. This also shows that G and H are not isomorphic.

Exercise 2.17. Show that every infinite graph can be embedded into a graph of the same size that is a model Φ_{random}.

Notice that the proof of Theorem 2.15 easily generalizes to all uncountable cardinals κ. Hence Φ_{random} is not κ-categorical for any $\kappa > \aleph_0$. This is not by accident. We have the following theorem.

Theorem 2.18 (Morley)**.** *If a countable first order theory is κ-categorical for some uncountable κ, then is is κ-categorical for all uncountable κ.*

We mention another important family of complete theories that are not $\aleph_0$-categorical but κ-categorical for all uncountable κ. For each prime number p let Φ_{ACF_p} be the theory of **algebraically closed fields of characteristic** p. Φ_{ACF_p} is the deductive closure of the axioms of field theory together with the axiom saying that the sum of p 1's is zero together with the infinitely many axioms that say that that every polynomial that is not constant has a zero. The following axiom says that every polynomial of degree $n > 0$ has a zero:

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

Here y^k is an abbreviation for the product $y \cdot \cdots \cdot y$ with k factors.

Similarly, let Φ_{ACF_0} be the theory of **algebraically closed fields of characteristic zero**, i.e., the deductive closure of the axioms of field theory together with the infinitely many axioms that say that the field is not of characteristic p for any prime number p together with the infinitely many axioms that say that every polynomial that is not constant has a zero.

Theorem 2.19. *If p is either 0 or a prime number, then Φ_{ACF_p} is not $\aleph_0$-categorical but κ-categorical for all uncountable κ. In particular, Φ_{ACF_p} is complete.*

Proof. Two algebraically closed fields of characteristic p are isomorphic iff they have the same transcendence degree over their minimal subfield, **the prime subfield**. If a field is uncountable, then its cardinality equals the transcendence degree over the prime subfield. $\square$

2.2. **Quantifier elimination.**

Definition 2.20. Let Φ be a first order theory over a vocabulary τ. Two τ-formulas $\varphi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ are Φ-equivalent if

$$\Phi \models \varphi(x_1, \ldots, x_n) \leftrightarrow \psi(x_1, \ldots, x_n).$$

Φ has **quantifier elimination** if for every τ-formula is Φ-equivalent to a quantifier-free formula. A τ-structure $\mathcal{A}$ has quantifier elimination if the theory $\mathrm{Th}(\mathcal{A})$ does.

Lemma 2.21. *A first order theory Φ has quantifier elimination iff for every quantifier-free formula $\varphi(x, y_1, \ldots, y_n)$ there is a quantifier-free formula $\psi(y_1, \ldots, y_n)$ such that*

$$T \models \exists x \varphi(x, y_1, \ldots, y_n) \leftrightarrow \psi(y_1, \ldots, y_n).$$

Proof. By induction on the complexity of formulas. □

Lemma 2.22. *Let τ be a finite vocabulary without function symbols, i.e., a finite* **relational vocabulary**. *A complete theory Φ over τ has quantifier elimination iff for every model $\mathcal{A}$ of Φ and all n-tupels $(a_1, \ldots, a_n), (b_1, \ldots, b_n) \in A^n$ the following holds:*

If $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ satisfy the same atomic formulas in $\mathcal{A}$, then for every $a_{n+1} \in A$ there is $b_{n+1} \in A$ such that $(a_1, \ldots, a_{n+1})$ and $(b_1, \ldots, b_{n+1})$ satisfy the same atomic formulas in $\mathcal{A}$.

Proof. We start with a couple observations. Let $\mathcal{A}$ be a τ-structure. Since τ is finite and relational, there are only finitely many atomic formulas in the variables $x_1, \ldots, x_n$. Moreover, if two n-tuples $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ satisfy the same atomic formulas in a structure $\mathcal{A}$, then they satisfy the same quantifier free formulas. Low let $a_1, \ldots, a_n \in \mathcal{A}$. Let $\chi_{(a_1,\ldots,a_n)}(x_1, \ldots, x_n)$ be the conjunction of all formulas $\psi(x_1, \ldots, x_n)$ that are atomic or negations of atomic formulas and such that $\mathcal{A} \models \psi(a_1, \ldots, a_n)$. Now an n-tuple $(b_1, \ldots, b_n) \in A^n$ satisfies the same atomic formulas as $(a_1, \ldots, a_n)$ iff

$$\mathcal{A} \models \chi_{(a_1,\ldots,a_n)}(b_1, \ldots, b_n).$$

We now start the actual proof of the lemma. First assume that Φ has quantifier elimination. Let $\mathcal{A}$ be a model of Φ and suppose that the n-tuples $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ satisfy the same atomic formulas in $\mathcal{A}$. Let $a_{n+1} \in A$. Now

$$\mathcal{A} \models \chi_{(a_1,\ldots,a_{n+1})}(a_1, \ldots, a_{n+1})$$

and therefore

$$\mathcal{A} \models (\exists x_{n+1} \chi_{(a_1,\ldots,a_{n+1})})(a_1, \ldots, a_n).$$

Since Φ has quantifier elimination, $\exists x_{n+1} \chi_{(a_1,\ldots,a_{n+1})})$ is Φ-equivalent to a quantifier free formula $\varphi(x_1, \ldots, x_n)$. But since $(a_1, \ldots, a_n)$ and

$(b_1, \dots, b_n)$ satisfy the same atomic formulas, they satisfy the same quantifier-free formulas. It follows that

$$\mathcal{A} \models \varphi(b_1, \dots, b_n)$$

and therefore

$$\mathcal{A} \models (\exists x_{n+1} \chi_{(a_1,\dots,a_{n+1})})(b_1, \dots, b_n).$$

Let $b_{n+1} \in A$ be such that

$$\mathcal{A} \models \chi_{(a_1,\dots,a_{n+1})}(b_1, \dots, b_{n+1}).$$

Now $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas.

On the other hand, suppose that for some model $\mathcal{A}$ of Φ we have that if $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ satisfy the same atomic formulas in $\mathcal{A}$, then for every $a_{n+1} \in A$ there is $b_{n+1} \in A$ such that $(a_1, \dots, a_{n+1})$ and $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas in $\mathcal{A}$.

Let $\varphi(x_1, \dots, x_{n+1})$ be a quantifier-free formula over τ. Let $(a_1, \dots, a_n) \in A^n$ be such that

$$\mathcal{A} \models (\exists x_{n+1}\varphi)(a_1, \dots, a_n).$$

Let $a_{n+1} \in A$ be such that $\mathcal{A} \models \varphi(a_1, \dots, a_{n+1})$. Now let $(b_1, \dots, b_n)$ be an n-tuple that satisfies the same atomic formulas in $\mathcal{A}$ as $(a_1, \dots, a_n)$. By our assumption, there is b_{n+1} such that $(b_1, \dots, b_{n+1})$ satisfies the same atomic formulas as $(a_1, \dots, a_{n+1})$. In other words,

$$\mathcal{A} \models (\exists x_{n+1}\varphi)(a_1, \dots, a_n) \to (\exists x_{n+1}\varphi)(b_1, \dots, b_n).$$

It follows that whether or not $\mathcal{A}$ satisfies $(\exists x_{n+1}\varphi)(a_1, \dots, a_n)$ only depends on the atomic formulas that $(a_1, \dots, a_n)$ satisfies. Let

$$X = \{\chi_{(a_1,\dots,a_n)}(x_1, \dots, x_n) : a_1, \dots, a_n \in A \text{ and } \mathcal{A} \models (\exists x_{n+1}\varphi)(a_1, \dots, a_n)\}.$$

Since there are only finitely many atomic formulas in the variables $x_1, \dots, x_{n+1}$, the set X is finite. Let

$$\psi(x_1, \dots, x_n) = \bigvee_{\chi \in X} \chi(x_1, \dots, x_n).$$

Since the validity of $(\exists x_{n+1}\varphi)(a_1, \dots, a_n)$ in $\mathcal{A}$ only depends on the atomic formulas satisfied by $(a_1, \dots, a_n)$, we have

$$\mathcal{A} \models \forall x_1, \dots, x_n(\exists x_{n+1}\varphi \leftrightarrow \psi).$$

Since Φ is a complete theory,

$$\Phi \models \forall x_1, \dots, x_n(\exists x_{n+1}\varphi \leftrightarrow \psi)$$

and therefore $(\exists x_{n+1}\varphi)(x_1, \dots, x_n)$ and $\psi(x_1, \dots, x_n)$ are Φ-equivalent. By Lemma 2.21, this shows that Φ has quantifier elimination. $\square$

Remark 2.23. Note that even though the previous lemma was formulated as “a complete theory Φ has quantifier elimination iff for every model $\mathcal{A}$ of Φ ...”, the proof of the lemma shows that the formulation “a complete theory Φ has quantifier elimination iff for some model $\mathcal{A}$ of Φ ...” works as well.

Corollary 2.24. *The theory of dense linear orders without endpoints has quantifier elimination.*

Proof. We use Lemma 2.22 and the remark following it. We consider the structure $\mathcal{Q} = (\mathbb{Q}, \leq)$. Let $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ be n-tuples satisfying the same atomic formulas in $\mathcal{Q}$. We may assume that $a_1 \leq \dots \leq a_n$. In this case $b_1 \leq \dots \leq b_n$. Now let $a_{n+1} \in \mathbb{Q}$. We have to find $b_{n+1} \in \mathbb{Q}$ such that $(a_1, \dots, a_{n+1})$ and $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas in $\mathcal{Q}$.

If there is $j \in \{1, \dots, n\}$ such that $a_{n+1} = a_j$, choose $b_{n+1} = b_j$. If $a_{n+1} < a_1$, choose $b_{n+1} \in \mathbb{Q}$ such that $b_{n+1} < b_1$. If $a_{n+1} > a_n$, choose $b_{n+1} > b_n$. If for some $j \in \{1, \dots, n-1\}$, $a_j < a_{n+1} < a_{j+1}$, choose $b_{n+1} \in \mathbb{Q}$ such that $b_j < b_{n+1} < b_{j+1}$. In any case, $(a_1, \dots, a_{n+1})$ and $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas in $\mathcal{Q}$. □

Corollary 2.25. Φ_{random} *has quantifier elimination.*

Proof. The argument is almost the same as in the proof of Corollary 2.24. Let $G = (V, E)$ be the random graph. Let $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ be n-tuples satisfying the same atomic formulas in G. Now let $a_{n+1} \in V$. We have to find $b_{n+1} \in V$ such that $(a_1, \dots, a_{n+1})$ and $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas in G.

If there is $j \in \{1, \dots, n\}$ such that $a_{n+1} = a_j$, choose $b_{n+1} = b_j$. Otherwise choose b_{n+1} different from all the b_j, $j \in \{1, \dots, n\}$, and such that for all $j \in \{1, \dots, n\}$, b_{n+1} is connected to b_j iff a_{n+1} is connected to a_j. This is possible by the properties of the random graph. Now clearly, $(a_1, \dots, a_{n+1})$ and $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas in G. □

Note that both the theory of dense linear orders without endpoints and Φ_{random} are $\aleph_0$-categorical. It turns out that there is a connection between categoricity and quantifier elemination.

Theorem 2.26. *Let Φ be a complete theory in a finite relational vocabulary τ that has infinite models and quantifier elimination. Then Φ is $\aleph_0$-categorical.*

Proof. Let $\mathcal{A}$ and $\mathcal{B}$ be countable models of Φ. We prove that $\mathcal{A}$ and $\mathcal{B}$ are isomorphic by using a back-and-forth argument. Let $(a_n)_{n\in\mathbb{N}}$ and $(b_n)_{n\in\mathbb{N}}$ be 1-1 enumerations of A and B, respectively. We construct enumerations $(p_n)_{n\in\mathbb{N}}$ and $(q_n)_{n\in\mathbb{N}}$ of A and B, respectively, such that for all n, $(p_0, \dots, p_n)$ and $(q_0, \dots, q_n)$ satisfy the same atomic formulas, the first n-tuple in $\mathcal{A}$, the second in $\mathcal{B}$. If we can accomplish this, then the function that maps p_n to q_n is an isomorphism between $\mathcal{A}$ and $\mathcal{B}$.

Let $p_0 = a_0$. Since τ is finite and relational, there are only finitely many atomic formulas in the variable x_1 and we can define the formula $\chi_{a_0}(x_0)$ as in the proof of Lemma 2.22. Now $\mathcal{A} \models \exists x_0 \chi_{a_0}$. Since Φ is complete, $\exists x_0 \chi_{a_0} \in \Phi$. Hence $\mathcal{B} \models \exists x_0 \chi_{a_0}$. Let $q_0 \in B$ be such that $\mathcal{B} \models \chi_{a_0}(q_0)$.

Now assume that n is even and we have already chosen $p_0, \dots, p_n \in A$ and $q_0, \dots, q_n \in B$ such that $(p_0, \dots, p_n)$ and $(q_0, \dots, q_n)$ satisfy the same atomic formulas in the respective structures. Let p_{n+1} be the $a_\ell \in A \setminus \{p_0, \dots, p_n\}$ of minimal index. Now

$$\mathcal{A} \models \chi_{(p_0,\dots,p_{n+1})}(p_0, \dots, p_{n+1})$$

and therefore

$$\mathcal{A} \models (\exists x_{n+1} \chi_{(p_0,\dots,p_{n+1})})(p_0, \dots, p_n).$$

By quantifier elimination,

$$(\exists x_{n+1} \chi_{(p_0,\dots,p_{n+1})})(x_0, \dots, x_n)$$

is Φ-equivalent to a quantifier-free formula $\varphi(x_0, \dots, x_n)$. Now

$$\mathcal{A} \models \varphi(p_0, \dots, p_n).$$

Since if $(p_0, \dots, p_n)$ and $(q_1, \dots, q_n)$ satisfy the same atomic formulas, they satisfy the same quantifier-free formulas,

$$\mathcal{B} \models \varphi(q_0, \dots, q_n).$$

It follows that

$$\mathcal{B} \models (\exists x_{n+1} \chi_{(p_0,\dots,p_{n+1})})(q_0, \dots, q_n).$$

Choose $q_{n+1} \in B$ such that

$$\mathcal{B} \models \chi_{(p_0,\dots,p_{n+1})})(q_0, \dots, q_{n+1}).$$

If n is odd, proceed in the same way with the roles of $\mathcal{A}$ and $\mathcal{B}$ switched. By the usual argument, $p_n \mapsto q_n$ defines an isomorphism between $\mathcal{A}$ and $\mathcal{B}$. $\square$

Lemma 2.27. *Let τ be any vocabulary and let Φ be a complete theory over τ. Then the following are equivalent:*

(1) *Φ has quantifier elimination.*
(2) *For every τ-structure $\mathcal{C}$, if $f : C \to M$ and $g : C \to M$ are embeddings of $\mathcal{C}$ into a model $\mathcal{M}$ of Φ, then for every quantifier-free formula τ-formula $\varphi(x, y_1, \dots, y_n)$ and every n-tuple $(c_1, \dots, c_n) \in C^n$ we have*

$$\mathcal{M} \models (\exists x\varphi)(f(c_1), \dots, f(c_n)) \quad \Leftrightarrow \quad \mathcal{M} \models (\exists x\varphi)(g(c_1), \dots, g(c_n)).$$

Proof. (1)⇒(2): Let C be a τ-structure and $f, g : C \to M$ embeddings into a model of Φ. Let $(c_1, \dots, c_n) \in C^n$. Let $\overline{a} = (a_1, \dots, a_n) = (f(c_1), \dots, f(c_n))$ and $\overline{b} = (b_1, \dots, b_n) = (g(c_1), \dots, g(c_n))$. Now $\overline{a}$ and $\overline{b}$ satisfy the same atomic formulas.

Let $\varphi(x, y_1, \dots, y_n)$ be a quantifier-free formula. Suppose

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

Since Φ has quantifier elimination, $(\exists x\varphi)(y_1, \dots, y_n)$ is Φ-equivalent to a quantifier-free formula $\psi(y_1, \dots, y_n)$. Since the n-tuples $\overline{a}$ and $\overline{b}$ satisfy the same atomic formulas, they satisfy the same quantifier-free formulas. It follows that

$$\mathcal{M} \models (\exists x\varphi)(\overline{a}) \Leftrightarrow \mathcal{M} \models \psi(\overline{a}) \Leftrightarrow \mathcal{M} \models \psi(\overline{b}) \Leftrightarrow \mathcal{M} \models (\exists x\varphi)(\overline{b}).$$

(2)⇒(1): Assume that Φ does not have quantifier elimination. By Lemma 2.21, there is a quantifier-free formula $\varphi(x, y_1, \dots, y_n)$ such that the formula $(\exists x\varphi)(y_1, \dots, y_n)$ is not Φ-equivalent to a quantifier-free formula.

Claim 2.28. There is a model $\mathcal{M}$ of Φ and there are n-tuples

$$(a_1, \dots, a_n), (b_1, \dots, b_n) \in M^n$$

that satisfy the same atomic formulas such that $\mathcal{M} \models (\exists x\varphi)(a_1, \dots, a_n)$ and $\mathcal{M} \models \neg(\exists x\varphi)(b_1, \dots, b_n)$.

For the proof of the claim, we add new constant symbols $d_1, \dots, d_n$ and $e_1, \dots, e_n$ to the vocabulary τ. Consider the theory

$$\Psi = \Phi \cup \{\psi(d_1, \dots, d_n) \leftrightarrow \psi(e_1, \dots, e_n) : \psi \text{ is an atomic formula}\} \\ \cup \{(\exists x\varphi)(d_1, \dots, d_n), \neg(\exists x\varphi)(e_1, \dots, e_n)\}.$$

We show that every finite subset Ψ_0 of Ψ has a model. Let $\mathcal{N}$ be a model of Φ and let $\Psi_0 \subseteq \Psi$ be finite. Then there is a finite set X of atomic formulas $\psi(y_1, \dots, y_n)$ such that

$$\Psi_0 \subseteq \Phi \cup \{\psi(d_1, \dots, d_n) \leftrightarrow \psi(e_1, \dots, e_n) : \psi \in X\} \\ \cup \{(\exists x\varphi)(d_1, \dots, d_n), \neg(\exists x\varphi)(e_1, \dots, e_n)\}.$$

Consider now the Boolean algebra $\mathcal{B}$ of subsets of N^n that is generated by the sets

$$\{(a_1, \dots, a_n) : \mathcal{N} \models \psi(a_1, \dots, a_n)\},$$

$\psi \in X$. Note that $\mathcal{B}$ is finite. A set $A \in \mathcal{B}$ is an **atom** if it is non-empty and does not have a proper subset in $\mathcal{B}$. Since Φ is complete, Φ knows that $(\exists x\varphi)(d_1, \dots, d_n)$ is not equivalent to a quantifier-free formula and hence not equivalent to a Boolean combination of formulas in X. It follows that the set

$$P = \{(a_1, \dots, a_n) \in N^n : \mathcal{N} \models (\exists x\varphi)(a_1, \dots, a_n)\}$$

is not an element of $\mathcal{B}$. But this implies that P is not the union of a set of atoms of $\mathcal{B}$. Hence, there is an atom $A \in \mathcal{B}$ containing n-tuples $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ with $(a_1, \dots, a_n) \in P$ and $(b_1, \dots, b_n) \notin P$. Since A is an atom of $\mathcal{B}$, $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ satisfy the same formulas in X. Interpreting d_i by a_i and e_i by b_i for all $i \in \{1, \dots, n\}$, we obtain a model of Ψ_0.

It follows that Ψ has a model $\mathcal{M}$. For all $i \in \{1, \dots, n\}$ let $a_i = d_i^{\mathcal{M}}$ and $b_i = e_i$. Now $\mathcal{M}$ and the two n-tuples $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ work for the claim. Let $\mathcal{M}_{\overline{a}}$ be the substructure of $\mathcal{M}$ generated by $a_1, \dots, a_n$ and let $\mathcal{M}_{\overline{b}}$ be the substructure generated by $b_1, \dots, b_n$. Since the two n-tuples satisfy the same atomic formulas, the two substructures are isomorphic by an isomorphism that maps each a_i, $i \in \{1, \dots, n\}$, to the corresponding b_i. However, since $\mathcal{M} \models (\exists x\varphi)(a_1, \dots, a_n)$ and $\mathcal{M} \models \neg(\exists x\varphi)(b_1, \dots, b_n)$, the condition in the lemma fails. □

Exercise 2.29. Show that a complete theory Φ over an arbitrary vocabulary τ has quantifier elimination iff the following condition holds:

(3) For every model $\mathcal{M}$ of T and every $n \in \mathbb{N}$, if $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ satisfy the same atomic formulas in $\mathcal{M}$, then for every $a_{n+1} \in M$ there are an elementary extension $\mathcal{N}$ of $\mathcal{M}$ and an element $b_{n+1} \in N$ such that $(a_1, \dots, a_{n+1})$ and $(b_1, \dots, b_{n+1})$ satisfy the same atomic formulas in $\mathcal{N}$.

Hint: Use compactness, elementary diagrams and the previous lemma.

We are now ready to show that the theory of algebraically closed fields of characteristic p, $p = 0$ or p a prime number, has quantifier elimination.

Theorem 2.30. *Let p be a prime number or $p = 0$. Then Φ_{ACF_p} has quantifier elimination.*

Proof. We use condition (2) of Lemma 2.27. Let $\mathcal{M}$ be a model of Φ_{ACF_p}, let $\mathcal{C}$ be a τ-structure and let $g, f : \mathcal{C} \to M$ be embeddings. Let $(c_1, \dots, c_n) \in C$. Let $\overline{a} = (a_1, \dots, a_n) = (f(c_1), \dots, f(c_n))$ and let $\overline{b} = (b_1, \dots, b_n) = (g(c_1), \dots, g(c_n))$. We have to show that for each quantifier-free formula $\varphi(x, y_1, \dots, y_n)$ we have

$$\mathcal{M} \models (\exists x\varphi)(\overline{a}) \quad \Leftrightarrow \quad \mathcal{M} \models (\exists x\varphi)(\overline{b}).$$

Assume that $\mathcal{M} \models (\exists x\varphi)(\overline{a})$ and let $A = \{a_1, \dots, a_n\}$. Let R be the subring of $\mathcal{M}$ generated by A. Every element of R can be written as a term in the elements of A. Note that $R \subseteq f[C]$. Also note that an element $a \in M$ is algebraic over R iff it is algebraic over the field generated by R. Now let $a \in M$ be such that $\mathcal{M} \models \varphi(a_1, \dots, a_n, a_{n+1})$.
Case 1. There is a polynomial $p(x)$ with coefficients in R such that $p(a) = 0$.

In this case we take $p(X)$ to be a polynomial of minimal degree with coefficients in R such that $p(a) = 0$. We write the coefficients of p as terms in the elements of A. Now replace every a_i that appears in $p(X)$ by the corresponding b_i in order to obtain the polynomial $q(X)$. Let S be the ring generated by $b_1, \dots, b_n$. Since the $(a_1, \dots, a_n)$ and $(b_1, \dots, b_n)$ satisfy the same atomic formulas, there is an isomorphism between R and S that maps a_i to b_i for each $i \in \{1, \dots, n\}$. This isomorphism extends to the fields generated by the two rings and maps the coefficients of $p(X)$ to the coefficients of $q(X)$.

Since $\mathcal{M}$ is algebraically closed, there is a root $b \in M$ of $q(X)$. Since $p(X)$ does not factor into polynomials of smaller degrees over R, $q(X)$ does not factor into polynomials of smaller degrees over S. It follows that there is an isomorphism from the field generated by R and a to the field generated by S and b that maps a to b and each a_i to the corresponding b_i. It follows that $(a, a_1, \dots, a_n)$ and $(b, b_1, \dots, b_n)$ satisfy the same atomic formulas. Since φ is quantifier-free, it follows that $\mathcal{M} \models \varphi(b, b_1, \dots, b_n)$, showing $\mathcal{M} \models (\exists x\varphi)(\bar{b})$.

Case 2. a is not algebraic over R.

Since $\bar{a}$ and $\bar{b}$ satisfy the same atomic formulas, the rings generated by $a_1, \dots, a_n$ and by $b_1, \dots, b_n$ are isomorphic. The fields generated by these rings are just the quotient fields of the respective rings and therefore are isomorphic, too. It follows that the fields generated by $a_1, \dots, a_n$ and by $b_1, \dots, b_n$ have the same (finite) transcendence degree over their prime subfield, which is also the prime subfield of $\mathcal{M}$.

Since a is transcendent over the field generated by $a_1, \dots, a_n$, the transcendence degree of $\mathcal{M}$ over the prime subfield is at least one greater than the transcendence degree of the field generated by $a_1, \dots, a_n$ and hence at least one greater than the transcendence degree of the field generated by $b_1, \dots, b_n$. It follows that there is some $b \in M$ that is transcendent over the field generated by $b_1, \dots, b_n$.

Now there is an isomorphism from the field generated by R and a to the field generated by S and b that maps a to b and each a_i to the corresponding b_i. It follows that $(a, a_1, \dots, a_n)$ and $(b, b_1, \dots, b_n)$ satisfy the same atomic formulas. Since φ is quantifier-free, $\mathcal{M} \models \varphi(b, b_1, \dots, b_n)$ and hence $\mathcal{M} \models (\exists x\varphi)(\bar{b})$. □

Definition 2.31. A theory Φ over τ is **model-complete** if for all models $\mathcal{M}$ of Φ and every substructure $\mathcal{N}$ of $\mathcal{M}$ that is a model of Φ, $\mathcal{N}$ is an elementary substructure of $\mathcal{M}$.

Lemma 2.32. *Let Φ be a complete theory with quantifier elimination. Then Φ is model-complete.*

Proof. Let $\mathcal{M}$ be a model of Φ and let $\mathcal{N}$ be a substructure of $\mathcal{M}$. By the Tarski-Vaught Criterion, it is enough to check that for every formula formula $\varphi(x, y_1, \dots, y_n)$ and parameters $b_1, \dots, b_n \in N$, if there is $a \in M$ such that $\mathcal{M} \models \varphi(a, b_1, \dots, b_n)$, then there is $b \in N$ such that $\mathcal{M} \models \varphi(b, b_1, \ldots, b_n)$. Since Φ has quantifier elimination, we may assume that φ is quantifier-free.

Since Φ has quantifier-elimination, $(\exists x\varphi)(y_1, \ldots, y_n)$ is Φ-equivalent to a quantifier-free formula $\psi(y_1, \ldots, y_n)$. Since $\mathcal{N}$ is a substructure of $\mathcal{M}$, $(b_1, \ldots, b_n)$ satisfies the same atomic formulas, and therefore the same quantifier-free formulas, in $\mathcal{N}$ and in $\mathcal{M}$. It follows that $\mathcal{N} \models \psi(b_1, \ldots, b_n)$ and hence $\mathcal{N} \models (\exists\varphi x)(b_1, \ldots, b_n)$. It follows that there is $b \in N$ such that $\mathbb{N} \models \varphi(b, b_1, \ldots, b_n)$. Since $\varphi(x, y_1, \ldots, y_n)$ is quantifier-free, we have $\mathcal{M} \models \varphi(b, b_1, \ldots, b_n)$. $\square$

Corollary 2.33. *For $p = 0$ or p a prime number, Φ_{ACF_p} is model-complete.*

The model-completeness of Φ_{ACF_p} has an important consequence for algebraic geometry, Hilbert's Nullstellensatz. Before we can prove the Nullstellensatz, we need a basic fact about rings of polynomials.

Theorem 2.34 (Hilbert's Basissatz)**.** *Let K be a field and consider the polynomial ring $K[X_1, \ldots, X_n]$ in n variables over K. Then every ideal of $K[X_1, \ldots, X_n]$ is generated by finitely many elements.*

Proof. Observe that $K[X_1, \ldots, X_{n+1}] = K[X_1, \ldots X_n][X_{n+1}]$. A commutative ring is **Noetherian** if every ideal is generated by finitely many elements. We show by induction that $K[X_1, \ldots, X_n]$ is Noetherian for every n. It is enough to show that for every commutative Noetherian ring R, the polynomial ring $R[X]$ is Noetherian.

So, assume that R is commutative and Noetherian. Let I be an ideal of $R[X]$. Suppose that I is not fintely generated. Then we can recursively construct a sequence $(f_i)_{i\in\mathbb{N}}$ of polynomials in I such that for every $n \in \mathbb{N}$, f_n is not in the ideal generated by $\{f_i : i < n\}$ and such that f_n is of minimal degree among all the polynomials in I that are not in the ideal generated by $\{f_i : i < n\}$. Note that the sequence of degrees of the f_n is non-decreasing.

For each $i \in \mathbb{N}$ let a_i be the leading coefficient of f_i. Let $n \in \omega$ be minimal such that the ideal generated by $\{a_i : i < n\}$ is equal to the ideal generated by $\{a_i : i \leq n\}$. Such an n exists since R is Noetherian. Then there are $u_0, \ldots, u_{n-1} \in R$ such that $a_n = u_0a_0 + \cdots + u_{n-1}a_{n-1}$.

Now consider the polynomial

$$g = u_0 f_0 X^{k_0} + \cdots + u_{n-1} f_{n-1} X^{k_{n-1}}$$

where $k_i = \deg(f_n) - \deg(f_i)$. Since $(\deg(f_i))_{i\in\mathbb{N}}$ is non-decreasing, each k_i is ≥ 0 and therefore g is indeed a polynomial. The degree of g is the same as the degree of f_n. Also, the leading coefficient of g is a_n. It follows that the degree of $f_n - g$ is strictly smaller than the degree of f_n. Also, since $f_n = g + (f_n - g)$, $f_n - g$ is in I but not in the ideal generated by $\{f_i : i < n\}$, contradicting the choice of f_n. $\square$

Theorem 2.35 (Hilbert's Nullstellensatz). *Let F be an algebraically closed field. If I is an ideal in $F[X_1, \dots, X_n]$ such that $1 \notin I$, then there is $(a_1, \dots, a_n) \in F^n$ such that for all $f \in I$, $f(a_1, \dots, a_n) = 0$.*

Proof. By the Basissatz, there are polynomials $f_1, \dots, f_k \in I$ such that I is the ideal generated by $f_1, \dots, f_k$. By Zorn's Lemma there is a maximal ideal $P \subseteq F[X_1, \dots, X_n]$ such that $I \subseteq P$ and $1 \notin P$. Since P is a maximal ideal, $G = F[X_1, \dots, X_n]/P$ is a field. We can consider F as a subfield of G via the embedding $a \mapsto a+P$. For each $i \in \{1, \dots, k\}$,

$$f_i(X_1 + P, \dots, X_n + P) = f_i(X_1, \dots, X_n) + P = 0 + P.$$

Let H be the algebraic closure of G. Now

$$H \models \exists x_1 \dots \exists x_n (f_1(x_1, \dots, x_n) = 0 \wedge \dots \wedge f_k(x_1, \dots, x_n) = 0).$$

Since F is an algebraically closed field and a substructure of H, we have $F \preccurlyeq H$. Therefore

$$F \models \exists x_1 \dots \exists x_n (f_1(x_1, \dots, x_n) = 0 \wedge \dots \wedge f_k(x_1, \dots, x_n) = 0).$$

Hence, the $f_1, \dots, f_k$ have a common zero $(a_1, \dots, a_n) \in F^n$. Since I is generated by $f_1, \dots, f_n$, for all $f \in I$ we have $f(a_1, \dots, a_n) = 0$. □

Exercise 2.36. Let $\mathcal{M}$ be a τ-structure for some vocabulary τ. A set $S \subseteq M^n$ is **definable** (with parameters) if there are a τ-formula $\varphi(x_1, \dots, x_n, y_1, \dots, y_m)$ and $b_1, \dots, b_m \in M$ such that

$$S = \{(a_1, \dots, a_n) \in M^n : \mathcal{M} \models \varphi(a_1, \dots, a_n, b_1, \dots, b_m)\}.$$

Let K be an algebraically closed field. The **zero-set** of a polynomial $f(X_1, \dots, X_n) \in K[X_1, \dots, X_n]$ is the set

$$\{(a_1, \dots, a_n) \in K^n : f(a_1, \dots, a_n) = 0\}.$$

Show that the definable subsets of K^n are precisely the Boolean combinations of zero-sets.

The theory of algebraically closed fields of characteristic 0 is the theory of $(\mathbb{C}, 0, 1, +, \cdot)$. Let us have a look at the theory of $(\mathbb{R}, 0, 1, +, \cdot, \leq)$, the **theory of real closed fields**.

Definition 2.37. Consider the vocabulary $\tau = \{0, 1, +, \cdot, \leq\}$. A τ-structure $(R, 0, 1, +, \cdot, \leq)$ is an **ordered ring** if $(R, 0, 1, +, \cdot)$ is a ring (with unit), $\leq$ is a linear order on R, and moreover, the following two axioms are satisfied:

(OR1) $\forall x \forall y \forall z (x \leq y \to x + z \leq y + z)$
(OR2) $\forall x \forall y \forall z ((x \leq y \wedge 0 \leq z) \to xz \leq yz)$

An **ordered field** is an ordered ring $(F, 0, 1, +, \cdot, \leq)$ such that $(F, 0, 1, +, \cdot)$ is a field.

A field $(F, 0, 1, +, \cdot)$ is **real closed** if there is an order $\leq$ on F which turns F into an ordered field such that every polynomial of odd degree with coefficients in F has at least one zero in F and such that every element ≥ 0 has a square root. Let Φ_{RCF} be the theory of real closed fields, i.e., the deductive closure of the axioms for real closed fields.

It is clear that $(\mathbb{R}, 0, 1, +, \cdot, \leq)$ is a real closed field. However, it is not clear that the theory of real closed fields is complete. In other words, there could be real closed fields that are not elementarily equivalent to the ordered field of real numbers. However, this is not the case. Since the theory of real closed fields is not κ-categorical for any κ, our previous strategy for showing the completeness of a first order theory fails. Another way of showing completeness is to show quantifier elimination first.

Lemma 2.38. *If Φ_{RCF} has elimination of quantifiers, then it is complete.*

Proof. Assume Φ_{RCF} has elimination of quantifiers. Let φ be a sentence. The only constants in the vocabulary are 0 and 1. If φ is equivalent to a quantifier-free formula, then it is equivalent to a Boolean combination of atomic formulas without variables. But Φ_{RCF} decides the validity of every atomic formula without variables. Hence Φ_{RCF} is complete. □

In order to prove quantifier elimination of the theory of real closed fields, we use a version of Lemma 2.27 that works for theories that have not been shown to be complete yet.

Lemma 2.39. *Let τ be any vocabulary and let Φ be a theory over τ. Then the following are equivalent:*

(1) *Φ has quantifier elimination.*
(2) *For every τ-structure $\mathcal{C}$, if $f : C \to M$ and $g : C \to N$ are embeddings of $\mathcal{C}$ into a models $\mathcal{M}$ and $\mathcal{N}$ of Φ, then for every quantifier-free formula τ-formula $\varphi(x, y_1, \dots, y_n)$ and every n-tuple $(c_1, \dots, c_n) \in C^n$ we have*

$$\mathcal{M} \models (\exists x\varphi)(f(c_1), \dots, f(c_n)) \quad \Leftrightarrow \quad \mathcal{N} \models (\exists x\varphi)(g(c_1), \dots, g(c_n)).$$

Proof. (1)⇒(2): This is practically the same as the proof of (1)⇒(2) in Lemma 2.27.

(2)⇒(1): Let $\varphi(x, y_1, \dots, y_n)$ be a quantifier-free τ-formula. Consider the vocabulary $\sigma = \tau \cup \{c_1, \dots, c_n\}$ where $c_1, \dots, c_n$ are new constant symbols. If $\mathcal{M}$ and $\mathcal{N}$ are two σ-structures that are models of Φ and satisfy the same quantifier-free σ-sentences, then the substructures generated by $c_1^{\mathcal{M}}, \dots, c_n^{\mathcal{M}}$ and $c_1^{\mathcal{N}}, \dots, c_n^{\mathcal{N}}$ are isomorphic and hence, by (2),

$$\mathcal{M} \models (\exists x\varphi)(c_1^{\mathcal{M}}, \dots, c_n^{\mathcal{M}}) \quad \Leftrightarrow \quad \mathcal{N} \models (\exists x\varphi)(c_1^{\mathcal{N}}, \dots, c_n^{\mathcal{N}}).$$

If no σ-structure that is a model of Φ is a model $(\exists x\varphi)(c_1, \dots, c_n)$, then $(\exists x\varphi)(y_1, \dots, y_n)$ is equivalent to a (quantifier-free) formula that is always false. Now, let $\mathcal{M}$ be a σ-structure that is a model of Φ and

$(\exists x\varphi)(c_1,\dots,c_n)$. Let $\Psi(\mathcal{M})$ be the set of quantifier-free σ-sentences that hold in $\mathcal{M}$. Since every model of $\Phi \cup \Psi(\mathcal{M})$ is also a model of $(\exists x\varphi)(c_1,\dots,c_n)$ by (2), $\Phi \cup \Psi(\mathcal{M}) \models (\exists x\varphi)(c_1,\dots,c_n)$. It follows that Φ together with a finite subset of $\Psi(\mathcal{M})$ implies $(\exists x\varphi)(c_1,\dots,c_n)$. Since $\Psi(\mathcal{M})$ is closed under finite conjunctions, Φ together with a single sentence $\psi_{\mathcal{M}} \in \Psi(\mathcal{M})$ implies $(\exists x\varphi)(c_1,\dots,c_n)$.

Now consider the theory

$$\Phi \cup \{(\exists x\varphi)(c_1,\dots,c_n)\} \cup \{\neg\psi_{\mathcal{M}} : \mathcal{M} \models (\exists x\varphi)(c_1,\dots,c_n)\}.$$

By the choice of the $\psi_{\mathcal{M}}$, this theory is inconsistent. Hence there are

$$\psi_1,\dots,\psi_k \in \{\neg\psi_{\mathcal{M}} : \mathcal{M} \models (\exists x\varphi)(c_1,\dots,c_n)\}$$

such that

$$\Phi \cup \{\neg\psi_1 \wedge \dots \wedge \neg\psi_k\} \models \neg(\exists x\varphi)(c_1,\dots,c_n).$$

It follows that $(\exists x\varphi)(c_1,\dots,c_n)$ is Φ-equivalent to $\psi_1 \vee \dots \vee \psi_k$. Hence $(\exists x\varphi)(y_1,\dots,y_n)$ is equivalent to a quantifier-free τ-formula. □

Exercise 2.40. In the proof of the previous lemma we implicitly used (and proved) the **separation lemma**:

Let τ be a vocabulary. Let Ψ be a collection of τ-sentences that is closed under conjunction and disjunction and contains $\bot$, the sentence that is always false, and $\top$, the sentence that is always true. Let Φ_0 and Φ_1 be theories over τ. Suppose for every model $\mathcal{M}$ of Φ_0 and every model $\mathcal{N}$ of Φ_1 there is $\psi \in \Psi$ such that $\mathcal{M} \models \psi$ and $\mathcal{N} \models \neg\psi$.

Show that there is $\psi \in \Psi$ such that $\Phi_0 \models \psi$ and $\Phi_1 \models \neg\psi$.

We collect some facts about real closed fields.

Definition 2.41. Let $(F,0,1,+,\cdot,\leq)$ be an ordered field. Then an ordered field $(K,0,1,+,\cdot,\leq)$ is the **real closure** of F if F is a substructure of K, K is a real closed field, and every element of K is algebraic over F.

Lemma 2.42. *a) If $(R,0,1,+,\cdot,\leq)$ is an ordered ring, then the order on R extends to the field of fractions of R, turning the field of fractions into an ordered field.*

b) A field $(F,0,1,+,\cdot)$ is **formally real** *if -1 is not the sum of squares. If F is formally real and $a \in F$, then F has an order $\leq$ with $a < 0$ iff a is not a sum of squares, i.e., not of the form $b_1^2 + \dots + b_n^2$ for some $b_1,\dots,b_n \in F$.*

b) If $(F,0,1,+,\cdot)$ is a real closed field, then the order witnessing this is unique and can be defined by letting $a \leq b$ if there is $c \in F$ such that $b = a + c^2$.

c) Every ordered field has a real closure. The real closure is unique up to isomorphism over the base field. (The **Artin-Schreier Theorem***.)*

d) If $(R, 0, 1, +, \cdot, \leq)$ *is a real closed field, then adjoining a zero of the polynomial* X^2+1 *to* $(R, 0, 1, +, \cdot)$ *yields an algebraically closed field* $R[\sqrt{-1}]$.

e) If $(R, 0, 1, +, \cdot, \leq)$ *is a real closed field, then every polynomial in* $R[X]$ *is the product of linear and quadratic factors of the form* $(X - d)^2 + e$ *with* $e > 0$.

Theorem 2.43. *The theory of real closed fields has quantifier elimination and is complete.*

Proof. Completeness follows from quantifier elimination by Lemma 2.38. We use Lemma 2.39 to show quantifier elimination. Let $\mathcal{C}$ be an ordered ring and let $f : C \to M$ and $g : C \to N$ be embeddings into real closed fields $\mathcal{M}$ and $\mathcal{N}$. Let $\varphi(x, y_1, \dots, y_n)$ be a quantifier-free formula and $c_1, \dots, c_n \in C$. Let

$$\overline{a} = (a_1, \dots, a_n) = (f(c_1), \dots, f(c_n))$$

and

$$\overline{b} = (b_1, \dots, b_n) = (g(c_1), \dots, g(c_n)).$$

Assume

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

Choose $a \in M$ such that $\mathcal{M} \models \varphi(a, \overline{a})$. Let F_0 be the real closed subfield of $\mathcal{M}$ generated by $a_1, \dots, a_n$ and let F_1 be the real closed subfield of $\mathcal{N}$ generated by $b_1, \dots, b_n$. Clearly, F_0 is isomorphic to F_1 via an isomorphism h that maps $f(c)$ to $g(c)$ for every $c \in C$..

If $a \in F_0$, then $\mathcal{N} \models \varphi(h(a), \overline{b})$. If $a \notin F_0$, then a is transcendent over F_0 and $F_0(a)$ is isomorphic to the field $F_0(X)$ of fractions of $F_0[X]$. Let F_0^ℓ denote the set of elements of F_0 that are less than a and let F_0^r denote the set of elements of F_0 that are greater than a. Let $F_1^\ell = h[F_0^\ell]$ and $F_1^r = h[F_0^r]$. Since ordered fields are dense linear orders, there is an elementary extension of $\mathcal{N}'$ of $\mathcal{N}$ that contains an element b which is greater than all elements of F_1^ℓ and smaller than all elements of F_1^r. Since F_1 is real closed an $b \notin F_1$, b is transcendent over F_1. Hence, algebraically h has a unique extension $\overline{h} : F_0(a) \to F_1(b)$ that maps a to b. We have to show that $\overline{h}$ is order-preserving.

It is enough to show that $\overline{h}$ is order-preserving on $F_0[a]$. But for this it is actually enough to show that $\overline{h}$ preserves positivity on $F_0[a]$. Let $p(X) \in F_0[X]$. Since F_0 is real closed, $p(X)$ factors as

$$p(X) = \alpha \cdot \prod_{i<n}(X - c_i) \cdot \prod_{j<m}((X - d_j)^2 + e_j),$$

with positive e_j's. Now, whether or not $p(a)$ is positive depends on how many of the factors α, $a - c_0$, ... , $a - c_{n-1}$ are negative. Whether or not $\overline{h}(p(a))$ is positive depends in the same way on the factors $\overline{h}(\alpha)$, $\overline{h}(a) - \overline{h}(c_0)$, ... , $\overline{h}(a) - \overline{h}(c_{n-1})$. But since $\overline{h}$ preserves the order on

F_0 and, by the choice of b, also between the elements of F_0 and a, it follows that $p(a)$ is positive iff $\overline{h}(p(a))$ is.

Since φ is quantifier-free, b witnesses $(\exists x\varphi)(\overline{b})$ in $\mathcal{N}'$. Since $\mathcal{N}'$ is an elementary extension of $\mathcal{N}$, $\mathcal{N} \models (\exists x\varphi)(\overline{b})$. This shows condition (2) in Lemma 2.39. Hence the theory of real closed fields has quantifier elimination. □

Since the theory of real closed fields has quantifier elimination, it is model-complete. We use this fact to prove a theorem that settles Hilbert's 17th problem.

Theorem 2.44. *Let K be a real closed field. A polynomial f in the polynomial ring $K[X_1, \dots, X_n]$ is a sum of squares*

$$f = g_1^2 + \cdots + g_k^2$$

of rational functions $g_1, \dots, g_k \in K(X_1, \dots, X_k)$ iff for all $a_1, \dots, a_n \in K$,

$$f(a_1, \dots, a_n) \geq 0.$$

Proof. Assume that

$$f = \frac{g_1^2}{h_1^2} + \cdots + \frac{g_k^2}{h_k^2}$$

for polynomials $g_1, \dots, g_k, h_1, \dots, h_k \in K[X_1, \dots, X_n]$, where the h_j are not constantly 0. Let D be the set of all $\overline{a} \in K^n$ such that for all $j \in \{1, \dots, k\}$, $h_j(\overline{a}) \neq 0$. The set D is dense and open in K^n. For each $\overline{a} \in D$ we have $f(\overline{a}) \geq 0$. The set of all $\overline{a} \in K^n$ with $f(\overline{a}) \geq 0$ is closed and includes D. It follows that $f(\overline{a}) \geq 0$ for all $\overline{a} \in K^n$.

Now suppose that f is not a sum of squares. K is an ordered field and -1 is negative with respect to the order on K. In particular, -1 is not a sum of squares in K. By the argument at the beginning of the proof, -1 is not a sum of squares in $K(X_1, \dots, X_n)$ either. It follows that $K(X_1, \dots, X_n)$ is an ordered field.

Since f is not a sum of squares in this field, there is an order with respect to which f is negative. Since the order of K is unique, the order on $K(X_1, \dots, X_n)$ extends the order on K. Let L be the real closure of $K(X_1, \dots, X_n)$. Since $f(X_1, \dots, X_n) < 0$ in L, we have

$$L \models \exists x_1 \dots \exists x_n (f(x_1, \dots, x_n) < 0).$$

Since the theory of real closed fields is model-complete, $K \preccurlyeq L$ and therefore

$$K \models \exists x_1 \dots \exists x_n (f(x_1, \dots, x_n) < 0).$$

It follows that there are $a_1, \dots, a_n \in K$ such that $f(a_1, \dots, a_n) < 0$. □

2.3. **Strongly minimal theories.**

Definition 2.45. Let τ be a vocabulary and $\mathcal{M}$ a τ-structure. Let B be a subset of M. A set $A \subseteq M^n$ is **definable over** B if there are a τ-formula $\varphi(x_1, \dots, x_n, y_1, \dots, y_m)$ and parameters $b_1, \dots, b_m \in B$ such that

$$A = \{(a_1, \dots, a_n) \in M^n : \mathcal{M} \models \varphi(a_1, \dots, a_n, b_1, \dots, b_m)\}.$$

A subset A of M^n is **definable** if it is definable over M. A is **definable without parameters** if A is definable over the empty set.

Lemma 2.46. *If $A \subseteq M^n$ is finite or cofinite, then A is definable.*

Proof. If A is finite, say

$$A = \{(a_1^i, \dots, a_n^i) : i < m\},$$

then A is defined by the formula

$$\bigvee_{i<m} (x_1 = y_1^i \wedge \dots \wedge x_n = y_n^i)$$

using the parameters $a_1^i, \dots, a_n^i$, $i < m$.

Clearly, the complement of every definable set is definable. □

Definition 2.47. An infinite structure $\mathcal{M}$ is **minimal** if all definable subsets of M are finite or cofinite. $\mathcal{M}$ is **strongly minimal** if every structure $\mathcal{N}$ that is elementarily equivalent to $\mathcal{M}$ is minimal. A theory is **strongly minimal** if all of its models are infinite and strongly minimal.

Observe that in the definition of minimality only definable subsets of M, and not of M^n are considered. The reason for this is that for every infinite structure $\mathcal{M}$ the set $\{(a, a) : a \in M\}$ is definable and both infinite and coinfinite.

Example 2.48. a) The structure $(\mathbb{Q}, \leq)$ is not minimal since for every $q \in \mathbb{Q}$ the set $\{a \in \mathbb{Q} : a \leq q\}$ is definable and both infinite and coinfinite.

b) The random graph is not minimal since for every vertex the set of neighbors is definable and both infinite and coinfinite.

Lemma 2.49. *Let Φ be a theory over τ that has quantifier elimination and only infinite models. Then Φ is strongly minimal iff for every model $\mathcal{M}$ of Φ the subsets of M defined by atomic formulas are finite or cofinite.*

Proof. Clearly, if Φ is strongly minimal, then for every model $\mathcal{M}$ of Φ, every subset of M definable by an atomic formula is finite or cofinite.

Now let $\mathcal{M}$ be a model of Φ and assume that every subsets of M defined by an atomic formulas is finite or cofinite. Since Φ has quantifier elimination, every definable subset of M is definable by a quantifier-free formula. Every quantifier-free formula is a Boolean combination of atomic formulas. A straight-forward induction shows that if every set definable by an atomic formula is finite or cofinite, then every set definable by a Boolean combination of atomic formulas if finite or cofinite. This shows that $\mathcal{M}$ is minimal. $\square$

Corollary 2.50. *The theory of algebraically closed fields of characteristic p, $p = 0$ or p a prime number, is strongly minimal.*

Proof. The theory has quantifier elimination and hence we can apply Lemma 4.19. Let $\mathcal{M}$ be an algebraically closed field of characteristic p. The atomic formulas in the language of fields are equations between polynomials. Every such equation in the variables $x, y_1, \dots, y_m$ is equivalent to an equation of the form $p(x, y_1, \dots, y_m) = 0$. If $b_1, \dots, b_m$ are parameters from M, then the equation $p(x, b_1, \dots, b_m)$ has only finitely many solutions $a \in M$. It follows that the corresponding definable subset of M is finite. $\square$

Definition 2.51. Let $\mathcal{M}$ be a τ-structure and $\varphi(x, y_1, \dots, y_n)$ a τ-formula. Let $b_1, \dots, b_n \in M$. Then $\varphi(x, b_1, \dots, b_n)$ is **algebraic** if $\{a \in M : \mathcal{M} \models \varphi(a, b_1, \dots, b_n)\}$ is finite. For $B \subseteq M$ and $a \in M$, a is **algebraic over** B if there are $b_1, \dots, b_n \in B$ and a formula $\varphi(x, y_1, \dots, y_n)$ such that $\varphi(x, b_1, \dots, b_n)$ is algebraic and $\mathcal{M} \models \varphi(a, b_1, \dots, b_n)$.

For $B \subseteq M$ let $\mathrm{acl}_{\mathcal{M}}(B)$ be the set of all $a \in M$ that are algebraic over B. The set $\mathrm{acl}_{\mathcal{M}}(B)$ is the **algebraic closure** of B. $A \subseteq M$ is **algebraically closed** if $A = \mathrm{acl}_{\mathcal{M}}(A)$.

Lemma 2.52. *a) Let $\mathcal{M}$ be a τ-structure and $A, B \subseteq M$. Then the following hold:*

(1) $A \subseteq \mathrm{acl}_{\mathcal{M}}(A)$ *(Reflexivity)*
(2) *If $A \subseteq B$, then $\mathrm{acl}_{\mathcal{M}}(A) \subseteq \mathrm{acl}_{\mathcal{M}}(B)$. (Monotonicity)*
(3) $\mathrm{acl}_{\mathcal{M}}(\mathrm{acl}_{\mathcal{M}}(A)) = \mathrm{acl}_{\mathcal{M}}(A)$ *(Idempotency)*
(4) *If $a \in \mathrm{acl}_{\mathcal{M}}(A)$, then $a \in \mathrm{acl}_{\mathcal{M}}(A_0)$ for some finite set $A_0 \subseteq A$. (Finite character)*
(5) $\mathrm{acl}_{\mathcal{M}}(A)$ *is the underlying set of a substructure of $\mathcal{M}$.*

c) If $\mathcal{M}$ is minimal, then for all $A \subseteq M$ and all $b, c \in M$ the follwing holds:

(6) *If $c \in \mathrm{acl}_{\mathcal{M}}(A \cup \{b\})$ and $c \notin \mathrm{acl}_{\mathcal{M}}(A)$, then $b \in \mathrm{acl}_{\mathcal{M}}(A \cup \{c\})$. (Exchange)*

Proof. (1) For every $a \in A$, $\{a\}$ is defined by $x = a$ and $x = a$ is algebraic.

(2) Every set definable over A is definable over B. If a is algebraic over A, then it is an element of a finite set that is definable over A. This finite set is definable over B and thus a is algebraic over B.

(3) By (1) and (2) we have $\mathrm{acl}_{\mathcal{M}}(A) \subseteq \mathrm{acl}_{\mathcal{M}}(\mathrm{acl}_{\mathcal{M}}(A))$. Now assume that a is algebraic over $\mathrm{acl}_{\mathcal{M}}(A)$. Then there are a formula

$\varphi(x, y_1, \dots, y_n)$ and $b_1, \dots, b_n \in \mathrm{acl}_{\mathcal{M}}(A)$ such that $\varphi(x, b_1, \dots, b_n)$ is algebraic and $\mathcal{M} \models \varphi(a, b_1, \dots, b_n)$. Let $\ell \in \mathbb{N}$ be the number of distinct x in M that satisfy $\varphi(x, b_1, \dots, b_n)$. For each $i \in \{1, \dots, n\}$ choose $\psi_i(y, z_1, \dots, z_{m_i})$ and $a_1^i, \dots, a_{m_i}^i \in A$ such that $\psi_i(y, a_1^i, \dots, a_{m_i}^i)$ is algebraic and $\mathcal{M} \models \psi_i(b_i, a_1^i, \dots, a_{m_i}^i)$.

For every formula χ let $\exists^{=\ell} y \chi$ be the formula that expresses "there are exactly ℓ distinct y that satisfy χ". Now

$$\exists y_1 \dots \exists y_n \Bigg(\varphi(x, y_1, \dots, y_n) \wedge \exists^{=\ell} z \varphi(z, y_1, \dots, y_n) \wedge \bigwedge_{i=1}^{n} \psi(y_i, a_1^i, \dots, a_{m_i}^i) \Bigg)$$

is algebraic and is satisfied in $\mathcal{M}$ by a. Hence $a \in \mathrm{acl}_{\mathcal{M}}(A)$.

(4) This is obvious.

(5) For every constant symbol c, $\{c^{\mathcal{M}}\}$ is definable without parameters. Hence $c^{\mathcal{M}} \in \mathrm{acl}_{\mathcal{M}}(A)$. If f is a n-ary function symbol and $b_1, \dots, b_n \in \mathrm{acl}_{\mathcal{M}}(A)$, then $\{f^{\mathcal{M}}(b_1, \dots, b_n)\}$ is definable over $\mathrm{acl}_{\mathcal{M}}(A)$ and therefore

$$f^{\mathcal{M}}(b_1, \dots, b_n) \in \mathrm{acl}_{\mathcal{M}}(\mathrm{acl}_{\mathcal{M}}(A)) = \mathrm{acl}_{\mathcal{M}}(A).$$

It follows that $\mathrm{acl}_{\mathcal{M}}(A)$ is closed under all the functions of $\mathcal{M}$. This shows (5).

(6) Let $c \in \mathrm{acl}_{\mathcal{M}}(A \cup \{b\})$. Let $\varphi(x, y, z_1, \dots, z_n)$ be a formula and $a_1, \dots, a_n \in A$ such that $\varphi(x, b, a_1, \dots, a_n)$ is algebraic and $\mathcal{M} \models \varphi(c, b, a_1, \dots, a_n)$. Let $k \in \mathbb{N}$ be such that $\varphi(x, b, a_1, \dots, a_n)$ is satisfied by exactly k distinct x in M. If $\varphi(c, y, a_1, \dots, a_n)$ is algebraic, then $b \in \mathrm{acl}_{\mathcal{M}}(A \cup \{c\})$.

Now assume that $\varphi(c, y, a_1, \dots, a_n)$ is not algebraic. Then by the minimality of $\mathcal{M}$, the set

$$\{d \in M : \mathcal{M} \models \varphi(c, d, a_1, \dots, a_n)\}$$

is cofinite, i.e., $\varphi(c, y, a_1, \dots, a_n)$ holds for almost all y in M.

Hence, there is $\ell \in \mathbb{N}$ such that $\neg\varphi(c, y, a_1, \dots, a_n)$ holds for exactly ℓ distinct y. Since $c \notin \mathrm{acl}_{\mathcal{M}}(A)$, $\exists^{=\ell} y \neg\varphi(x, y, a_1, \dots, a_n)$ holds for almost all x in M. But this means that for almost all $d \in M$, $\varphi(d, y, a_1, \dots, a_n)$ is satisfied by all but at most ℓ distinct y in M.

Suppose that there are distinct $e_1, \dots, e_{\ell+1} \in M$ such that for each $i \in \{1, \dots, \ell + 1\}$, $\varphi(x, e_i, a_1, \dots, a_n)$ fails for almost all x in M. Since finite intersections of cofinite sets are again cofinite, we can conclude that for almost all x, $\varphi(x, e_i, a_1, \dots, a_n)$ fails for all $i \in \{1, \dots, \ell + 1\}$, contradicting the previous statement.

Hence, for all but at most ℓ distinct $e \in M$, $\varphi(x, e, a_1, \dots, a_n)$ is satisfied by almost all x in M.

In particular, $\exists^{=k}x\varphi(x, y, a_1, \dots, a_n)$ fails for almost all y. But this shows that $\exists^{=k}x\varphi(x, y, a_1, \dots, a_n)$ is algebraic and therefore

$$b \in \mathrm{acl}_{\mathcal{M}}(A) \subseteq \mathrm{acl}_{\mathcal{M}}(A \cup \{c\}).$$

□

Definition 2.53. Let $\mathcal{M}$ be a structure and $A, C \subseteq M$. Then A is **independent over** C if for every $a \in A$, $a \notin \mathrm{acl}_{\mathcal{M}}((A \cup C) \setminus \{a\})$. A is **independent** if A is independent over $\emptyset$.

$B \subseteq A$ **generates** A **over** C iff $A \subseteq \mathrm{acl}_{\mathcal{M}}(B \cup C)$. Note that this is equivalent to $\mathrm{acl}_{\mathcal{M}}(A \cup C) = \mathrm{acl}_{\mathcal{M}}(B \cup C)$. B **generates** A iff it generates A over $\emptyset$.

$B \subseteq A$ is a **basis for** A **over** C iff B is independent and generates A over C. B is a **basis for** A if B is a basis for A over $\emptyset$.

Exercise 2.54. Show that the theory of infinite vector spaces over a fixed field F has quantifier elimination.

Exercise 2.55. Show that the theory of infinite vector spaces over a fixed field F is strongly minimal. Note that the model theoretic notions "independent" and "basis" coincide with their counter parts from linear algebra.

We will now show that minimal structures allow the definition of a dimension of sets.

Lemma 2.56. *Let $\mathcal{M}$ be a minimal structure, let $C \subseteq M$ and let $E \subseteq M$ be finite and independent over C. If $F \subseteq M$ is such that $|E| = |F|$ and $E \subseteq \mathrm{acl}_{\mathcal{M}}(F \cup C)$, then $F \subseteq \mathrm{acl}_{\mathcal{M}}(E \cup C)$.*

Proof. We prove the lemma by induction on $n = |E| = |F|$. Let us start with the following claim:

Claim 2.57. If the lemma holds for n, then whenever $E \subseteq M$ is finite and independent over C and $F \subseteq M$ is finite and such that $E \subseteq \mathrm{acl}_{\mathcal{M}}(F \cup C)$ and $|F| = n$, then $|E| \leq n$.

For the proof of the claim suppose that $|E| > n$. Let $E' \subseteq E$ be of size n. Clearly, $E' \subseteq \mathrm{acl}_{\mathcal{M}}(F \cup C)$. By the lemma for n, $F \subseteq \mathrm{acl}_{\mathcal{M}}(E' \cup C)$. It follows that $E \subseteq \mathrm{acl}_{\mathcal{M}}(E' \cup C)$, contradicting the independence of E. This shows the claim.

Let us return to the proof of the lemma itself. First assume that $n = 1$. Let f be the unique element of F and let e be the unique element of E. Since $E = \{e\}$ is independent over C, $e \notin \mathrm{acl}_{\mathcal{M}}(C)$. By (6) of Lemma 2.52 (exchange), $f \in \mathrm{acl}_{\mathcal{M}}(\{e\} \cup C)$, i.e., $F \subseteq \mathrm{acl}_{\mathcal{M}}(E \cup C)$

Now let $n > 1$ and assume the lemma is true for sets of size less than n. Let $E \subseteq M$ be independent over C and $F \subseteq M$ such that $n = |E| = |F|$ and $E \subseteq \mathrm{acl}_{\mathcal{M}}(F \cup C)$. We write $F = \{f_1, \dots, f_n\}$.

By the induction hypothesis together with the claim,

$$E \not\subseteq \mathrm{acl}_{\mathcal{M}}(\{f_2, \dots, f_n\} \cup C).$$

Hence, there is $e_1 \in E \setminus \mathrm{acl}_{\mathcal{M}}(\{f_2, \dots, f_n\}) \cup C)$. By exchange,

$$f_1 \in \mathrm{acl}_{\mathcal{M}}(\{e_1, f_2, \dots, f_n\} \cup C).$$

Now suppose that for some k we have

$$\{f_1, \dots, f_k\} \subseteq \mathrm{acl}_{\mathcal{M}}(\{e_1, \dots, e_k, f_{k+1}, \dots, f_n\} \cup C)$$

with pairwise distinct $e_1, \dots, e_k \in E$. By the induction hypothesis together with the claim,

$$E \not\subseteq \mathrm{acl}_{\mathcal{M}}((\{e_1, \dots, e_k, f_{k+1}, \dots, f_n\} \setminus \{f_{k+1}\}) \cup C).$$

Let

$$e_{k+1} \in E \setminus \mathrm{acl}_{\mathcal{M}}((\{e_1, \dots, e_k, f_{k+1}, \dots, f_n\} \setminus \{f_{k+1}\}) \cup C).$$

Again by exchange,

$$f_{k+1} \in \mathrm{acl}_{\mathcal{M}}((\{e_1, \dots, e_{k+1}, f_{k+1} \dots, f_n\} \setminus \{f_{k+1}\}) \cup C).$$

This process terminates after n steps and we have

$$f_1, \dots, f_n \in \mathrm{acl}_{\mathcal{M}}(\{e_1, \dots, e_n\} \cup C),$$

finishing the proof of the lemma. □

Corollary 2.58. *Let $\mathcal{M}$ be a minimal structure and $A, C \subseteq M$. If A is generated over C by a finite set, then A has a finite basis over C and any two bases of A over C are of the same size.*

Proof. Suppose A is generated over C by a set of size n. By Claim 2.57, every set $B \subseteq A$ that is independent over C is of size at most n. Let $B \subseteq A$ be maximally independent over C. We show that B generates A over C.

Let $a \in A$. By the choice of B, $B \cup \{a\}$ is not independent over C. Hence we have $a \in \mathrm{acl}_{\mathcal{M}}(B \cup C)$ or for some $b \in B$, $b \in \mathrm{acl}_{\mathcal{M}}((B \setminus \{b\}) \cup \{a\} \cup C)$. In the latter case, by exchange, $a \in \mathrm{acl}_{\mathcal{M}}(B \cup C)$. It follows that $A \subseteq \mathrm{acl}_{\mathcal{M}}(B \cup C)$. Hence B is a basis for A over C.

Let B' be another basis for A over C. Since B' is independent and $B' \subseteq \mathrm{acl}_{\mathcal{M}}(B \cup C)$, Lemma 2.57 implies that $|B'| \leq |B|$. With the roles of B and B' reversed, the same argument yields $|B| \leq |B'|$. Hence every basis for A over C has the same size as B. □

Lemma 2.59. *Let $\mathcal{M}$ be a minimal structure and $A, C \subseteq M$. Then A has a basis over C and any two bases for A over C have the same size.*

Proof. By Zorn's Lemma, there is a maximal set $B \subseteq A$ that is independent over C. Note that the proof of this fact uses the finite character of our notion of independence (a set is independent iff every finite subset is) and hence of the algebraic closure. By the same argument as in the corollary above, B is a basis for A over C.

If A has a finite basis over C, then the lemma follows from the previous corollary. Now let B and B' be two infinite bases for A over C and assume $|B| < |B'|$. We have $B' \subseteq \mathrm{acl}_{\mathcal{M}}(B \cup C)$. By the finite character of the algebraic closure, for every $b \in B'$ there is a finite set $F_b \subseteq B$ such that $b \in \mathrm{acl}_{\mathcal{M}}(F_b \cup C)$. Since B is infinite, B has $|B|$ finite subsets. Since $|B'| > |B|$, there is a finite set $F \subseteq B$ such that $F_b = F$ for infinitely many $b \in B'$. But now $\mathrm{acl}_{\mathcal{M}}(F \cup C)$ contains an infinite set that is independent over C, contradicting Claim 2.57. $\square$

Exercise 2.60. Let $\mathcal{M}$ be a minimal structure and $A, C \subseteq M$. Show that the following are equivalent for all $B \subseteq A$:

(1) B is a basis for A over C.
(2) B is a maximal independent set over C.
(3) B is a minimal set that generates A over C.

Definition 2.61. Let $\mathcal{M}$ be a minimal structure and $A, C \subseteq M$. The **dimension of** A **over** C is the size $\dim(A/C)$ of a basis of A over C. The **dimension of** A is the size $\dim(A)$ of a basis of A over $\emptyset$.

Lemma 2.62. *Let $\mathcal{M}$ and $\mathcal{N}$ be elementary equivalent minimal structures, and $A \subseteq M$ and $C \subseteq N$. If $\dim(A) = \dim(C)$, then $\mathrm{acl}_{\mathcal{M}}(A) \cong \mathrm{acl}_{\mathcal{N}}(C)$.*

Proof. Let B_A be a basis of A and let B_C be a basis of C. Choose a bijection $f : B_A \to B_C$. We first show that f is elementary in the sense that for all formulas $\varphi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in B_A$,

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

Next, we show that f can be extended to an isomorphism from $\mathrm{acl}_{\mathcal{M}}(A)$ onto $\mathrm{acl}_{\mathcal{N}}(C)$.

Claim 2.63. For every formula $\varphi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in B_A$, if $\mathcal{M} \models \varphi(a_1, \dots, a_n)$, then $\mathcal{N} \models \varphi(f(a_1), \dots, f(a_n))$.

We prove the claim by induction on the number of free variables of φ. If φ has no free variables, i.e., if $n = 0$, then the claim is satisfied since $\mathcal{M}$ and $\mathcal{N}$ satisfy the same sentences.

Now suppose we have proved the claim for n. Consider a formula $\varphi(x_1, \dots, x_{n+1})$ and $a_1, \dots, a_{n+1} \in B_A$. We may assume that the a_i are pairwise distinct. Since B_A is independent, $\varphi(a_1, \dots, a_n, x_{n+1})$ is not algebraic. Since $\mathcal{M}$ is minimal, $\varphi(a_1, \dots, a_n, x_{n+1})$ is satisfied by cofinitely many x_{n+1} in M. In other words, $\neg\varphi(a_1, \dots, a_n, x_{n+1})$ is algebraic. Hence, for some $\ell \in \mathbb{N}$,

$$\mathcal{M} \models \exists^{=\ell} x_{n+1} \neg\varphi(a_1, \dots, a_n, x_{n+1}).$$

By our induction hypothesis,

$$\mathcal{N} \models \exists^{=\ell} x_{n+1} \neg\varphi(f(a_1), \dots, f(a_n), x_{n+1}).$$

It follows that $\neg\varphi(f(a_1), \dots, f(a_n), x_{n+1})$ is algebraic. Since B_C is independent, $f(a_{n+1})$ does not satisfy an algebraic formula with parameters from $B_C \setminus \{f(a_{n+1})\}$. It follows that $\mathcal{N} \models \varphi(f(a_1), \dots, f(a_{n+1}))$. This proves the claim. Hence f is $\mathcal{M}$-elementary.

Claim 2.64. If $E \subseteq \mathrm{acl}_{\mathcal{M}}(A)$, $B_A \subseteq E$, and $a \in \mathrm{acl}_{\mathcal{M}}(A) \setminus E$, then any elementary function $g : E \to \mathrm{acl}_{\mathcal{N}}(C)$ extends to an elementary function $g' : E \cup \{a\} \to \mathrm{acl}_{\mathcal{N}}(C)$.

Since a is algebraic over E, there is a formula $\theta(x, y_1, \ldots, y_n)$ and parameters $e_1, \ldots, e_n \in E$ such that $\mathcal{M} \models \theta(a, e_1, \ldots, e_n)$ and such that for some $\ell \in \mathbb{N}$, $\theta(x, e_1, \ldots, e_n)$ is satisfied by exactly ℓ x in M. We choose θ and $e_1, \ldots, e_n$ such that ℓ is as small as possible. This means that for every formula $\psi(x, z_1, \ldots, z_m)$ and all parameters $d_1, \ldots, d_m \in E$ with $\mathcal{M} \models \psi(a, d_1, \ldots, d_m)$,

$$\mathcal{M} \models \forall x(\theta(x, e_1, \ldots, e_n) \to \psi(x, d_1, \ldots, d_m)),$$

since otherwise $\theta(x, e_1, \ldots, e_n) \wedge \psi(x, d_1, \ldots, d_m)$ witnesses that a is algebraic over E and has less solutions than $\theta(x, e_1, \ldots, e_n)$.

We have $\mathcal{M} \models \exists^{=\ell} x \theta(x, e_1, \ldots, e_n)$. Since g is elementary,

$$\mathcal{N} \models \exists^{=\ell} x \theta(x, g(e_1), \ldots, g(e_n)).$$

In particular, there is $b \in \mathrm{acl}_{\mathcal{N}}(C)$ such that $\mathcal{N} \models \theta(b, g(e_1), \ldots, g(e_n))$. We extend g to $E \cup \{a\}$ by letting $g'(a) = b$ and $g' \restriction E = g$. Now, let $\varphi(x, y_1, \ldots, y_m)$ be a formula and $d_1, \ldots, d_m \in E$. Suppose that $\mathcal{M} \models \varphi(a, d_1, \ldots, d_m)$. By the minimaliy of ℓ,

$$\mathcal{M} \models \forall x(\theta(x, e_1, \ldots, e_n) \to \varphi(x, d_1, \ldots, d_m)).$$

Since g is elementary, we have

$$\mathcal{N} \models \forall x(\theta(x, g(e_1), \ldots, g(e_n)) \to \varphi(x, g(d_1), \ldots, g(d_m))).$$

But since $\mathcal{N} \models \theta(b, g(e_1), \ldots, g(e_n))$, this implies

$$\mathcal{N} \models \varphi(b, g(d_1), \ldots, g(d_m)).$$

It follows that g' is elementary, as we wanted to show.

This second claim shows that by transfinite recursion, f can be extended to an elementary function $f' : \mathrm{acl}_{\mathcal{M}}(A) \to \mathrm{acl}_{\mathcal{N}}(C)$. It is clear that f' is an isomorphism onto its image. It remains to show that f' is onto $\mathrm{acl}_{\mathcal{N}}(C)$.

Let $c \in \mathrm{acl}_{\mathcal{N}}(C)$. Then there is an algebraic formula $\varphi(x, c_1, \ldots, c_n)$ with parameters in B_C witnessing that c is algebraic over B_C. Let $\ell \in \mathbb{N}$ be such that $\mathcal{N} \models \exists^{=\ell} x \varphi(x, c_1, \ldots, c_n)$. For each c_i let $a_i = f^{-1}(c_i)$. By elementarity of f, $\mathcal{M} \models \exists^{=\ell} x \varphi(x, a_1, \ldots, a_n)$. Now the ℓ distinct elements of $\mathrm{acl}_{\mathcal{M}}(A)$ satisfying $\varphi(x, a_1, \ldots, a_n)$ are mapped by f' to the ℓ distinct elements of $\mathrm{acl}_{\mathcal{N}}(C)$ that satisfy $\varphi(x, c_1, \ldots, c_n)$. In particular, there is $a \in \mathrm{acl}_{\mathcal{M}}(A)$ such that $f'(a) = c$. It follows that $f' : \mathrm{acl}_{\mathcal{M}}(A) \to \mathrm{acl}_{\mathcal{N}}(C)$ is an isomorphism. □

Corollary 2.65. *Let Φ be a countable complete theory that is strongly minimal. Then it is κ-categorical for every uncountable cardinal κ.*

Proof. By Lemma 2.62, it is enough to show that any two models of Φ of size κ have the same dimension. Clearly, the dimension of a model of Φ of size κ is at most κ. On the other hand, if $\mathcal{M}$ is a model of Φ and B is a basis for M, then $M = \mathrm{acl}_{\mathcal{M}}(B)$. But if Φ is countable, then Φ only uses a countable vocabulary τ. There are only countably many formulas over τ. It follows that there are at most as many algebraic formulas with parameters from B as B has finite subsets. Since each algebraic formula is only satisfied by finitely many elements of M, this implies that B is infinite and therefore we only have $|B|$ elements of M that are algebraic over B. Hence $\kappa = |M| \leq |B|$. It follows that $\dim(M) = \kappa$. □

Let us briefly mention a variation of strong minimality.

Definition 2.66. Let $\mathcal{M}$ be a τ-structure for a vocabulary τ that contains the binary relation symbol $\leq$. $\mathcal{M}$ is **o-minimal** if $\mathcal{M}$ is totally ordered by $\leq$ and every definable subset of $\mathcal{M}$ is a finite unition of intervals and singletons.

A theory is o-minimal if every model of the theory is o-minimal.

While the theory of a minimal structure is not necessarily strongly minimal, the theory of an o-minimal structure is o-minimal. Examples of o-minimal structures are real closed fields, dense linear orders, the ordered field of real numbers with exponentiation. Algebraically closed substructures of o-minimal structures have the exchange property. In particular, in o-minimal structure it is possible to define dimensions. However, o-minimal structures are not uncountably categorical.

3. Types

Definition 3.1. Fix a vocabulary τ. Let $\mathcal{M}$ be a τ-structure and $B \subseteq M$. Let $n \in \mathbb{N}$. Let Γ be a set of formulas $\varphi(x_1, \dots, x_n, b_1, \dots, b_m)$ that have free variables among $x_1, \dots, x_n$ and additional parameters from B. Γ is an **n-type over** B if it is **consistent** in the sense that any finitely many formulas from Γ are simultaneously satisfied by an n-tuple $(a_1, \dots, a_n) \in M^n$. Γ is an **n-type** if it is an n-type over $\emptyset$.

We write $\Gamma(x_1, \dots, x_n)$ instead of just Γ to indicate that Γ is an n-type and that the formulas in Γ have free variables among $x_1, \dots, x_n$. Γ is **realized by** $(a_1, \dots, a_n) \in M^n$ if for all $\varphi(x_1, \dots, x_n, b_1, \dots, b_m) \in \Gamma$, $\mathcal{M} \models \varphi(a_1, \dots, a_n, b_1, \dots, b_m)$. If no n-tuple in M^n realizes Γ, we say that $\mathcal{M}$ **omits** Γ.

Γ is a complete n-type over B if Γ is a maximal n-type over B. Otherwise, Γ is a partial type.

Example 3.2. a) The only 0-type over B is the set

$$\{\varphi(b_1, \dots, b_m) : \varphi(y_1, \dots, y_m) \text{ is a } \tau \text{ formula},\\ b_1, \dots, b_m \in B \text{ and } \mathcal{M} \models \varphi(b_1, \dots, b_m)\}.$$

b) For all $a_1, \dots, a_n \in M^n$,

$$\mathrm{tp}_{\mathcal{M}}((a_1, \dots, a_n)/B) = \{\varphi(x_1, \dots, x_n, b_1, \dots, b_m) :\\ \varphi(x_1, \dots, x_n, y_1, \dots, y_m) \text{ is a } \tau \text{ formula},\\ b_1, \dots, b_m \in B \text{ and } \mathcal{M} \models \varphi(a_1, \dots, a_n, b_1, \dots, b_m)\}$$

is a complete n-type, the **type of** $(a_1, \dots, a_n)$ **over** B.

c) If $\mathcal{M}$ is an elementary substructure of $\mathcal{N}$ and $(a_1, \dots, a_n) \in N^n$, then the type of $(a_1, \dots, a_n)$ in $\mathcal{N}$ is an n-type over B in $\mathcal{M}$.

Exercise 3.3. Determine all the complete 1-types in the structure $(\mathbb{Q}, \leq)$ over the set $\mathbb{Q}$.

Hint: Consider the real numbers (and more!).

Lemma 3.4. *Let $\mathcal{M}$ be a structure, $A \subseteq M$ and let Γ be an n-type over A in $\mathcal{M}$. Then $\mathcal{M}$ has an elementary extension $\mathcal{N}$ in which Γ is realized.*

Proof. Consider the elementary diagram eldiag($\mathcal{M}$) of $\mathcal{M}$. The vocabulary of eldiag($\mathcal{M}$) has a constant symbol c_a for every $a \in M$. We introduce new constant symbols $d_1, \dots, d_n$. Let Γ' be the theory obtained from Γ by replacing every formula $\varphi(x_1, \dots, x_n, b_1, \dots, b_m) \in \Gamma$ by $\varphi(d_1, \dots, d_n, c_{b_1}, \dots, c_{b_m})$.

Since every finite subset of Γ is realized by an n-tuple in M, the theory eldiag($\mathcal{M}$)$\cup\Gamma'$ is consistent. Now, if $\mathcal{N}$ is a model of eldiag($\mathcal{M}$)$\cup$ Γ', we can consider $\mathcal{N}$ as an elementary extension of $\mathcal{M}$. Clearly, $(d_1^{\mathcal{N}}, \dots, d_n^{\mathcal{N}})$ realizes Γ in $\mathcal{N}$. □

Definition 3.5. Given a theory Φ over the vocabulary τ, let $S_n(\Phi)$ be the set of all complete n-types (over $\emptyset$) that are realized in some model of Φ. Let $S(\Phi) = \bigcup_{n\in\mathbb{N}} S_n(\Phi)$.

$S_n(\Phi)$ carries a natural topology. $O \subseteq S_n(\Phi)$ is open if for all $\Gamma \in O$ there is a formula $\varphi(x_1, \dots, x_n, b_1, \dots, b_m) \in \Gamma$ such that all $\Delta \in S_n(\Phi)$ with $\varphi(x_1, \dots, x_n, b_1, \dots, b_m) \in \Delta$ are types in O.

Lemma 3.6. *For every theory Φ and every $n \in \mathbb{N}$ the space $S_n(\Phi)$ is compact.*

Observe that this lemma makes sense even for $n = 0$. In this case $S_n(\Phi)$ is the set of all complete theories that extend Φ.

Proof. We first show that $S_n(\Phi)$ is Hausdorff. Let $p, q \in S_n(\Phi)$. If $p \neq q$, there is a formula $\varphi(x_1, \dots, x_n)$ such that $\varphi \in p$ and $\varphi \notin q$. Since q is a complete type, $\varphi \notin q$ implies $\neg\varphi \in q$.

Let U be the set of all n-types that contain φ and let V be the set of all n-types that contain $\neg\varphi$. Now $p \in U$ and $q \in V$ and U and V are disjoint and open. This shows that $S_n(\Phi)$ is Hausdorff.

Now let $\mathcal{U}$ be a collection of open subsets of $S_n(\Phi)$ such that $S_n(\Phi) = \bigcup\mathcal{U}$. For each type $p \in S_n(\Phi)$ there is some $U \in \mathcal{U}$ such that $p \in U$. Since U is open, there is a formula $\varphi_p(x_1, \dots, x_n) \in p$ such that all n-types that contain this formula are elements of U. In order to show that $S_n(\Phi)$ is the union of finitely many sets from $\mathcal{U}$, it is enough to show that there are finitely many types $p_1, \dots, p_m \in S_n(\Phi)$ such that every type in $S_n(\Phi)$ contains at least one of the formulas $\varphi_{p_1}, \dots, \varphi_{p_m}$.

Suppose this is not the case. Then whenever $F \subseteq S_n(\Phi)$ is finite, there is an n-type q such that for all $p \in F$, $\varphi_p \notin q$. Since q is complete, $\{\neg\varphi_p : p \in F\} \subseteq q$. This shows that every finite subset of $\{\neg\varphi_p : p \in S_n(\Phi)\}$ can be realized. It follows that $\{\neg\varphi_p : p \in S_n(\Phi)\}$ is a type. This type extends to a complete type q that does not contain any formula of the form φ_p, $p \in S_n(\Phi)$, a contradiction. □

Exercise 3.7. Let $\mathcal{M}$ be a structure and $A \subseteq M$. For $n \in \mathbb{N}$ consider the space $S_n(\mathcal{M}, A)$ of complete n-types over A in $\mathcal{M}$. $S_n(\mathcal{M}, A)$ is topologized in analogy to $S_n(\Phi)$. Show that $S_n(\mathcal{M}, A)$ is compact.

3.1. **Isolated types.**

Definition 3.8. Let Φ be a theory. A type $p \in S_n(\Phi)$ is **isolated** if the set $\{p\}$ is open in $S_n(\Phi)$, i.e., if there is a formula $\varphi(x_1, \dots, x_n) \in p$ such that p is the only type in $S_n(\Phi)$ that contains φ. Similarly, we defined isolated types in $S_n(\mathcal{M}, A)$ for a structure $\mathcal{M}$ and $A \subseteq M$.

Lemma 3.9. *Let Φ be a complete theory and let $p \in S_n(\Phi)$ be isolated. Then p is realized in every model of Φ.*

Proof. Let $\varphi(x_1, \dots, x_n)$ be a formula such that p is the only type in $S_n(\Phi)$ that contains φ. Let $\mathcal{M}$ be a model of Φ in which p is realized.

Then $\mathcal{M} \models \exists x_1 \dots \exists x_n \varphi$. Since Φ is complete, we have $\exists x_1 \dots \exists x_n \varphi \in \Phi$. But this implies that in every model $\mathcal{N}$ of Φ there is an n-tuple $(a_1, \dots, a_n)$ such that $\mathcal{N} \models \varphi(a_1, \dots, a_n)$. Since p is isolated and since this is witnessed by φ, $\mathrm{tp}_{\mathcal{N}}(a_1, \dots, a_n) = p$. □

Theorem 3.10 (Omitting Types Theorem)**.** *Let Φ be a countable complete theory. If $p \in S_n(\Phi)$ is not isolated, then there is a countable model of Φ that omits p.*

Proof. Let $(c_i)_{i \in \mathbb{N}}$ be a sequence of new constant symbols. Let τ' be the vocabulary τ of Φ together with the new constant symbols c_i, $i \in \mathbb{N}$. We construct a maximal consistent theory Φ^+ over τ' with the following properties:

(1) $\Phi \subseteq \Phi^+$
(2) For every formula in Φ^+ of the form $\exists x \theta(x)$, there is there is some $i \in \mathbb{N}$ such that $\theta(c_i) \in \Phi^+$.
(3) For all $(i_1, \dots, i_n) \in \mathbb{N}^n$, there is a formula $\varphi(x_1, \dots, x_n) \in p$ such that $\neg\varphi(c_{i_1}, \dots, c_{i_n}) \in \Phi^+$.

We construct Φ^+ as the union of countably many approxiations Φ_m, $m \in \mathbb{N}$. Let $\Phi_0 = \Phi$, let $(\psi_i)_{i \in \mathbb{N}}$ be an enumeration of all sentences over τ' and let $(d_1^m, \dots, d_n^m)_{m \in \mathbb{N}}$ be an enumeration of the set of n-tuples from the set $C = \{c_i : i \in \mathbb{N}\}$. Since every Φ_m will be obtained by adding only finitely many sentences to Φ_0, for each m, only finitely many of the constants c_i occur in Φ_m.

Suppose Φ_m has been defined. If $\Phi_m \cup \{\psi_m\}$ is not consistent, let $\Phi'_m = \Phi_m \cup \{\neg\psi_m\}$. If $\Phi_m \cup \{\psi_m\}$ is consistent and ψ_m is of the form $\exists x \theta(x)$, choose some $i \in \mathbb{N}$ such that c_i does not occur in $\Phi_m \cup \{\exists x \theta(x)\}$ and let $\Phi'_m = \Phi_m \cup \{\exists x \theta(x), \theta(c_i)\}$. Otherwise let $\Phi'_m = \Phi_m \cup \{\psi_m\}$. This makes sure that (2) is satisfied for Φ^+.

Now, let $\varphi(d_1^m, \dots, d_n^m, \overline{c})$ be the conjunction of the finitely many sentences in Φ'_m that use constant symbols from C, where $\overline{c}$ is a tuple consisting of all elements of C that are used in Φ'_m and are not in the n-tuple $(d_1^m, \dots, d_n^m)$. Note that $\Phi \cup \{\varphi(d_1^m, \dots, d_n^m, \overline{c})\}$ is equivalent to Φ'_m in the sense that the first theory implies every sentence in the second and the second theory implies every sentence in the first. We consider the formula $\varphi(x_1, \dots, x_n, \overline{y})$ obtained from $\varphi(d_1^m, \dots, d_n^m, \overline{c})$ by replacing each occurence of d_i by x_i and each occurence of c_j by y_j.

Clearly, $\Phi'_m \models \exists \overline{y} \varphi(d_1, \dots, d_n, \overline{y})$. If $\exists \overline{y} \varphi(x_1, \dots, x_n, \overline{y}) \notin p$, then $\neg\exists \overline{y} \varphi(x_1, \dots, x_n, \overline{y}) \in p$ since p is complete. In this case let

$$\Phi_{m+1} = \Phi'_m \cup \{\exists \overline{y} \varphi(d_1, \dots, d_n, \overline{y})\}.$$

This takes care of (3). If

$$\exists \overline{y} \varphi(x_1, \dots, x_n, \overline{y}) \in p,$$

then, since p is not isolated, there is another type q such that

$$\exists \overline{y} \varphi(x_1, \dots, x_n, \overline{y}) \in q.$$

Let $\psi(x_1,\dots,x_n)$ be any formula in $q\setminus p$. Then $\neg\psi(x_1,\dots,x_n)\in p$.

Claim 3.11. $\Phi'_m\cup\{\psi(d^m_1,\dots,d^m_n)\}$ is consistent.

Since $\exists\overline{y}\varphi(x_1,\dots,x_n,\overline{y})\in q$, $q\cup\{\varphi(x_1,\dots,x_n,\overline{y})$ can be realized. Since the variables from C do not occur in $q=q(x_1,\dots,x_n)$, this shows that $q(d^m_1,\dots,d^m_n)\cup\{\varphi(d^m_1,\dots,d^m_n,\overline{c})$ is consistent. Let $\mathcal{M}$ be a model of $q(d^m_1,\dots,d^m_n)\cup\{\varphi(d^m_1,\dots,d^m_n,\overline{c})\}$. Since $q\in S(\Phi)$, $\Phi\subseteq q$. It follows that $\mathcal{M}$ is a model of Φ. Since $\psi(x_1,\dots,x_n)\in q$,

$$\mathcal{M}\models\Phi\cup\{\varphi(d^m_1,\dots,d^m_n,\overline{c}),\psi(d^m_1,\dots,d^m_n)\}.$$

But $\Phi\cup\{\varphi(d^m_1,\dots,d^m_n,\overline{c})\}\models\Phi'_m$. Hence $\mathcal{M}$ is a model of $\Phi'_m\cup\{\psi(d^m_1,\dots,d^m_n)\}$, which proves the claim.

Let $\Phi_{m+1}=\Phi'_m\cup\{\psi(d^m_1,\dots,d^m_n)\}$. The choice of $\psi(x_1,\dots,x_n)$ takes care of (3).

This finishes the construction of the theories Φ_m. Let $\Phi^+=\bigcup_{m\in\mathbb{N}}\Phi_m$. Since every Φ_m is consistent, so is Φ^+. Let $\mathcal{N}$ be a model of Φ^+. Let $M=\{c_i^{\mathcal{N}}:i\in\mathbb{N}\}$. (2) implies that M is the underlying set of an elementary substructure $\mathcal{M}$ of $\mathcal{N}$. (3) guarantees that $\mathcal{M}$ omits p. □

The proof of the omitting types theorem together with some book keeping can be used to show that countably many non-isolated types can be omitted simultaneously, in a countable model.

Corollary 3.12. *Let Φ be a countable complete theory. A type in $S(\Phi)$ is isolated iff it is realized in every countable model of Φ.*

Theorem 3.13. *Let Φ be a countable complete theory having only infinite models. Then the following are equivalent:*

(1) *$S(\Phi)$ is countable and every type in $S(\Phi)$ is isolated.*
(2) *Each $S_n(\Phi)$ is finite.*
(3) *Φ is $\aleph_0$-categorical.*

Proof. (1)⇒(2): Let $n\in\mathbb{N}$. If every type in $S_n(\Phi)$ is isolated, then every singleton in $S_n(\Phi)$ is open. Since $S_n(\Phi)$ is compact, the space is a finite union of singletons and therefore finite.

(2)⇒(3): Let $\mathcal{M}$ and $\mathcal{N}$ be two countable models of Φ. We use back-and-forth argument that is practically the same as the one used to show the uniqueness of the random graph. The crucial step is this:

Suppose we have chosen elements $p_1,\dots,p_n$ of M and $q_1,\dots,q_n$ of N such that $\mathrm{tp}_{\mathcal{M}}(p_1,\dots,p_n)=\mathrm{tp}_{\mathcal{N}}(q_1,\dots,q_n)$. Note that for $n=0$ this just says that $\mathcal{M}$ and $\mathcal{N}$ are elementarily equivalent, which they are. Let $p_{n+1}\in N\setminus\{p_1,\dots,p_n\}$. The type $\mathrm{tp}_{\mathcal{M}}(p_1,\dots,p_{n+1})$ is isolated and hence implied by a single formula $\varphi(x_1,\dots,x_{n+1})$. We have $\exists x_{n+1}\varphi(x_1,\dots,x_{n+1})\in\mathrm{tp}_{\mathcal{M}}(p_1,\dots,p_n)$. Hence $\exists x_{n+1}\varphi(x_1,\dots,x_{n+1})\in\mathrm{tp}_{\mathcal{N}}(q_1,\dots,q_n)$. It follows that there is some $q_{n+1}\in N$ such that $\mathcal{N}\models\varphi(q_1,\dots,q_{n+1})$. Now $\mathrm{tp}_{\mathcal{N}}(q_1,\dots,q_{n+1})$ is the type implied by $\varphi(x_1,\dots,x_{n+1})$ and hence equal to $\mathrm{tp}_{\mathcal{M}}(p_1,\dots,p_{n+1})$.

This allows it to recursively construct an isomorphism between $\mathcal{M}$ and $\mathcal{N}$.

(3)⇒(1): If $p \in S(\Phi)$ is not isolated, there is a countable model $\mathcal{M}$ of Φ that omits p. Let $\mathcal{N}$ be a model of Φ in which p is realized by some n-tuple $(a_1, \dots, a_n)$. By the Löwenheim-Skolem Theorem, $\mathcal{N}$ has a countable elementary submodel that contains $a_1, \dots, a_n$. It follows that p is realized in some countable model of Φ. This model is not isomorphic to $\mathcal{M}$. Hence Φ is not $\aleph_0$-categorical.

This shows that if Φ is $\aleph_0$-categorical, then all types are isolated. If all types are isolated, then, again by compactness, each $S_n(\Phi)$ is finite. Therefore $S(\Phi)$ is countable in this case. □

Exercise 3.14. Show that the theory of algebraically closed fields of characteristic p, $p = 0$ or p a prime number, has exactly one 1-type that is not isolated, namely the type of an element that is transcendent over the prime subfield.

Exercise 3.15. How many n-types does the theory of dense linear orders without endpoints have?

3.2. Saturation.

Definition 3.16. Let κ be a cardinal. A structure $\mathcal{M}$ is κ-saturated if for all $A \subseteq M$ with $|A| < \kappa$, every 1-type over A is realized in $\mathcal{M}$. $\mathcal{M}$ is **saturated** if it is $|M|$-saturated.

Lemma 3.17. *Let $\mathcal{M}$ be a structure for a countable vocabulary and let κ be an infinite cardinal. Suppose every 1-type over every set $A \subseteq M$ of size $< \kappa$ is realized in $\mathcal{M}$. Then for every $n \in \mathbb{N}$, every n-type over every set $A \subseteq M$ of size $< \kappa$ is realized in $\mathcal{M}$.*

Proof. We prove the lemma by induction on n. Suppose every every n-type and every 1-type over every set $A \subseteq M$ of size $< \kappa$ is realized in $\mathcal{M}$. Let $A \subseteq M$ be a set of size $< \kappa$ and let $\Gamma(x_1, \dots, x_{n+1})$ be a complete $n+1$-type over A. Let $\Gamma_n(x_1, \dots, x_n)$ be the subset of Γ consisting of formulas that only have the free variables $x_1, \dots, x_n$.

By the inductive hypothesis, there is an n-tuple $(a_1, \dots, a_n) \in M^n$ realizing Γ_n. We claim that $\Gamma(a_1, \dots, a_n, x_{n+1})$ is a 1-type over the set $A \cup \{a_1, \dots, a_n\}$.

To see that $\Gamma(a_1, \dots, a_n, x_{n+1})$ is a type let $\varphi_1, \dots, \varphi_k$ be formulas in Γ. Since Γ is complete,

$$\psi(x_1, \dots, x_{n+1}) = \varphi_1 \wedge \dots \wedge \varphi_k \in \Gamma.$$

The formula $\exists x_{n+1}\psi$ is consistent with any finite set of formulas in Γ. Since Γ is complete, $\exists x_{n+1}\psi \in \Gamma$. Hence $\exists x_{n+1}\psi \in \Gamma_n$. It follows that $\mathcal{M} \models \exists x_{n+1}\psi(a_1, \dots, a_n)$. Let $a \in M$ be a witness to this. Now a simultaneously realizes $\varphi_1(a_1, \dots, a_n, x_{n+1}), \dots, \varphi_k(a_1, \dots, a_n, x_{n+1})$, showing that $\Gamma(a_1, \dots, a_n, x_{n+1})$ is a 1-type.

By our assumption there is $a_{n+1} \in M$ realizing $\Gamma(a_1, \ldots, a_n, x_{n+1})$. Now $(a_1, \ldots, a_{n+1})$ realizes Γ. This finishes the inductive argument. $\square$

Example 3.18. a) The random graph is saturated:

Since the Φ_{random} has quantifier elimination, every type is equivalent to a type that consists of quantifier free formulas. Hence a complete 1-type over a finite set either says that the vertex realizing this type is equal to a certain vertex in the finite set or specifies to which vertices from the finite set the vertex realizing the type is connected. The extension axioms precisely state that every 1-type over a finite set is realized.

b) $(\mathbb{Q}, \leq)$ is saturated:

Again by quantifier elimination, we only have to consider types consisting of quantifier free formulas. But quantifier free 1-types over finite sets essentially only describe the relative position of a new point to these finitely many points. Since the order is dense and there are no endpoints, all these types are realized.

Theorem 3.19. *Let Φ be a countable complete theory that has only infinite models.*

a) If $\mathcal{M}$ is a κ-saturated model of Φ, then every model $\mathcal{N}$ of Φ of size $\leq \kappa$ embeds into $\mathcal{M}$ elementarily.

b) If $\mathcal{M}$ and $\mathcal{N}$ are both saturated models of Φ of the same size, then they are isomorphic.

Proof. a) Let $\{a_\alpha : \alpha < \kappa\}$ be a enumeration of N. We define an elementary embedding e of $\mathcal{N}$ into $\mathcal{M}$ by recursion. Suppose for some $\beta < \kappa$ we have defined $e(a_\alpha)$ for all $\alpha < \beta$. Consider the type $\text{tp}_{\mathcal{N}}(a_\beta/\{a_\alpha : \alpha < \beta\})$. This type corresponds to the type p in $\mathcal{M}$ over the set $\{e(a_\alpha) : \alpha < \beta\}$ that is obtained by replacing every formula of the form $\varphi(x, a_{i_1}, \ldots, a_{i_n})$ by $\varphi(x, e(a_{i_1}), \ldots, e(a_{i_n}))$. Since $\mathcal{M}$ is κ-saturated, there is some $b \in M$ that realizes p. Let $e(a_\beta) = b$. This finishes the recursive definition of e.

It is easily checked that $e : N \to M$ is an elementary embedding.

b) We use the same approach as in the proof of a), but in a back-and-forth way. Let $(a_\alpha)_{\alpha<\kappa}$ and $(b_\alpha)_{\alpha<\kappa}$ be enumerations of M and N, respectively. We call an ordinal α **even** if it is of the form $\beta + 2n$ for some $n \in \mathbb{N}$ and some β that is either 0 or a limit ordinal. Otherwise α is **odd**.

Now suppose that for some $\beta < \kappa$ we have defined two sequences $(p_\alpha)_{\alpha<\beta}$ and $(q_\alpha)_{\alpha<\beta}$ in M, respectively in N. If β is even, let p_β be the $a_\gamma \in M \setminus \{p_\alpha : \alpha < \beta\}$ of the smallest possible index. Choose $q_\beta \in N$ such that it realizes the type over $\{q_\alpha : \alpha < \beta\}$ that corresponds to $\text{tp}_{\mathcal{M}}(p_\beta/\{p_\alpha : \alpha < \kappa\})$. If β is odd, proceed in the same way with the roles of $\mathcal{M}$ and $\mathcal{N}$ switched. This finishes the recursive definition of two sequences $(p_\alpha)_{\alpha<\kappa}$ and $(q_\alpha)_{\alpha<\kappa}$ in M, respectively N.

It is easily checked that the map $p_\alpha \mapsto q_\alpha$ is an isomorphism between $\mathcal{M}$ and $\mathcal{N}$. $\square$

Definition 3.20. Let Φ be a complete theory without finite models over a countable vocabulary τ. Let κ be an infinite cardinal. A model $\mathcal{M}$ is κ-**universal** if every model $\mathcal{N}$ of Φ of size $< \kappa$ elementary embeds into $\mathcal{M}$. $\mathcal{M}$ is **universal** if it is $|M|^+$-universal.

A map $f : A \to M$ is **partial elementary** if it preserves all τ-formulas. $\mathcal{M}$ is κ-**homogenous** if for all $A \subseteq M$ with $|A| < \kappa$, every partial elementary map $f : A \to M$, and every $a \in M$ there is a partial elementary extension $\overline{f} : A \cup \{a\} \to M$ of f. $\mathcal{M}$ is **homogeneous** if it is $|M|$-homogeneous.

Note that by the usual back-and-forth argument, if $\mathcal{M}$ is homogeneous and $A \subseteq M$ is of size $< |M|$, then every partial elementary map $f : A \to M$ extends to an automorphism of $\mathcal{M}$.

Theorem 3.21. *Let κ be an infinite cardinal. Then a structure $\mathcal{M}$ is κ-saturated iff it is κ^+-universal and κ-homogeneous. If $\kappa > \aleph_0$, then $\mathcal{M}$ is κ-saturated iff it is κ-universal and κ-homogeneous.*

Proof. Let κ be an infinite cardinal. If $\mathcal{M}$ is κ-saturated, then it is κ^+-universal and κ-homogeneous by the arguments in the proof of Theorem 3.19. Clearly, κ^+ universality is stronger than κ-universality.

On the other hand, assume that $\mathcal{M}$ is κ^+-universal and κ-homogeneous. We show that $\mathcal{M}$ is κ-saturated.

Let $A \subseteq M$ be a set of size $< \kappa$ and let Γ be a 1-type over A. By the Löwenheim-Skolem theorem, there is an elementary substructure $\mathcal{M}_0$ of $\mathcal{M}$ of size κ containing A.

Let $\mathcal{N}$ be an elementary extension of $\mathcal{M}_0$ in which Γ is realized by some element b. We can choose $\mathcal{N}$ of size κ. By the κ^+-universality of $\mathcal{M}$, there is an elementary embedding $f : \mathcal{N} \to \mathcal{M}$.

Now consider the map $g : f[A] \to \mathcal{M}; f(a) \to a$. The map g is partial elementary. By the κ-homogeneity of $\mathcal{M}$, g extends to a partial elementary map $\overline{g} : f[A] \cup \{f(b)\} \to \mathcal{M}$. Now $g(f(b))$ realizes the type Γ.

If $\kappa > \aleph_0$, we can choose the elementary submodel $\mathcal{M}_0$ of size $< \kappa$. Also the elementary extension $\mathcal{N}$ can be chosen of size $< \kappa$. The rest of the argument goes through as before, giving the second part of the theorem. $\square$

Definition 3.22. For an infinite cardinal κ, $2^{<\kappa}$ denotes the cardinal $\sup\{2^\lambda : \lambda < \kappa\}$.

Recall that the **Generalized Continuum Hypothesis** (GCH) is the statement that for every infinite κ, 2^κ is the cardinal successor κ^+ of κ. GCH is equivalent to the statement that for all infinite κ, $2^{<\kappa} = \kappa$.

Theorem 3.23. *a) Let Φ be a complete countable theory with only infinite models. Then for every cardinal κ, Φ has a κ-saturated model.*

b) GCH implies that Φ has arbitrarily large saturated models.

Proof. a) By enlarging κ if necessary, we may assume that κ is a successor cardinal, i.e., $\kappa = \lambda^+$ for some infinite cardinal λ. We define an elementary chain $(\mathcal{M}_\alpha)_{\alpha<\lambda^+}$ of models of Φ. Let $\mathcal{M}_0$ be any model of Φ of size κ. Suppose β is some ordinal $< \lambda^+$ and $\mathcal{M}_\alpha$ has been defined for all $\alpha < \beta$.

Let $\mathcal{N}_0^\beta$ be the limit of the elementary chain $(\mathcal{M}_\alpha)_{\alpha<\beta}$. Let μ_β be the number of 1-types in $\mathcal{N}_0^\beta$ over sets of size $< \kappa$ and let $(p_\gamma)_{\gamma<\mu_\beta}$ be an enumeration of those types. We construct an elementary chain $(\mathcal{N}_\gamma^\beta)_{\gamma<\mu_\beta}$ as follows: Suppose for some $\nu < \mu_\beta$, $\mathcal{N}_\gamma^\beta$ has been defined for all $\gamma < \nu$. Let $\mathcal{L}_\nu^\beta$ be the limit of $(\mathcal{N}_\gamma^\beta)_{\gamma<\nu}$. Since $\mathcal{L}_\nu^\beta$ is an elementary extension of N_0^β, p_ν is a type in $\mathcal{L}_\nu^\beta$ over a set of size $< \kappa$. Let $\mathcal{N}_\nu^\beta$ be an elementary extension of $\mathcal{L}_\nu^\beta$ in which p_ν is realized. This finishes the definition of the elementary chain $(\mathcal{N}_\gamma^\beta)_{\gamma<\mu_\beta}$.

Now let $\mathcal{M}_\beta$ be the limit of $(\mathcal{N}_\gamma^\beta)_{\gamma<\mu_\beta}$. This finishes the definition of the elementary chain $(\mathcal{M}_\alpha)_{\alpha<\lambda^+}$. Finally, let $\mathcal{M}$ be the limit of $(\mathcal{M}_\alpha)_{\alpha<\lambda^+}$. Let $A \subseteq M$ be a set of size $< \kappa = \lambda^+$. Then for some $\alpha < \lambda^+$, $A \subseteq M_\alpha$. If p is a type in $\mathcal{M}$ over A, then by our construction, p is realized in $\mathcal{M}_{\alpha+1}$ and hence in $\mathcal{M}$.

b) The proof is exactly the same as for a), only we keep track of the sizes of the models $\mathcal{M}_\alpha$, $\alpha < \lambda^+$. If $\mathcal{N}$ is of size λ^+, then N has $(\lambda^+)^\lambda = (2^\lambda)^\lambda = 2^\lambda = \lambda^+$ subsets of size λ. Since Φ is countable, for each set $A \subseteq N$ of size λ, there are λ 1-types in $\mathcal{N}$ over A. It follows that in $\mathcal{N}$ there are λ^+ 1-types over sets of size $< \lambda^+$. It follows that all the structures that occur in the proof of a) can be chosen of size κ. In particular, the final model is of size κ and hence saturated. □

3.3. **Stability.**

Definition 3.24. Let κ be an infinite cardinal. A complete countable theory Φ is κ-**stable** if for every model $\mathcal{M}$ of Φ and every $A \subseteq M$ with $|A| = \kappa$ there are only κ many complete types over A.

A structure $\mathcal{M}$ is κ-**stable** if $\mathrm{Th}(\mathcal{M})$ is.

We write ω-stable for $\aleph_0$-stable.

Note that the existence of only few types enables us to construct small saturated models.

Theorem 3.25. *Let κ be an uncountable cardinal and suppose that Φ is a countable complete theory with only infinite models. If Φ is κ-stable, then for every $\lambda < \kappa$, Φ has a λ^+-saturated model of size κ. In particular, if κ is a successor cardinal, Φ has a saturated model of size κ.*

Proof. Using the method in the proof of Theorem 3.23, we construct an elementary chain $(\mathcal{M}_\alpha)_{\alpha<\lambda^+}$ of models of Φ of size κ such that for all $\beta<\kappa$ all types over $\bigcup_{\alpha<\beta} M_\alpha$ are realized in $\mathcal{M}_\beta$. This is possible since by κ-stability there are only κ types over $\bigcup_{\alpha<\beta} M_\alpha$. It is easily checked that the union of the $\mathcal{M}_\alpha$ is λ^+-saturated and of size κ. $\square$

Example 3.26. a) For p a prime or $p=0$ and for all infinite κ, the theory Φ_{ACF_p} is κ-stable.

b) The theory of dense linear orders without endpoints is not ω-stable.

Theorem 3.27. *If a countable complete theory Φ is ω-stable, then it is κ-stable for every infinite κ.*

Proof. Suppose that Φ is not κ-stable. Then for some model $\mathcal{M}$ of Φ and some n there is $A\subseteq M$ such that $|A|=\kappa$ and $|\mathcal{S}_n(\mathcal{M},A)|>\kappa$. We say that a formula $\varphi(x_1,\dots,x_n)$ with parameters in A (not shown) is **small** if it is contained in not more than κ complete types over A. Otherwise the formula is **large**.

Since there are only κ formulas with parameters in A and every small formula is contained in only κ many types over A, there are at most κ types that contain small formulas. It follows that every large formula is contained in at least two complete types that do not contain any small formulas.

Let $2^{<\omega}=\bigcup_{n\in\mathbb{N}} 2^n$. We construct a family

$$(\varphi_\sigma(x_1,\dots,x_n))_{\sigma\in 2^{<\omega}}$$

of formulas with parameters in A such that the following hold:

(1) Each φ_σ is large.
(2) For all $\sigma\in 2^{<\omega}$, φ_σ is equivalent to $\varphi_{\sigma^\frown 0}\vee\varphi_{\sigma^\frown 1}$.
(3) For all $\sigma\in 2^{<\omega}$, $\varphi_{\sigma^\frown 0}\wedge\varphi_{\sigma^\frown 1}$ is false.

Let $\varphi_\emptyset$ be any large formula. Suppose for some $\sigma\in 2^{<\omega}$ we have already chosen φ_σ. Choose two distinct complete types p and q containing φ_σ that contain only large formulas. Let $\psi\in p$ be such that $\neg\psi q$. Let $\varphi_{\sigma^\frown 0}=\varphi_\sigma\wedge\psi$ and $\varphi_{\sigma^\frown 1}=\varphi_\sigma\wedge\neg\psi$. Since p and q are complete, $\varphi_{\sigma^\frown 0}\in p$ and $\varphi_{\sigma^\frown 1}\in q$. Since p and q contain only large formulas, $\varphi_{\sigma^\frown 0}$ and $\varphi_{\sigma^\frown 1}$ are large.

Let $A_0\subseteq A$ be the set of parameters used in formulas of the form φ_σ, $\sigma\in 2^{<\omega}$. Clearly, A_0 is countable. For each $x\in 2^\omega$, $\{\varphi_{x\restriction k}:k\in\mathbb{N}\}$ is a type over A_0 and hence contained in a maximal n-type p_x over A_0. But for any distinct $x,y\in 2^\omega$, $p_x\neq p_y$. It follows that there are $2^{\aleph_0}$ complete n-types over the countable set A_0. Hence Φ is not ω-stable. $\square$

3.4. **Indiscernibles.**

Definition 3.28. Let $\mathcal{M}$ be a structure over a vocabulary τ and let $(I,<)$ be a linear order. A family $(a_i)_{i\in I}$ of elements of $\mathcal{M}$ is a sequence of **indiscernibles** if the a_i are pairwise distinct and for all τ-formulas $\varphi(x_1, \ldots, x_n)$ and strictly increasing n-tuples $(i_1, \ldots, i_n), (j_1, \ldots, j_n) \in I^n$

$$\mathcal{M} \models \varphi(a_{i_1}, \ldots, a_{i_n}) \leftrightarrow \varphi(a_{j_1}, \ldots, a_{j_n}).$$

Theorem 3.29. *If a theory Φ has an infinite model, then for every linear order $(I, <)$ there is a model of Φ with a sequence of indiscernibles indexed by I.*

The proof of this theorem uses the infinite Ramsey theorem in a more general form than the previously mentioned theorem for graphs. For a set X and a cardinal n let $[X]^n$ denote the collection of n-element subsets of X.

Theorem 3.30. *For all $n, k \in \mathbb{N}$ with $n, k > 0$, all infinite sets X, and all colorings $c : [X]^n \to k$ there is an infinite set $H \subseteq X$ such that c is constant on $[H]^n$.*

Proof of Theorem 3.29. For each $i \in I$ we introduce a new constant symbol a_i. Now consider the theory

$$\begin{aligned}\Phi_I = \Phi \cup \{a_i \neq a_j : i, j \in I, i \neq j\} \\ \cup \{\varphi(a_{i_1}, \ldots, a_{i_n}) \leftrightarrow \varphi(a_{j_1}, \ldots, a_{j_n}) : \\ \varphi \text{ is a } \tau\text{-formula and } (i_1, \ldots, i_n) \text{ and } (j_1, \ldots, j_n) \\ \text{are strictly increasing } n\text{-tuples in } I\}.\end{aligned}$$

Is is clear that any model of Φ_I is a model of Φ with a sequence of indiscernibles indexed by I.

We show that Φ_I is consistent. Let $\mathcal{M}$ be an infinite model of Φ and let $L \subseteq M$ be a countably infinite set. Choose a linear order $<$ of order type ω on L. Now let Φ_0 be a finite subset of Φ_I. We show that Φ_0 is consistent by interpreting the a_i in $\mathcal{M}$ such that the resulting structure satisfies Φ_0.

Let Δ be the set of all τ-formulas φ such that for some n and strictly increasing n-tuples $(i_1, \ldots, i_n)$ and $(j_1, \ldots, j_n)$ in I we have $\varphi = \varphi(x_1, \ldots, x_n)$ and $\varphi(a_{i_1}, \ldots, a_{i_n}) \leftrightarrow \varphi(a_{j_1}, \ldots, a_{j_n}) \in \Phi_0$. Let n be such that all $\varphi \in \Delta$ are formulas in the free variables $x_1, \ldots, x_n$.

Given a strictly increasing n-tuple $(b_1, \ldots, b_n)$ in L, let

$$c(\{b_1, \ldots, b_n\}) = \{\varphi \in \Delta : \mathcal{M} \models \varphi(b_1, \ldots, b_n)\}.$$

By Ramsey's theorem, there is an infinite set $H \subseteq L$ such that c is constant on $[H]^n$. Let I_0 be a finite subset of I such that all a_i that occur in Φ_0 have $i \in I_0$. Let $e : I_0 \to L$ be order preserving. For each $i \in I_0$ interpret a_i by $e(i)$. All other a_i can be interpreted arbitrarily. It is easily checked that this yields a model of Φ_0. □

Lemma 3.31. *If $\mathcal{M}$ is a structure for a countable vocabulary and $\mathcal{M}$ is generated by a well-ordered sequence of indiscernibles, then for each* *countable set $A \subseteq M$ only countably many types over A are realized in $\mathcal{M}$.*

Proof. Let $(a_i)_{i\in I}$ be a well-ordered sequence of indiscernibles generating $\mathcal{M}$ and let $S \subseteq M$ be countable. Then there is a countable set $A_0 \subseteq A$ such that every element of S can be obtained by evaluating a term in n variables at an n-tuple from A_0. Now the type of a tuple $\overline{b}$ from M over S only depends on its type over A_0.

In fact, if $\overline{b} = (b_1, \dots, b_n)$ is an n-tuple from M, then $b_1, \dots, b_n$ are already contained in the substructure of $\mathcal{M}$ generated by a finite set $C = \{a_{i_1}, \dots, a_{i_k}\} \subseteq A$. The type of $\overline{b}$ over A_0, and therefore over S, only depends on the type of $(a_{i_1}, \dots, a_{i_k})$ over A_0. But since $(a_i)_{i\in I}$ is a sequence of indiscernibles, the type of $(a_{i_1}, \dots, a_{i_k})$ over A_0 only depends on the quantifier free type that $(i_1, \dots, i_k)$ has in I over the set $I_0 = \{i \in I : a_i \in A_0\}$. Here the **quantifier free type** consists of all quantifier free formulas in the complete type.

We may assume that $i_1 < \dots < i_k$. The quantifier free type of $(i_1, \dots, i_k)$ over I_0 describes the relative positions of $i_1, \dots, i_k$ to the elements of I_0. For each $i \in I$ there are the following possible positions relative to I_0:

(1) i greater than all elements of I_0,
(2) i is equal to one element of I_0, or
(3) $i \notin I$ and there is some $j \in I_0$ that is the smallest element of I_0 above i.

Since I_0 is countable, there are only countably many possible positions of an element of I relative to I_0. It follows that there are countably many possible positions of $i_1, \dots, i_k$ relative to I_0. Hence in $\mathcal{M}$ there are only countably many types over S. □

Theorem 3.32. *Let κ be an infinite cardinal and let Φ be a theory with infinite models over a countable vocabulary τ. Then Φ has a model $\mathcal{M}$ of size κ such that for each countable set $S \subseteq M$ only countably many complete types over S are realized in $\mathcal{M}$.*

Proof. We define an increasing sequence $(\tau_n)_{n\in\mathbb{N}}$ of vocabularies and $(\Phi_n)_{n\in\mathbb{N}}$ of theories such that each Φ_n is a theory over the vocabulary τ_n. Let $\tau_0 = \tau$ and $\Phi_0 = \Phi$. Suppose we have already defined τ_n and Φ_n.

For each τ_n-formula $\varphi(x, x_1, \dots, x_n)$ we choose a new n-ary function symbol f_φ and let

$$\tau_{n+1} = \tau_n \cup \{f_\varphi : \varphi \text{ is a } \tau_n\text{-formula}\}.$$

Let

$$\Phi_{n+1} = \Phi_n \cup \{\forall x_1, \dots, x_n(\exists x\varphi \to \varphi(f_\varphi(x_1, \dots, x_n), x_1, \dots, x_n)) : \\ \varphi(x, x_1, \dots, x_n) \text{ is a } \tau_n\text{-formula}\}.$$

Now let $\tau' = \bigcup_{n\in\mathbb{N}} \tau_n$ and $\Phi' = \bigcup_{n\in\mathbb{N}} \Phi_n$. By induction on n, using Skolem functions it is easily checked that each model of Φ can be expanded to a model of Φ_n. Hence Φ' has an infinite model.

By Theorem 3.29, Φ' has a model $\mathcal{N}$ generated by a sequence $(a_i)_{i\in\kappa}$ of indiscernibles. By the construction of Φ', it has **built in Skolem functions**:

For every τ'-formula $\varphi(x, x_1, \dots, x_n)$ there is an n-ary function symbol $f_\varphi \in \tau'$ such that

$$\forall x_1, \dots, x_n(\exists x\varphi \to \varphi(f_\varphi(x_1, \dots, x_n), x_1, \dots, x_n)) \in \Phi'.$$

By the Tarski-Vaught criterion, this implies that every substructure of a model of Φ' is an elementary substructure.

Now let $\mathcal{M}$ be the substructure of $\mathcal{N}$ generated by the set $\{a_i : i \in \kappa\}$. Note that since τ' is countable, M is of size κ. The sequence $(a_i)_{i\in\kappa}$ is a sequence of indiscernibles for $\mathcal{M}$ and $\mathcal{M}$ is generated by the a_i.

By Lemma 3.31, for each countable set $S \subseteq N$, there are only countably many complete types over S realized in $\mathcal{N}$. Now, this is with respect to the vocabulary τ'. But every complete type of τ' formulas contains a unique complete type with respect to the vocabulary τ. It follows that for every countable set $S \subseteq N$ there are only countably many complete types with respect to τ over S realized in $\mathcal{N}$. This finishes the proof of the theorem. □

Corollary 3.33. *Let κ be an uncountable cardinal. If a theory Φ over a countable vocabulary is κ-categorical, then it is ω-stable.*

Proof. By Theorem 3.32, let $\mathcal{M}$ be a model of Φ of size κ such that over every countable set $S \subseteq M$ only countably many types are realized in $\mathcal{M}$. By the κ-categoricity of Φ, every model of Φ of size κ has the property that only countably many types are realized over countable subsets. If Φ is not ω-stable, we can construct a model $\mathcal{N}$ of Φ of size $\aleph_1$ such that over some countable subset of N there are uncountably many types realized in $\mathcal{N}$. We can then find an elementary extension of $\mathcal{N}$ of size κ. A contradiction. □

Corollary 3.34. *Let κ be an infinite cardinal. A complete countable theory Φ with only infinite models is κ-categorical iff every model of Φ of size κ is saturated.*

Proof. If every model of Φ of size κ is saturated, then these models are pairwise isomorphic and hence Φ is κ-categorical.

Now assume that Φ is κ-categorical. We first consider the case $\kappa > \aleph_0$. By Corollary 3.33, Φ is ω-stable. By Theorem 3.27, Φ is κ-stable. By Theorem 3.25, Φ has λ^+-saturated models of size κ for all $\lambda < \kappa$. Since Φ has only one model of size κ up to isomorphism, this model is λ^+-saturated for all $\lambda < \kappa$. It follows that the model is κ-saturated.

Now consider the case $\kappa = \aleph_0$. By Theorem 3.13, since Φ is $\aleph_0$-categorical, for each $n \in \mathbb{N}$, there are only finitely many n-types and each n-type is isolated. But this implies that every n-type over a finite subset of a model of Φ is generated by a single formula. Hence the type is realised. It follows that every model of Φ is $\aleph_0$-saturated. □

4. The categoricity theorem

In this section we give a full proof of Theorem 2.18.

4.1. **Prime models.** Fix a countable complete theory Φ without finite models.

Definition 4.1. A model $\mathcal{M}$ of Φ is a **prime model** of Φ if it embeds into every other model of Φ elementarily.

A model $\mathcal{M}$ of Φ is prime over a set $A \subseteq M$ if for every model $\mathcal{N}$ of Φ every partial elementary map $f : A \to \mathcal{N}$ extends to an elementary embedding

Theorem 4.2. *a) A model $\mathcal{M}$ of Φ is a prime model iff it is countable and exactly the isolated types are realized in $\mathcal{M}$.*

b) Any two prime models of Φ are isomorphic.

Proof. a) Since Φ has countable models, a prime model of Φ must be countable. Also, since every non-isolated type is omitted in some countable model, a prime model of Φ must omit all non-isolated types. On the other hand, all isolated types are realized in every countable model. It follows that prime models realize exactly the isolated types.

Now assume that $\mathcal{M}$ is a countable model of Φ in which only the isolated types are realized. Let $\mathcal{N}$ be a model of Φ. We construct an elementary embedding $f : \mathcal{M} \to \mathcal{N}$.

Let $(a_n)_{n\in\mathbb{N}}$ be an enumeration of M. Suppose for some $n \in \mathbb{N}$ and all $i < n$ we have already chosen $f(a_i) \in N$ such that f is partial elementary on $\{a_i : i < n\}$. Let $\varphi(x_0, \dots, x_n)$ be a formula that generates the type of $(a_0, \dots, a_n)$.

Since

$$\mathcal{M} \models \exists x_{n+1}\varphi(a_0, \dots, a_n, x_{n+1}),$$

$\exists x_{n+1}\varphi(x_0, \dots, x_n, x_{n+1})$ is in the type of $(a_0, \dots, a_n)$. Hence it is in the type of $(f(a_0), \dots, f(a_n))$. Choose $f(a_{n+1})$ such that $\mathcal{N} \models \varphi(f(a_0), \dots, f(a_(n+1))$. Now the n-tuple $(f(a_0), \dots, f(a_{n+1}))$ has the same type in $\mathcal{N}$ as $(a_0, \dots, a_{n+1})$ has in $\mathcal{M}$, namely the type generated by $\varphi(x_0, \dots, x_{n+1})$. Hence f is partial elementary on $\{a_1, \dots, a_{n+1}\}$. This finishes the recursive definition of f. Since f is partial elementary on on all of M, it is an elementary embedding of $\mathcal{M}$ into $\mathcal{N}$.

b) The proof uses the back-and-forth method using the argument in the proof of a) for extending partial elementary maps. □

Theorem 4.3. *Φ has a prime model iff for all $n \in \mathbb{N}$, the isolated types are dense in $S_n(\Phi)$.*

Proof. If the isolated types are not dense in $S_n(\Phi)$, then there is a formula $\varphi(x_1, \dots, x_n)$ that is consistent, i.e., realized in some model of $\mathcal{M}$, but not contained in any isolated type. We call such a formula **perfect**. Since Φ is complete, the formula $\exists x_1, \dots, x_n \varphi$ is in Φ. Hence every model of Φ contains an n-tuple satisfying φ. But the type of such an n-tuple is never isolated. Hence Φ has no prime model.

Now assume that the isolated types are dense in each $S_n(\Phi)$. We use a variation of the proof of the Omitting Types Theorem to construct a countable complete Henkin theory Φ^+ whose canonical model realizes only isolated types.

Let τ be the vocabulary of Φ. Let τ' be the vocabulary τ with additional constant symbols c_i, $i \in \mathbb{N}$. We construct a complete theory Φ^+ over τ' with the following properties:

(1) $\Phi \subseteq \Phi^+$.
(2) For every formula $\exists x \theta(x) \in \Phi^+$ there is $i \in \mathbb{N}$ such that $\theta(c_i) \in \Phi^+$.
(3) For all $(i_1, \dots, i_n) \in \mathbb{N}^n$, there is a complete formula $\varphi(x_1, \dots, x_n)$ such that $\varphi(c_{i_1}, \dots, c_{i_n}) \in \Phi^+$.

The construction of Φ^+ is as in the proof of the omitting types theorem. The crucial step is the following: we have constructed an extension Φ' of Φ such that $\Phi' \setminus \Phi$ is finite and we are given an n-tuple $(d_1, \dots, d_n)$ of constant symbols. Our task is to find a complete formula $\varphi(x_1, \dots, x_n)$ such that $\Phi' \cup \{\varphi(d_1, \dots, d_n)\}$ is consistent.

Let $\theta(d_1, \dots, d_n, \overline{c})$ be the conjunction of the sentences in $\Phi' \setminus \Phi$, where for some k, $\overline{c}$ is a k-tuple consisting of all the new constant symbols used in Φ' that are not among $d_1, \dots, d_n$. Since the isolated types are dense in $S_{n+k}(\Phi)$, $\theta(x_1, \dots, x_n, y_1, \dots, y_k)$ is contained in some isolated $n+k$-type p. The type p is generated by some complete formula $\varphi(x_1, \dots, x_n, y_1, \dots, y_k)$. But if $\varphi(x_1, \dots, x_n, y_1, \dots, y_k)$ is complete, then so is $\exists y_1, \dots, y_k \varphi$.

By the choice of φ, the theory $\Phi' \cup \{\exists y_1, \dots, y_k \varphi(d_1, \dots, d_n, y_1, \dots, y_k)\}$ is consistent. Φ^+ is an extension of this theory and therefore, the type of $(d_1, \dots, d_n)$ in a canonical model of Φ^+ we is isolated.

It follows that the canonical model of Φ^+ is a prime model of Φ since it is countable and realizes only isolated types. □

Lemma 4.4. *If Φ has only countably many complete types, then the isolated types are dense. In particular, Φ has a prime model.*

Proof. If for some $n \in \mathbb{N}$ the isolated types are not dense in $S_n(\Phi)$, then there is a perfect formula $\varphi(x_1, \dots, x_n)$, i.e., a consistent formula that is not contained in an isolated type. Starting with $\varphi_\emptyset = \varphi$ we choose a tree $(\varphi_\sigma)_{\sigma \in 2^{<\omega}}$ of consistent (with Φ) formulas with free variables among $x_1, \dots, x_n$ such that for all $\sigma \in 2^{<\omega}$ there is a formula ψ_σ with $\varphi_{\sigma^\frown 0} = \varphi_\sigma \wedge \psi_\sigma$ and $\varphi_{\sigma^\frown 1} = \varphi_\sigma \wedge \neg\psi_\sigma$.

This can be done recursively since whenever $\varphi(x_1,\dots,x_n)$ is perfect, there is a formula $\psi(x_1,\dots,x_n)$ such that φ implies neither ψ nor $\neg\psi$ and hence $\varphi\wedge\psi$ and $\varphi\wedge\neg\psi$ are both consistent.

Given the tree $(\varphi_\sigma)_{\sigma\in 2^{<\omega}}$ we get $2^{\aleph_0}$ pairwise distinct complete n-types as in the proof of Theorem 3.27. Hence, if Φ has only countably many complete types, then the isolated types are dense and Φ has a prime model. □

Note that this lemma implies that every ω-stable theory has only countably many types and has a prime model. However, we can do a bit better.

Lemma 4.5. *Let Φ be ω-stable and let $\mathcal{M}\models\Phi$. For all $A\subseteq M$ and all $n\in\mathbb{N}$, the isolated types are dense in $S_n(\mathcal{M},A)$.*

Proof. If the isolated types are not dense in $S_n(\mathcal{M},A)$, then there is a formula $\varphi(x_1,\dots,x_n)$ with parameters in A that is not contained in an isolated type. As in the proof of the previous lemma, we can build a tree $(\varphi_\sigma)_{\sigma\in 2^{<\omega}}$ of consistent formulas with parameters in A that are not contained in any isolated type. There is a countable set $A_0\subseteq A$ such that all φ_σ have only parameters from A_0. As in the proof of the previous lemma, we can conclude that there are $2^{\aleph_0}$ complete types over the set A_0, contradicting the ω-stability of Φ. □

This lemma has an important consequence.

Theorem 4.6. *If Φ is ω-stable, $\mathcal{M}\models\Phi$, and $A\subseteq\mathcal{M}$, then there is an elementary substructure $\mathcal{M}_0$ of $\mathcal{M}$ such that $\mathcal{M}_0$ is prime over A and every n-tuple from M_0 realizes an isolated type over A.*

Proof. For some ordinal γ we build a sequence $(A_\alpha)_{\alpha\leq\gamma}$ of subsets of M as follows:

Let $A_0=A$. If δ is a limit ordinal and A_α has been chosen for all $\alpha<\delta$, let $A_\delta=\bigcup_{\alpha<\delta}A_\alpha$. If A_α has been chosen and no $a\in M\setminus A_\alpha$ realizes an isolated type over A_α, then we stop and let $\gamma=\alpha$. If there is some $a\in M\setminus A_\alpha$ realizing an isolated type over A_α, we let $a_\alpha=a$ and $A_{\alpha+1}=A_\alpha\cup\{a_\alpha\}$.

Let $\mathcal{M}_0$ be the substructure of $\mathcal{M}$ on the set A_γ. Of course, at this point it is not clear that A_γ is the underlying set of a substructure of $\mathcal{M}$. However, if $\varphi(x,a_1,\dots,a_n)$ a formula with parameters in A_γ and there is some $a\in M$ such that $\mathcal{M}\models\varphi(a,a_1,\dots,a_n)$,

By Lemma 4.5, the isolated types are dense in $S_1(\mathcal{M},A_\gamma)$. Hence there is $b\in M$ such that $\mathcal{M}\models\varphi(b,a_1,\dots,a_n)$ and the type of b over A_γ is isolated. By the choice of γ, $b\in A_\gamma$.

By the Tarski-Vaught-criterion, $\mathcal{M}_0$ is indeed an elementary substructure of $\mathcal{M}$. Now let $\mathcal{N}\models\Phi$ and assume that $f:A\to\mathcal{N}$ is partial elementary. We define an extension $\overline{f}$ of f to all of A_γ.

Suppose $\overline{f}$ has been defined on some A_α and is partial elementary. Then a_α realizes some isolated type over A_α. Let $\varphi(x,b_1,\dots,b_n)$ be a formula with parameters in A_α that isolates this type. Since $\overline{f}$ is partial elementary on A_α, there is some $c \in N$ such that

$$\mathcal{N} \models \varphi(c, \overline{f}(b_1), \ldots, \overline{f}(b_n)).$$

Now we can extend $\overline{f}$ by letting $\overline{f}(a_\alpha) = c$. It is clear that $\overline{f}$ is an elementary embedding of $\mathcal{M}_0$ into $\mathcal{N}$ that extends f.

It remains to show that only isolated types over A are realized in M_0. We show by induction on $\alpha \leq \gamma$ that every n-tuple from A_α realizes an isolated type over A. The limit step is trivial. The successor step follows from the following lemma.

Lemma 4.7. *Suppose $A \subseteq B \subseteq M$, $\mathcal{M} \models \Phi$ and every n-tuple $\overline{b} \in B^n$ realizes an isolated type over A. Suppose that $\overline{a} \in M^m$ realizes an isolated type over B. Then $\overline{a}$ realizes an isolated type over A.*

Let $\varphi(\overline{x}, \overline{y})$ be a formula and let $\overline{b} \in B^m$ be such that $\varphi(\overline{x}, \overline{b})$ isolates the type of $\overline{a}$ over B. Let $\theta(\overline{y})$ be a formula with parameters in A that isolates the type of $\overline{b}$ over A. We claim that $\varphi(\overline{x}, \overline{y}) \wedge \theta(\overline{y})$ isolates the type of $(\overline{a}, \overline{b})$ over A.

Suppose that $\mathcal{M} \models \psi(\overline{a}, \overline{b})$. Let $\mathcal{M}_A$ denote the structure $\mathcal{M}$ expanded by the natural interpretations of new constant symbols for all elements of A. Since $\varphi(\overline{x}, \overline{b})$ isolates the type of $\overline{a}$ over B,

$$\mathrm{Th}(\mathcal{M}_A) \models \varphi(\overline{x}, \overline{b}) \rightarrow \psi(\overline{x}, \overline{b}).$$

Since $\theta(\overline{y})$ isolates the type of $\overline{b}$ over A,

$$\mathrm{Th}(\mathcal{M}_A) \models \theta(\overline{y}) \rightarrow (\varphi(\overline{x}, \overline{y}) \rightarrow \psi(\overline{x}, \overline{y}))$$

and

$$\mathrm{Th}(\mathcal{M}_A) \models (\theta(\overline{y}) \wedge \varphi(\overline{x}, \overline{y})) \rightarrow \psi(\overline{x}, \overline{y}).$$

This shows that the type of $(\overline{a}, \overline{b})$ over A is isolated. Now $\exists \overline{y}(\theta(\overline{y}) \wedge \varphi(\overline{x}, \overline{y}))$ isolates the type of $\overline{a}$ over A. □

4.2. **Vaughtian pairs.**

Definition 4.8. Let τ be a countable vocabulary. If $\mathcal{M}$ is a τ-structure and $\varphi(x_1, \ldots, x_n)$ is a τ-formula with parameters in M, we write $\varphi(\mathcal{M})$ for the set $\{\overline{a} \in M^n : \mathcal{M} \models \varphi(\overline{a})\}$.

A pair $(\mathcal{N}, \mathcal{M})$ is a **Vaughtian pair** if $\mathcal{M}$ is a proper elementary substructure of $\mathcal{N}$ and there is a formula φ with parameters in M such that $\varphi(\mathcal{M})$ is infinite and $\varphi(\mathcal{M}) = \varphi(\mathcal{N})$.

A theory Φ has a Vaughtian pair if there is a Vaughtian pair $(\mathcal{N}, \mathcal{M})$ of models of Φ. Φ has a (κ, λ)**-model** if there is a a model $\mathcal{M}$ of Φ of size κ and there is a formula $\varphi(x_1, \ldots, x_n)$ such that $\varphi(\mathcal{M})$ is of size λ.

Lemma 4.9. *If a countable theory Φ has a (κ, λ)-model where $\kappa > \lambda \geq \aleph_0$, then it has a Vaughtian pair.*

Proof. Let $\mathcal{N}$ be a (κ,λ)-model of Φ. Let φ be a formula such that $|\varphi(\mathcal{N})| = \lambda$. By the Löwenheim-Skolem theorem, there is an elementary submodel $\mathcal{M}$ of $\mathcal{N}$ of size λ such that $\varphi(\mathcal{N}) \subseteq \mathcal{M}$. Now $\varphi(M) = \varphi(N)$, but $\mathcal{M} \neq \mathcal{N}$. Hence $(\mathcal{N},\mathcal{M})$ is a Vaughtian pair. □

Lemma 4.10. *If a countable complete theory Φ has a Vaughtian pair, then it has a Vaughtian pair consisting of countable structures.*

Proof. Let $(\mathcal{N},\mathcal{M})$ be a Vaughtian pair of Φ and let $\varphi(x_1,\dots,x_n)$ be a formula with parameters in M such that $\varphi(\mathcal{M})$ is infinite and equals $\varphi(\mathcal{N})$. We introduce a new unary relation symbol U and interpret U in $\mathcal{N}$ by M. Slightly abusing notation, we write $U^{\mathcal{N}}$ for the interpretation of U in $\mathcal{N}$.

Let $(\mathcal{N}_0, U^{\mathcal{N}_0})$ be a countable elementary submodel of $(\mathcal{N}, U^{\mathcal{N}})$ containing the parameters used in φ. Then $U^{\mathcal{N}_0}$ carries an elementary substructure $\mathcal{M}_0$ of $\mathcal{N}_0$. Since Φ is complete, every model of Φ has infinitely many n-tuples satisfying φ.

Since

$$(\mathcal{N}, U^{\mathcal{N}}) \models \forall x_1,\dots,x_n(\varphi \leftrightarrow (\varphi \wedge U(x_1) \wedge \dots \wedge U(x_n))),$$

the same is true in $\mathcal{N}_0$. Also, $\mathcal{N} \models \exists x \neg U(x)$. Hence $\mathcal{N}_0 \models \exists x \neg U(x)$ and thus $\mathcal{M}_0 \neq \mathcal{N}_0$. It follows that $(\mathcal{N}_0,\mathcal{M}_0)$ is a Vaughtian pair. □

Lemma 4.11. *Let $\mathcal{M}_0 \preccurlyeq \mathcal{N}_0$ and suppose $\mathcal{N}_0$ is a countable model of Φ. Then there is a countable elementary extension $(\mathcal{N}, M)$ of $(\mathcal{N}_0, M_0)$ such that $\mathcal{N}$ and $\mathcal{M}$ are homogeneous and, for all $n \in \mathbb{N}$, realize the same types of $S_n(\Phi)$. In particular, $\mathcal{N}$ and $\mathcal{M}$ are isomorphic.*

Proof. We first observe that if $\overline{a}$ is an n-tuple in M_0 and p is a complete m-type over $\overline{a}$ that is realized in $\mathcal{N}_0$, then there is an elementary extension $(\mathcal{N}', M')$ of $(\mathcal{N}_0, M_0)$ such that p is realized in $\mathcal{M}'$. This follows from an easy compactness argument.

Iterating this observation, we find a countable elementary extension $(\mathcal{N}^*, M^*)$ of $(\mathcal{N}_0, M_0)$ such that for all finite tuples $\overline{a}$ from M_0 and all complete types p over $\overline{a}$, if p is realized in $\mathcal{N}_0$, then it is realized in $\mathcal{M}^*$.

A second observation is that for every n-tuple $\overline{b}$ in $\mathcal{N}_0$ and every complete m-type p over $\overline{b}$ there is an elementary extension $(\mathcal{N}', M')$ of $(\mathcal{N}_0, \mathcal{M}_0)$ such that p is realized in $\mathcal{N}'$.

Using this we construct a chain

$$(\mathcal{N}_0, M_0) \preccurlyeq (\mathcal{N}_1, M_1) \preccurlyeq \dots$$

of countable structures such that

1. (1) if $p \in S_n(\Phi)$ is realized in $\mathcal{N}_{3i}$, then it is realized in $\mathcal{M}_{3i+1}$,
2. (2) if $\overline{a}$, $\overline{b}$, and c are from $\mathcal{M}_{3i+1}$ and $\overline{a}$ and $\overline{b}$ have the same type in $\mathcal{M}_{3i+1}$, then there is $d \in M_{3i+2}$ such that $(\overline{a}, c)$ and $(\overline{b}, d)$ have the same type in $\mathcal{M}_{3i+2}$.

(3) if $\overline{a}$, $\overline{b}$, and c are from $\mathcal{N}_{3i+2}$ and $\overline{a}$ and $\overline{b}$ have the same type in $\mathcal{N}_{3i+2}$, then there is $d \in N_{3i+3}$ such that $(\overline{a}, c)$ and $(\overline{b}, d)$ have the same type in $\mathcal{N}_{3i+3}$.

For (1) and (2) we use the first observation, for (3) the second.

Now let $(\mathcal{N}, M)$ be the direct limit of the system $((\mathcal{N}_i, M_i))_{i\in\mathbb{N}}$. Then $(\mathcal{N}, \mathcal{M})$ is a Vaughtian pair of countable structures. By (1), $\mathcal{M}$ and $\mathcal{N}$ realize the same types. By (2) and (3), they are homogeneous. It follows that they are isomorphic. □

Theorem 4.12 (Vaught's two cardinal theorem)**.** *If a countable complete theory Φ has a (κ, λ)-model for $\kappa > \lambda \geq \aleph_0$, then Φ has an $(\aleph_1, \aleph_0)$-model.*

Proof. If Φ has a (κ, λ)-model, then it has a countable Vaughtian pair $(\mathcal{N}, M)$ such that $\mathcal{N}$ and $\mathcal{M}$ are homogeneous and realize the same types. Let $\varphi(x_1, \ldots, x_n)$ be a formula with parameters in M such that $\varphi(\mathcal{M})$ is infinite and equals $\varphi(\mathcal{N})$.

We build an elementary chain $(\mathcal{N}_\alpha)_{\alpha<\omega_1}$ such that each $\mathcal{N}_\alpha$ is isomorphic to $\mathcal{N}$ and $(\mathcal{N}_{\alpha+1}, N_\alpha)$ is isomorphic to $(\mathcal{N}, M)$ for all $\alpha < \omega_1$. In particular, no n-tuples outside N_0 satisfy φ.

Let $\mathcal{N}_0 = \mathcal{N}$. We take unions at limit stages. For a limit ordinal δ, $\mathcal{N}_\delta$ is the union of a chain of elementary submodels isomorphic to $\mathcal{N}$. It follows that $\mathcal{N}_\delta$ is homogeneous and realizes the same types as $\mathcal{N}$. Hence $\mathcal{N}_\delta$ is isomorphic to $\mathcal{N}$.

Given $\mathcal{N}_\alpha \cong \mathcal{N}$, since $\mathcal{N} \cong \mathcal{M}$ there is an elementary extension $\mathcal{N}_{\alpha+1}$ of $\mathcal{N}_\alpha$ such that $(\mathcal{N}_{\alpha+1}, N_\alpha)$ is isomorphic to $(\mathcal{N}, M)$. Clearly $\mathcal{N}_\alpha \cong \mathcal{N}$.

Let $\mathcal{N}^*$ be the union of the $\mathcal{N}_\alpha$, $\alpha < \omega_1$. Now $\mathcal{N}^*$ is of size $\aleph_1$ and if $\overline{a}$ satisfies φ in $\mathcal{N}^*$, then $\overline{a} \in M^n$. It follows that $\mathcal{N}^*$ is an $(\aleph_1, \aleph_0)$-model. □

Corollary 4.13. *If a countable complete theory Φ is $\aleph_1$-categorical, then it has no Vaughtian pairs and no (κ, λ)-models for $\kappa > \lambda \geq \aleph_0$.*

In the context of ω-stable theories, we can prove a converse of Vaught's two cardinal theorem.

Lemma 4.14. *Suppose that the countable complete theory Φ is ω-stable. If $\mathcal{M} \models \Phi$ and $|M| \geq \aleph_1$, then there is a proper elementary extension of $\mathcal{M}$ such that if $\Gamma(x_1, \ldots, x_n)$ is a countable type over M realized in $\mathcal{N}$, then $\Gamma(\overline{x})$ is realized in $\mathcal{M}$.*

Proof. We first show that there is a formula $\varphi(x)$ with parameters in M such that for all formulas $\psi(x)$ with parameters in M either $(\varphi \wedge \psi)(\mathcal{M})$ or $(\varphi \wedge \neg\psi)(\mathcal{M})$ is countable.

If there were no such formula φ, we could construct a binary tree $(\varphi_\sigma(x))_{\sigma\in 2^{<\omega}}$ of formulas with parameters in M such that for each $\sigma \in 2^{<\omega}$, $\varphi_\sigma(\mathcal{M})$ is uncountable and $\varphi_{\sigma^\frown 0}(\mathcal{M})$ and $\varphi_{\sigma^\frown 1}(\mathcal{M})$ are disjoint subsets of $\varphi_\sigma(\mathcal{M})$. As in previous arguments, the existence of such a tree of formulas contradicts the ω-stability of Φ.

Let $\varphi(x)$ be as above. Let

$$p = \{\psi(x) : \psi \text{ is a formula with parameters in } M \text{ and } |(\varphi \wedge \neg\psi)(\mathcal{M})| \leq \aleph_0\}.$$

It is easily checked that p is closed under finite conjunctions. It follows that p is a type. Since for each formula $\psi(x)$ with parameters in M the type p contains either ψ or $\neg\psi$, p is complete.

Let $\mathcal{M}'$ be an elementary extension of $\mathcal{M}$ containing an element c that realizes p. Since Φ is ω-stable, by Lemma 4.6 there is an elementary substructure $\mathcal{N}$ of $\mathcal{M}'$ that is prime over the set $M \cup \{c\}$ such that every n-tuple in $\mathcal{N}$ realizes an isolated type over $M \cup \{c\}$.

Now let $\Gamma(\overline{y})$ be a countable type over M realized by some $\overline{b}$ in $\mathcal{N}$. There is a formula $\theta(\overline{y}, x)$ with parameters from M such that $\theta(\overline{y}, c)$ isolates the type of $\overline{b}$ over $M \cup \{c\}$. We have $\exists\overline{y}\theta(\overline{y}, x) \in p$ and for all $\gamma \in \Gamma$,

$$\forall\overline{y}(\theta(\overline{y}, x) \to \gamma(\overline{y})) \in p.$$

Let

$$\Delta = \{\exists\overline{y}\theta(\overline{y}, x)\} \cup \{\forall\overline{y}(\theta(\overline{y}, x) \to \gamma(\overline{y})) : \gamma \in \Gamma\}.$$

Then $\Delta \subseteq p$ is countable, and if c' realizes Δ, then $\exists\overline{y}\theta(\overline{y}, c')$, and if $\theta(\overline{b}', c')$, then $\overline{b}'$ realizes Γ.

Choose an enumeration $(\delta_i(x))_{i\in\mathbb{N}}$ of Δ. For all $n \in \mathbb{N}$,

$$|(\varphi \wedge \neg(\delta_0 \wedge \cdots \wedge \delta_n))(\mathcal{M})| \leq \aleph_0$$

It follows that the set of all $a \in M$ that satisfy φ and realize Δ is uncountable. Let $c' \in M$ realize Δ and choose $\overline{b}'$ such that $\mathcal{M} \models \theta(\overline{b}', c')$. Then $\overline{b}'$ realizes Γ in $\mathcal{M}$. □

Theorem 4.15. *If Φ is a countable, complete, ω-stable theory with an $(\aleph_1, \aleph_0)$-model, then Φ has $(\kappa, \aleph_0)$-models for all $\kappa > \aleph_1$.*

Proof. Let $\mathcal{M}$ be an uncountable model of Φ and let $\varphi(x)$ be a formula with parameters in M such that $\varphi(\mathcal{M})$ is countable. Let $\mathcal{N}$ be a proper elementary extension of $\mathcal{M}$ such that every countable type realized in $\mathcal{N}$ is realized in $\mathcal{M}$. Then

$$\Gamma(x) = \{\varphi(x)\} \cup \{x \neq m : m \in M \text{ and } \mathcal{M} \models \varphi(m)\}.$$

is a countable type not realized in $\mathcal{M}$ and therefore also not in $\mathcal{N}$. Hence $\varphi(\mathcal{M}) = \varphi(\mathcal{N})$.

Iterating this construction, we can build an elementary chain $(\mathcal{M}_\alpha)_{\alpha<\kappa}$ such that $\mathcal{M}_0 = \mathcal{M}$ and $\mathcal{M}_{\alpha+1} \neq \mathcal{M}_\alpha$ but $\varphi(\mathcal{M}_\alpha) = \varphi(\mathcal{M}_0)$. If $\mathcal{N} = \bigcup_{\alpha<\kappa} \mathcal{M}_\alpha$, then $\mathcal{N}$ is a $(\kappa, \aleph_0)$-model of Φ. □

Corollary 4.16. *Let κ be an uncountable cardinal. If a countable theory Φ is κ-categorical, then it has no Vaughtian pairs.*

Proof. Assume that Φ has a Vaughtian pair. By the previous results, Φ has a $(\kappa, \aleph_0)$-model. An easy compactness argument shows that Φ has a model of size κ where every infinite definable set is of size κ. Clearly, these two models of Φ are not isomorphic. It follows that Φ is not κ-categorical. $\square$

4.3. **The Baldwin-Lachlan theorem.** We prove the following theorem, which immediately implies Morley's categoricity theorem.

Theorem 4.17 (Baldwin-Lachlan)**.** *Let Φ be a complete theory over a countable vocabulary with infinite models. Let κ be an uncountable cardinal. Then Φ is κ-categorical iff it is ω-stable and has no Vaughtian pairs.*

Recall that if Φ is κ-categorical, then by Corollary 4.16 Φ has no Vaughtian pair and by Corollary 3.33 it is ω-stable. It remains to show that if Φ is ω-stable and has no Vaughtian pairs, then it is κ-categorical.

Definition 4.18. Let τ be a countable vocabulary and let $\mathcal{M}$ be a τ-structure. A formula $\varphi(x_1, \dots, x_n)$ with parameters in M is **minimal** if all definable subsets of $\varphi(\mathcal{M})$ are finite or cofinite. The formula φ is **strongly minimal** if φ is minimal in every elementary extension of $\mathcal{M}$.

Lemma 4.19. *Let Φ be an ω-stable theory. Then every model of Φ has a minimal formula.*

Proof. Let $\mathcal{M}$ be an infinite model of Φ. If no formula is minimal, then we can build a tree $(\varphi_\sigma(x))_{\sigma \in 2^{<\omega}}$ of formulas such that for each $\sigma \in 2^{<\omega}$, $\varphi_\sigma(\mathcal{M})$ is infinite and for some formula $\psi(x)$ with parameters from M, $\varphi_{\sigma^\frown 0} = \varphi_\sigma \wedge \psi$ and $\varphi_{\sigma^\frown 1} = \varphi_\sigma \wedge \neg\psi$. The existence of this tree of formulas contradicts the ω-stability of Φ. $\square$

Lemma 4.20. *Suppose Φ is a theory without Vaughtian pairs over a countable vocabulary τ.. Let $\mathcal{M}$ be a model of Φ and let*

$$\varphi(x_1, \dots, x_m, y_1, \dots, y_k)$$

be a formula. Then there is a number n such that if $\overline{a}$ is a tuple from M and $|\varphi(\mathcal{M}, \overline{a})| > n$, then $\varphi(\mathcal{M}, \overline{a})$ is infinite.

Proof. If not, for every $n \in \mathbb{N}$ we can find a tuple $\overline{a}_n$ in M such that $\varphi(\mathcal{M}, \overline{a}_n)$ is of size at least n. We add a new unary relation symbol U to the vocabulary τ. Let $\Gamma(y_1, \dots, y_k)$ be the type saying

1. U supports an elementary submodel of the full structure,
2. $U(y_1) \wedge \cdots \wedge U(y_k)$,
3. there are infinitely many $\overline{x}$ satisfying $\varphi(\overline{x}, \overline{y})$,
4. $\forall \overline{x}(\varphi(\overline{x}, \overline{y}) \rightarrow (U(x_1) \wedge \cdots \wedge U(x_m)))$.

Let $\mathcal{N}$ be an elementary extension of $\mathcal{M}$. Since $\varphi(\mathcal{M}, \overline{a}_n)$ is finite, $\varphi(\mathcal{M}, \overline{a}_n) = \varphi(\mathcal{N}, \overline{a}_n)$. If $\Delta \subseteq \Gamma(y_1, \dots, y_k)$ is finite, then $\overline{a}_n$ realizes

Δ in $(\mathcal{N}, M)$, provided n is sufficiently large. Hence Γ can be realized in some structure $(\mathcal{N}', M')$ where $\mathcal{M}' \models \Phi$ and $\mathcal{N}'$ is an elementary extension of $\mathcal{M}'$.

Now let $\overline{a}$ realize Γ in $(\mathcal{N}', M')$. Then $\varphi(\mathcal{M}', \overline{a})$ is infinite and $\varphi(\mathcal{M}', \overline{a}) = \varphi(\mathcal{N}', \overline{a})$, contradicting the fact that Φ has no Vaughtian pairs. $\square$

Corollary 4.21. *If Φ has no Vaughtian pairs, then every minimal formula is strongly minimal.*

Proof. Let $\mathcal{M} \models \Phi$ and let $\varphi(\overline{x})$ be a minimal formula with parameters in M. Suppose there are an elementary extension $\mathcal{N}$ of $\mathcal{M}$, parameters $\overline{b}$ in N, and a formula $\psi(\overline{x}, \overline{y})$ such that $\varphi(\mathcal{N}) \cap \psi(\mathcal{N}, \overline{b})$ is a subset of $\varphi(\mathcal{N})$ that is neither finite nor cofinite.

By Lemma 4.20, there is $n \in \mathbb{N}$ such that for all parameters $\overline{a}$ in M, each of the two sets

$$|\psi(\mathcal{M}, \overline{a}) \cap \varphi(\mathcal{M})|$$

and

$$|\neg\psi(\mathcal{M}, \overline{a}) \cap \varphi(\mathcal{M})|$$

is infinite iff it is of size larger than n. Since $\varphi(\overline{x})$ is minimal, the following statement is true:

(∗) For all parameters $\overline{a}$ in M, either

$$|\psi(\mathcal{M}, \overline{a}) \cap \varphi(\mathcal{M})| \leq n$$

or

$$|\neg\psi(\mathcal{M}, \overline{a}) \cap \varphi(\mathcal{M})| \leq n.$$

But this statement can be expressed as $\mathcal{M} \models \theta$ for some sentence θ. Hence the statement holds for $\mathcal{N}$ instead of $\mathcal{M}$. But this contradicts our assumption that $\varphi(\mathcal{N}) \cap \psi(\mathcal{N}, \overline{b})$ is an infinite, coinfinite subset of $\varphi(N)$. $\square$

Lemma 4.22. *If Φ has no Vaughtian pairs, $\mathcal{M} \models \Phi$, and $X \subseteq M^n$ is infinite and definable, then no proper elementary submodel of $\mathcal{M}$ contains X. If Φ is also ω-stable, then $\mathcal{M}$ is prime over X.*

Proof. If $\mathcal{N}$ is a proper elementary submodel of $\mathcal{M}$ that contains X, then $(\mathcal{M}, \mathcal{N})$ is a Vaughtian pair. If Φ is ω-stable, then there is an elementary submodel $\mathcal{N}$ of $\mathcal{M}$ that is prime over X. Since Φ has no Vaughtian pairs, $\mathcal{M} = \mathcal{N}$. Hence $\mathcal{M}$ is prime over X. $\square$

Proof of Theorem 4.17. We have already argued that a κ-categorical theory Φ is ω-stable and has no Vaughtian pairs.

Now assume that Φ has no Vaughtian pairs and is ω-stable. Since Φ is ω-stable, it has a prime model $\mathcal{M}_0$. By Lemma 4.19, there is a minimal formula $\varphi(x)$ with parameters from $\mathcal{M}_0$. By Corollary 4.21, φ is strongly minimal.

Now let $\mathcal{M}$ and $\mathcal{N}$ be models of Φ of size κ. Both structure can be considered as elementary extensions of $\mathcal{M}_0$. We can define a notion of dimension for the minimal sets $\varphi(\mathcal{M})$ and $\varphi(\mathcal{N})$ just as in the case of minimal theories. Since Φ has no Vaughtian pairs, Φ has no (κ, λ)-models for any infinite $\lambda < \kappa$. Hence $\varphi(\mathcal{M})$ and $\varphi(\mathcal{N})$ are both of size κ. It follows that the dimension of both sets is κ.

Also as in the case of minimal theories, there is a partial elementary bijection $f : \varphi(\mathcal{M}) \to \varphi(\mathcal{N})$. $\mathcal{M}$ is prime over $\varphi(\mathcal{M})$ and hence we can extend f to an elementary embedding $\overline{f} : M \to N$. Since $\mathcal{N}$ has no proper elementary submodels containing $\varphi(\mathcal{N})$, $\overline{f}$ is onto. It follows that $\mathcal{M}$ and $\mathcal{N}$ are isomorphic. $\square$