Math 225A – Model Theory – Draft

Speirs, Martin

Autumn 2013

Contents

General Information

These notes are based on a course in *Metamathematics* taught by Professor Thomas Scanlon at UC Berkeley in the Autumn of 2013. The course will focus on Model Theory and the course book is Hodges' *a shorter model theory* (see [1]).

As with any such notes, these may contain errors and typos. I take full responsibility for such occurrences. If you find any errors or typos (no matter how trivial!) please let me know at mps@berkeley.edu.

Lecture 1

Model Theory is the study of the interrelation between structures and syntax. We shall start by defining the structures. To define a structure we need the data of a *signature* and then an *interpretation* of the signature.

1.1 The Category of τ-Structures

Definition. A **signature** τ consists of three (disjoint) sets $\mathcal{C}_\tau, \mathcal{F}_\tau, \mathcal{R}_\tau$ together with a function

$$\text{arity} : \mathcal{F}_\tau \cup \mathcal{R}_\tau \to \mathbb{Z}_+$$

The sets $\mathcal{C}_\tau, \mathcal{F}_\tau, \mathcal{R}_\tau$ will contain the constant symbols, function symbols and relation symbols, respectively. The arity function assigns to each function symbol and each relation symbol some positive integer thought of as the number of arguments that the function (respectively, the relation) takes. Note that we do not allow arities of functions and relations to be zero.

Definition. A **τ-structure** $\mathfrak{A}$ is given by a set A and interpretations of the elements of $\mathcal{C}_\tau \cup \mathcal{F}_\tau \cup \mathcal{R}_\tau$, i.e.:

- each $c \in \mathcal{C}_\tau$ is interpreted as an element $c^{\mathfrak{A}}$ in A.
- each $f \in \mathcal{F}_\tau$ is interpreted as a function $f^{\mathfrak{A}} : A^{\text{arity}(f)} \to A$.
- each $R \in \mathcal{R}_\tau$ is interpreted as a set $R^{\mathfrak{A}} \subseteq A^{\text{arity}(R)}$.

The set A is called the *domain of* $\mathfrak{A}$ and also denoted $\text{dom}(\mathfrak{A})$. We also use the notation $R(\mathfrak{A})$ for $R^{\mathfrak{A}}$ in anticipation of definable sets. For $a \in R^{\mathfrak{A}}$ we may also write $R(a)$.

Remark. In this course we do not require that A be nonempty!

Note that in order for $\varnothing$ to be a structure there can be no constant symbols (i.e. $\mathcal{C}_\tau = \varnothing$).

Remark. There is also a notion of a *sorted* signature, in which we would have another set $\mathcal{S}_\tau$ and in which "arity" would be replaced by giving the sort of each constant symbol, the domain and target of each function symbol and the field of each relation symbol. This is relevant in many situations for example when describing a vector space over a field (so we need two sorts: vectors and scalars) and in computer science.

Example (Groups). A group G may be regarded as a structure. The signature is in this case $\mathcal{C}_\tau = \{1\}$, $\mathcal{F}_\tau = \{\cdot\}$ and $\mathcal{R}_\tau = \varnothing$, and $\operatorname{arity}(\cdot) = 2$.

As an interpretation we might let 1^G be the identity element of G and $\cdot^G : G \times G \to G$ the group multiplication.

Example (Graphs). A graph G is a triple (V, E, I) of vertices, edges and an incidence relation, such that for $e \in E$ and $v, w \in V$ we have $I(v, w, e)$ is $(v, w) \in e$ (i.e. if e is an edge between v and w).

There are two natural signatures to use that do give different notions of graphs as structures.

- Let τ be given by $\mathcal{C}_\tau = \varnothing$, $\mathcal{F}_\tau = \varnothing$ and $\mathcal{R}_\tau = \{V, E, I\}$ where $\operatorname{arity}(V) = \operatorname{arity}(E) = 1$ and $\operatorname{arity}(I) = 3$. With this signature we can now set $\operatorname{dom}(G) = V \cup E$, $V^G = V$, $E^G = E$ and
$$I^G = \{(v, w, e) | (v, w) \in e\}.$$
- Let σ be the signature given by $\mathcal{C}_\sigma = \varnothing$, $\mathcal{F}_\sigma = \varnothing$ and $R_\sigma = \{E\}$ with $\operatorname{arity}(E) = 2$. Now $\operatorname{dom}(G) = V$ and
$$E^G = \{(v, w) \mid \exists e \in E \text{ such that } (v, w) \in e\}.$$

Now both signatures can be used to describe graphs but they are different. In the first case there can be multiple edges between the same vertices, while in the second there cannot. However in the first case we can have an edge $e \in E$ which is not connected to any vertices. So it makes a difference which language one uses!

We want to turn the collection of τ-structures into a category. For this we need morphisms.

Definition. A **homomorphism** $f : \mathfrak{A} \to \mathfrak{B}$ of τ-structures $\mathfrak{A}$ and $\mathfrak{B}$, is given by a function
$$f : \operatorname{dom}(\mathfrak{A}) \to \operatorname{dom}(\mathfrak{B})$$
which respects all the "extra structure". More precisely

- for all $c \in \mathcal{C}_\tau$ we have $f(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$
- for all $g \in \mathcal{F}_\tau$ (with say $n = \operatorname{arity}(g)$) and $a_1, \dots, a_n \in \operatorname{dom}(\mathfrak{A})$ then
$$f(g^{\mathfrak{A}}(a_1, \dots, a_n)) = g^{\mathfrak{B}}(f(a_1), \dots, f(a_n)).$$

- for all $R \in \mathcal{R}_\tau$ (with say $n = \text{arity}(R)$) if $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$ then

$$(f(a_1), \dots, f(a_n)) \in R^{\mathfrak{B}}.$$

Note that the notion of homomorphism depends on the choice of signature. For instance when defining homomorphisms of rings if we use a signature which has a constant symbol for the unit element then we get unit-preserving homomorphisms. If the signature does not have a constant symbol for the unit then homomorphisms of rings need not preserve the unit.

Proposition. *If* $f : \mathfrak{A} \to \mathfrak{B}$ *and* $g : \mathfrak{B} \to \mathfrak{C}$ *are homomorphisms of* τ*-structures then* $g \circ f$ *is a homomorphism from* $\mathfrak{A}$ *to* $\mathfrak{C}$*. Furthermore the identity map* $1_A : A \to A$ *gives a homomorphism of* τ*-structures* $1_{\mathfrak{A}} : \mathfrak{A} \to \mathfrak{A}$*.*

Proof. Exercise (purely formal). □

Thus, the collection of τ-structures together with homomorphisms between them form a category, called $\text{Str}(\tau)$.

Definition. For $\mathfrak{A}$ and $\mathfrak{B}$ τ-structures (with $\text{dom}(A) = A$ and $\text{dom}(\mathfrak{B}) = B$) then $\mathfrak{A}$ is a **substructure** of $\mathfrak{B}$, written $\mathfrak{A} \subseteq \mathfrak{B}$ if, $A \subseteq B$ and

- for all $c \in \mathcal{C}_\tau$, $c^{\mathfrak{A}} = c^{\mathfrak{B}}$.
- for all $f \in \mathcal{F}_\tau$, $f^{\mathfrak{A}} = f^{\mathfrak{B}}|_{A^n}$ where $n = \text{arity}(f)$.
- for all $R \in \mathcal{R}_\tau$, $R^{\mathfrak{A}} = R^{\mathfrak{B}} \cap A^n$ where $n = \text{arity}(R)$.

Proposition. *For* τ*-structures* $\mathfrak{A}$ *and* $\mathfrak{B}$*. If* $\mathfrak{A} \subseteq \mathfrak{B}$ *then the inclusion map* $\iota : \text{dom}(\mathfrak{A}) \to \text{dom}(\mathfrak{B})$ *is a homomorphism.*

Proof. We have a function from $\text{dom}(\mathfrak{A})$ to $\text{dom}(\mathfrak{B})$. We check all three conditions

- let $c \in \mathcal{C}_\tau$ then $\iota(c^{\mathfrak{A}}) = c^{\mathfrak{A}} = c^{\mathfrak{B}}$ since $\mathfrak{A}$ is a substructure of $\mathfrak{B}$.
- ι commutes with all interpretations of the function symbols since $\iota = 1_B|_A$.
- if $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$ then by the substructure property $(\iota(a_1), \dots, \iota(a_n)) = (a_1, \dots, a_n) \in R^{\mathfrak{B}}$.

so ι is indeed a homomorphism. □

Warning!. The converse of the above proposition is not true in general! There exist τ-structures $\mathfrak{A}$ and $\mathfrak{B}$ such that $\text{dom}(\mathfrak{A}) \subseteq \text{dom}(\mathfrak{B})$ and such that the inclusion map is a homomorphism and yet $\mathfrak{A}$ is *not* a substructure of $\mathfrak{B}$.

As an example, let τ be the signature given by $\mathcal{C}_\tau = \mathcal{F}_\tau = \varnothing$ and $\mathcal{R}_\tau = \{P\}$ with $\text{arity}(P) = 1$. Let $B = \mathbb{R}$ and $A = \mathbb{R}$ and consider these as τ-structures where $P^B = \mathbb{R}$ and $P^A = \varnothing$, respectively. Now the inclusion map $1_{\mathbb{R}} : A \to B$ is a homomorphism (we need only to check the condition on relation symbols, which is vacuous since $P^A = \varnothing$). However A is not a substructure of B since $P^A \neq P^B \cap A$.

Our notion of substructure is in some sense not the categorically correct notion. It is too restrictive. We want the homomorphisms (of τ-structures) to be the morphisms of $\mathrm{Str}(\tau)$ but the above example shows that a subobject is not given by a monic morphism.

Proposition. *For τ-structures $\mathfrak{A}, \mathfrak{B}$ and $\mathfrak{C}$. If $\mathfrak{A}, \mathfrak{B} \subseteq \mathfrak{C}$ and $\mathrm{dom}(\mathfrak{A}) = \mathrm{dom}(\mathfrak{B})$ then $\mathfrak{A} = \mathfrak{B}$*

Proof. By hypothesis $\mathrm{dom}(\mathfrak{A}) = \mathrm{dom}(\mathfrak{B})$. For $c \in \mathfrak{C}_\tau$ we have $c^{\mathfrak{A}} = c^{\mathfrak{C}} = c^{\mathfrak{B}}$. Let $f \in \mathcal{F}_\tau$ and $R \in \mathcal{R}_\tau$ of arity n. We have $f^{\mathfrak{A}} = f^{\mathfrak{C}}|_{A^n} = f^{\mathfrak{C}}|_{B^n} = f^{\mathfrak{B}}$ by the substructure property. Finally $R^{\mathfrak{A}} = R^{\mathfrak{C}} \cap A^n = R^{\mathfrak{C}} \cap B^n = R^{\mathfrak{B}}$ again by the substructure property. $\square$

So substructures are completely determined by their domain.

If $\{\mathfrak{B}_i\}_{i \in I}$ is a collection of substructures (with $B_i = \mathrm{dom}(\mathfrak{B}_i)$) of $\mathfrak{A}$ then $\bigcap \mathfrak{B}_i$ will denote the τ-structure whose domain is $\bigcap \mathrm{dom}(\mathfrak{B}_i)$ and where constants, functions and relations are interpreted on the intersection as before.

Proposition. *For $\mathfrak{A}$ a τ-structure and $X \subseteq \mathrm{dom}(\mathfrak{A})$ there exists a smallest (with respect to inclusion) substructure $\langle X \rangle \subseteq \mathfrak{A}$ such that $X \subseteq \mathrm{dom}(\langle X \rangle)$.*

Proof. We claim that if $\{\mathfrak{B}_i\}_{i \in I}$ is a collection of substructures of $\mathfrak{A}$ then $\bigcap \mathfrak{B}_i$ is a substructure of $\mathfrak{A}$.

If c is a constant symbol then $\forall i \in I$, $c^{\mathfrak{A}} \in B_i$ so $c^{\mathfrak{A}} \in \bigcap B_i$. For f a function symbol with arity n and $\bar{a} \in (\bigcap B_i)^n$ then $\forall j \in I$, $f^{\mathfrak{A}}(\bar{a}) \in B_j$ and so $f^{\mathfrak{A}}(\bar{a}) \in \bigcap B_i$. Thus $f^{\mathfrak{A}}|_{(\bigcap B_i)^n}$ is a function from $(\bigcap B_i)^n$ to $\bigcap B_i$. Likewise for R a relation symbol of arity n, define $R^{\bigcap B_i} := R^{\mathfrak{A}} \cap (\bigcap B_i)^n$. Now by definition $\bigcap \mathfrak{B}_i$ is a substructure of $\mathfrak{A}$.

With this in hand we now define the set of all substructures of $\mathfrak{A}$ that contain the subset X,

$$\mathfrak{X} := \{\mathfrak{B} : \mathfrak{B} \subseteq \mathfrak{A} \text{ and } X \subseteq \mathrm{dom}(\mathfrak{B})\}.$$

Note that $\mathfrak{A} \in \mathfrak{X}$. By the above claim, $\bigcap \mathfrak{X}$ is a substructure of $\mathfrak{A}$, and by definition $X \subseteq \mathrm{dom}(\bigcap \mathfrak{X})$. Furthermore if $X \subseteq \mathrm{dom}(\mathfrak{B})$ for some substructure $\mathfrak{B}$ then since $\mathfrak{B} \in \mathfrak{X}$ we have $\mathrm{dom}(\bigcap \mathfrak{X}) \subseteq \mathrm{dom}(\mathfrak{B})$. Thus setting $\langle X \rangle = \bigcap \mathfrak{X}$ completes the proof. $\square$

Remark. The above proposition is true as stated since we allow structures with empty domains.

Lemma. *(The Squash Lemma) If $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$ are τ-structures and* $\mathrm{dom}(\mathfrak{A}) \subseteq \mathrm{dom}(\mathfrak{B}) \subseteq \mathrm{dom}(\mathfrak{C})$*, and if both $\mathfrak{A} \subseteq \mathfrak{C}$ and $\mathfrak{B} \subseteq \mathfrak{C}$, then $\mathfrak{A} \subseteq \mathfrak{B}$.*

Proof. For $c \in \mathcal{C}_\tau$ we have $c^{\mathfrak{A}} = c^{\mathfrak{C}} = c^{\mathfrak{B}}$ by use of the given substructures. Similarly for function and relation symbols. □

The above construction of $\langle X \rangle$ doesn't actually show how to build $\langle X \rangle$. It works from above, since we know that there is some substructure containing X. Now we think of $\langle X \rangle$ as the substructure *generated* by X and we should be able to build $\langle X \rangle$ from below, simply using X.

In our attempt to build the substructure $\langle X \rangle$ we must first look at the constant symbols. For each $c \in \mathcal{C}_\tau$ if $c^{\mathfrak{A}} \notin X$ then we must add it. Futhermore for all function symbols and all tuples from X if f applied to these tuples is not in X then we must add these values. Thus we get a bigger set, and we can start over, and keep going until we finish. To make sense of this we introduce *terms*.

Definition. To a signature τ we have a set $\mathcal{T}(\tau)$ containing the **closed τ-terms**. $\mathcal{T}$ is given by recursion.

- $\mathcal{T}(\tau)$ contains all constant symbols
- if $f \in \mathcal{F}_\tau$ with $\mathrm{arity}(f) = n$ and $t_1, \ldots, t_n \in \mathcal{T}(\tau)$ then $f(t_1, \ldots, t_n) \in \mathcal{T}(\tau)$.

Remark. The above definition of $\mathcal{T}(\tau)$ has two subtle issues. For one we did not specify exactly what a term is. Secondly it is not clear that the above recursive definition actually defines a set. To actually justify these details requires quite a bit of set theory.

Lecture 2

Last time we introduced closed τ-terms. Before we start trying to make sense of how to interpret terms we must define the notions of expansions and reducts. Given two signatures τ and σ we write $\tau \subseteq \sigma$ when $\mathcal{C}_\tau \subseteq \mathcal{C}_\sigma$, $\mathcal{F}_\tau \subseteq \mathcal{F}_\sigma$ and $\mathcal{R}_\tau \subseteq \mathcal{R}_\sigma$.

Definition. Given signatures τ and σ with $\tau \subseteq \sigma$ and a σ-structure $\mathfrak{B}$ we may define a τ-structure $\mathfrak{A} = \mathfrak{B}|_\tau$. $\mathfrak{A}$ is the τ-structure given by $\mathrm{dom}(\mathfrak{A}) = \mathrm{dom}(\mathfrak{B})$ and for $x \in \mathcal{C}_\tau \cup \mathcal{F}_\tau \cup \mathcal{R}_\tau$ we set $x^{\mathfrak{A}} = x^{\mathfrak{B}}$. We call $\mathfrak{A}$ the **τ-reduct of** $\mathfrak{B}$ and say that $\mathfrak{B}$ is *an* **expansion of** $\mathfrak{A}$ **to** σ.

This gives a functor from $\mathrm{Str}(\sigma)$ (the category of σ-structures) to $\mathrm{Str}(\tau)$[1].

As long as either $\mathrm{dom}(\mathfrak{A}) \neq \varnothing$ or $\mathcal{C}_\sigma = \varnothing$ then the τ-structure $\mathfrak{A}$ admits some expansion to σ.

Example. We may think of $\mathbb{R}$ as an ordered field $(\mathbb{R}, +, \cdot, \leqslant, 0, 1)$[2]. Then the signature of this structure is $\{+, \cdot, \leqslant, 0, 1\}$. Now we may form the reduct to, say, the language of groups $\{+, 0\}$. This yields the structure $(\mathbb{R}, +, 0)$. Of course there are many ways to expand the group structure on the reals.

Recall that for a signature τ, the set $\mathcal{T}(\tau)$ of all **closed τ-terms** is the smallest set of finite sequences from $\mathcal{C}_\tau \cup \mathcal{F}_\tau \cup \{(\} \cup \{,\} \cup \{)\}$, such that for all $c \in \mathcal{C}_\tau$ then $c \in \mathcal{T}(\tau)$ and such that if $t_1, \ldots, t_n \in \mathcal{T}(\tau)$ and $f \in \mathcal{F}_\tau$ (with $\mathrm{arity}(f) = n$) then $f(t_1, \ldots, t_n) \in \mathcal{T}(\tau)$.

Remark. To show that $\mathcal{T}(\tau)$ is actually a set one uses weak recursion: given a set X an element $a \in X$ and a function $I : X \to X$ then there is a unique function $f : \omega \to X$ such that $f(0) = a$ and for all $n \in \omega$ we have $f(n+1) = I(f(n))$. This is a theorem, which will be proven in the homework.

[1]morphisms of σ-structures respect of the σ-structure and so they will also respect all the τ-structure.

[2]This notation for a structure means that $\mathbb{R}$ is the domain, $+$ and $\cdot$ are the interpretations of the function symbols, $\leqslant$ is the interpretation of the relation symbol, and $0, 1$ are the interpretations of the constant symbols.

Definition. For τ a signature we define the **free term τ-structure** $\widetilde{\mathcal{T}(\tau)}$ to be the τ-structure with domain $\mathcal{T}(\tau)$ and with interpretations as follows:

- for $c \in \mathcal{C}_\tau$ we set $c^{\widetilde{\mathcal{T}(\tau)}} = c$
- for $f \in \mathcal{F}_\tau$ with $\text{arity}(f) = n$ and with $a_1, \ldots, a_n \in \mathcal{T}(\tau)$, we set

$$f^{\widetilde{\mathcal{T}(\tau)}}(a_1, \ldots, a_n) = f(a_1, \ldots, a_n).$$

- for $R \in \mathcal{R}_\tau$ we let $R^{\widetilde{\mathcal{T}(\tau)}} = \varnothing$.

It is clear that $\widetilde{\mathcal{T}(\tau)}$ is in fact a τ-structure. The real content of this fact is just the fact that $\mathcal{T}(\tau)$ actually forms a set.

Furthermore $\widetilde{\mathcal{T}(\tau)}$ has a universal property.

Proposition. *For any τ-structure $\mathfrak{A}$ then there exists a unique homomorphism of τ-structures $\rho : \widetilde{\mathcal{T}(\tau)} \to \mathfrak{A}$.*

Proof. We first define the map $\rho : \widetilde{\mathcal{T}(\tau)} \to \mathfrak{A}$ by recursion on the construction on $\mathcal{T}(\tau)$.

- for $c \in \mathcal{C}_\tau$ we let $\rho(c^{\widetilde{\mathcal{T}(\tau)}}) = c^{\mathfrak{A}}$.
- if $t \in \mathcal{T}(\tau)$ has the form $f(t_1, \ldots, t_n)$ then $\rho(t) = f^{\mathfrak{A}}(\rho(t_1), \ldots, \rho(t_n))$.

This is well-defined since we have a unique parsing lemma for terms [Hodges Problem 5, of Section 1.3]. Furthermore ρ is clearly a homomorphism. On constant and function symbols it is defined as is should be and for relation symbols the claim is vacuous since $R^{\widetilde{\mathcal{T}(\tau)}} = \varnothing$. This takes care of the existence.

For the uniqueness we use induction on the complexity of terms. Suppose $\rho, \xi : \widetilde{\mathcal{T}(\tau)} \to \mathfrak{A}$ are homomorphisms. Then

- for all $c \in \mathcal{C}_\tau$ we have $\rho(c) = c^{\mathfrak{A}} = \xi(c)$,
- if $t \in \mathcal{T}(\tau)$ has the form $f(t_1, \ldots, t_n)$ then $\rho(t) = f^{\mathfrak{A}}(\rho(t_1), \ldots, \rho(t_n)) = f^{\mathfrak{A}}(\xi(t_1), \ldots, \xi(t_n)) = \xi(t)$ since the t_i's have lower complexity than t.

Thus $\rho = \xi$ and this finishes the proof. $\square$

The proposition shows that the free term structure, $\widetilde{\mathcal{T}(\tau)}$, is an initial object in $\text{Str}(\tau)$. [As such it is the unique (up to isomorphism) τ-structure which satisfies the property of the proposition.].

This way of constructing structures out of their own names will be done several times during the course. If one wants to show that certain sentences are consistent or that it is possible to have some specific kind of structure, then one can try to write down what one wants to exist and then the description is *itself* the structure.

We now introduce one of the most important kinds of expansions.

Definition. Given a τ-structure $\mathcal{M}$ and a subset $A \subseteq \mathrm{dom}(\mathcal{M})$, then τ_A is the signature with

$$\mathcal{C}_{\tau_A} = \mathcal{C}_\tau \dot{\cup} A \qquad \text{(disjoint union)}$$

and with $\mathcal{F}_{\tau_A} = \mathcal{F}_\tau$, and $\mathcal{R}_{\tau_A} = \mathcal{R}_\tau$. We define $\mathcal{M}_A$ to be the expansion of $\mathcal{M}$ to τ_A by interpreting $a \in A \subseteq \mathcal{C}_{\tau_A}$ as $a^{\mathcal{M}_A} = a$.

The expansion $\mathcal{M}_A$ has names (in the form of constant symbols) for all the elements of A.

We sometimes want to talk about having *variables*. Variables should be thought of as constant symbols that we don't know how to interpret yet.

Definition. Given a signature τ, a **term** over τ is an element of $\mathcal{T}(\tau_X)$ where $X = \{x_i : i \in \omega\}$.

[strictly speaking we have not defined what τ_X means in the context where there is no τ-structure, i.e. no domain.]

Now let $\mathfrak{A}$ be a τ-structure. We shall interpret the terms in $\mathcal{L}(\tau_X)$ in $\mathfrak{A}$. Let t be a term in which only the variables x_i for $i < n$ occur, so that $t \in \mathcal{L}(\tau_{\{x_i : i<n\}})$. Then $t^{\mathfrak{A}} : A^n \to A$ is the function given by sending, for each $\bar{a} = (a_0, \dots, a_{n-1}) \in A^n$ to the image of t under the unique $\tau_{\{x_i:i<n\}}$-homomorphism

$$\mathcal{T}(\widetilde{\tau_{\{x_i:i<n\}}}) \longrightarrow \mathfrak{A}_{\bar{a}}$$

where $x_i^{\mathfrak{A}_A} = a_i$.

Remark. If $A = \varnothing$ then there are no n-tuples $\bar{a} \in A^n$ and so we interpret t as the empty function.

2.1 Logics

Definition. Given a signature τ, an **atomic formula** is a finite sequence from the set

$$\mathcal{C}_\tau \cup \mathcal{F}_\tau \cup \mathcal{R}_\tau \cup \{(\} \cup \{,\} \cup \{)\} \cup \{=\}$$

of the form

- $t = s$, or
- $R(t_1, \dots, t_n)$

where t, s and $t_1, \dots, t_n$ are τ-terms, and $R \in \mathcal{R}_\tau$.

The set of all τ-formulae, $\mathcal{L}(\tau)$, is the smallest set of finite sequences in

$$\mathcal{C}_\tau \cup \mathcal{F}_\tau \cup \mathcal{R}_\tau \cup \{(\} \cup \{,\} \cup \{)\} \cup \{=\} \cup \{\vee\} \cup \{\wedge\} \cup \{\rightarrow\} \cup \{\leftrightarrow\} \cup \{\forall\} \cup \{\exists\} \cup \{x_i : i \in \omega\}$$

(where all unions are disjoint) such that

- every atomic formula belongs to $\mathcal{L}(\tau)$,
- if φ is in $\mathcal{L}(\tau)$, then $\neg(\varphi)$ is in $\mathcal{L}(\tau)$,
- if φ and ψ are in $\mathcal{L}(\tau)$, then $(\varphi \vee \psi), (\varphi \wedge \psi), (\varphi \to \psi)$ and $(\varphi \leftrightarrow \psi)$ are in $\mathcal{L}(\tau)$
- if φ is in $\mathcal{L}(\tau)$ and $i \in \omega$ then $(\exists x_i)\varphi$ and $(\forall x_i)\varphi$ are in $\mathcal{L}(\tau)$.

an element of $\mathcal{L}(\tau)$ is called a **formula**, and $\mathcal{L}(\tau)$ is called the **language** of τ.

Remark. Each of the four conditions on the set $\mathcal{L}(\tau)$ may be thought of as closure properties of $\mathcal{L}(\tau)$. For instance $\mathcal{L}(\tau)$ is closed under taking negation, i.e. if $\varphi \in \mathcal{L}(\tau)$ then $\neg(\varphi) \in \mathcal{L}(\tau)$. By the way we have defined $\mathcal{L}(\tau)$ it is clear that the set actually exists. This is since there is at least one set satisfying all the closure properties (namely the set of all sequences in the given symbols) and since each condition in the definition is such that for any collection of sets of sequences satisfying the given condition, their intersection will also satisfy it. Thus taking the intersection of all sets satisfying the conditions we get the smallest set, namely $\mathcal{L}(\tau)$.

This construction of $\mathcal{L}(\tau)$ is "from above". A more useful way to construct $\mathcal{L}(\tau)$ would be "from below" namely using weak recursion as in the above construction of the closed τ-terms.

2.2 Free and bound variables

We would like to say that a variable is free (or bound) in a formula φ but really we can only say that a particular instance of the given variable is free (or bound).

We define free and bound variables by recursion on the construction of formulae.

- In an atomic formula all variables are free, including variables not appearing in the atomic formula.
- In $\neg(\varphi)$, $(\varphi \wedge \psi)$ and $(\varphi \vee \psi)$ the free (respectively bound) instances of variables are what they where in the constituent formulae. For clarity let us be more precise in the case $(\varphi \wedge \psi)$. Now $(\varphi \wedge \psi)$ is a sequence of length, $3 + \text{length}(\varphi) + \text{length}(\psi)$. For $i < \text{length}(\varphi) + 1$ then the i^{th} coordinate is a free (respectively bound) variable if the $(i-1)^{th}$ coordinate of φ is free (respectively bound). For $2 + \text{length}(\varphi) \leqslant j < 3 + \text{length}(\varphi) + \text{length}(\psi)$ then the j^{th} coordinate is a free (respectively bound) variable if the $(j - (2 + \text{length}(\varphi)))^{th}$ coordinate is free (respectively bound) in ψ.
- In $(\forall x_i)\varphi$ and $(\exists x_i)\varphi$ no instance of x_i is free (i.e. all such instances are bound) and all other variables remain how they were (free or bound) in φ.

Warning!. The same variable can appear twice in the same formula as both a free and bound variable! For example in the formula $(\exists x_1(\neg(x_1 = x_2)) \wedge x_1 = x_3)$, the variable x_1 is bound in the first instance and free in the second. Of course it is not a good idea to do this, but it is allowed.

We can now understand how to interpret formulae. Given a τ-structure $\mathfrak{A}$ (with domain A) and φ a formula we interpret φ as follows

- If φ is atomic and equal to $t = s$ then $\varphi(\mathfrak{A}) := \{a \in \mathfrak{A}^\omega : t^{\mathfrak{A}}(a) = s^{\mathfrak{A}}(a)\}$.
- If φ is atomic and equal to $R(t_1, \dots, t_n)$ then $\varphi(\mathfrak{A}) := \{a \in A^\omega : (t_1^{\mathfrak{A}}(a), \dots, t_n^{\mathfrak{A}}(a)) \in R^{\mathfrak{A}}\}$.
- If φ is $\neg(\psi)$ then $\varphi(\mathfrak{A}) := A^\omega \setminus \varphi(\mathfrak{A})$.
- If φ is $(\psi \wedge \theta)$ then $\varphi(\mathfrak{A}) := \psi(\mathfrak{A}) \cap \theta(\mathfrak{A})$.
- If φ is $(\psi \vee \theta)$ then $\varphi(\mathfrak{A}) := \psi(\mathfrak{A}) \cup \theta(\mathfrak{A})$.
- If φ is $(\exists x_i)\psi$ then

  $\varphi(\mathfrak{A}) := \{a \in A^\omega : \exists b_i \in A \text{ such that } \tilde{a} \in \psi(\mathfrak{A}) \text{ where } (\tilde{a})_j = b_i \text{ if } j = i \text{ and } (\tilde{a})_j = a_j \text{ otherwise}\}$

- If φ is $(\forall x_i)\psi$ then

  $\varphi(\mathfrak{A}) := \{a \in A^\omega : \forall b_i \in A \text{ such that } \tilde{a} \in \psi(\mathfrak{A}) \text{ where } (\tilde{a})_j = b_i \text{ if } j = i \text{ and } (\tilde{a})_j = a_j \text{ otherwise}\}$

Remark. Often when $\varphi \in \mathcal{L}(\tau)$ is a formula and the free variables of φ are taken from $\{x_i : i < n\}$, we think of $\varphi(\mathfrak{A})$ as a subset of A^n. *This is a mistake.* By the above definition $\varphi(\mathfrak{A})$ is a subset of A^ω.

Definition. A formula $\varphi \in \mathcal{L}(\tau)$ is a **sentence** if no free variables appear in φ

Note that all variables not *appearing* in φ are free. So for a formula to be a sentence we only care about the variables actually appearing.

Definition. For φ a sentence we say that $\mathfrak{A}$ models φ, written $\mathfrak{A} \models \varphi$, if $\varphi(\mathfrak{A}) = A^\omega$.

Lecture 3

We continue studying the connection between language and signature. Last time we discussed how one might interpret a language in a structure. Today we will go the other way and associate to each structure a certain set of formulae which describe the structure.

3.1 Diagrams

We let τ be a signature and $\mathfrak{A}$ a τ-structure with domain A. Recall that a *sentence* in $\mathcal{L}(\tau)$ is a formula with no free variables.

Definition. The **theory of** $\mathfrak{A}$, written $\mathrm{Th}(A)$ (or $\mathrm{Th}_{\mathcal{L}(\tau)}(\mathfrak{A})$ to emphasize the signature), is the set of sentences ψ in $\mathcal{L}(\tau)$ that such that $\mathfrak{A} \models \psi$. I.e.

$$\mathrm{Th}(\mathfrak{A}) = \{\psi \in \mathcal{L}(\tau) : \psi \text{ is a sentence and } \mathfrak{A} \models \psi\}.$$

The theory of $\mathfrak{A}$ contains all that can be said about the structure $\mathfrak{A}$ using the language $\mathcal{L}(\tau)$. There are some subclasses of $\mathrm{Th}(\mathfrak{A})$ which are also of interest, for instance we might look at all *quantifier-free* sentences true in $\mathfrak{A}$ or all *existential* sentences true in $\mathfrak{A}$.

If we want to describe the basic structure $\mathfrak{A}$ itself, (i.e. answer questions such as; What are the relations? What are the functions? What are the constants inside of $\mathfrak{A}$?) then we look at the *diagram* of $\mathfrak{A}$. To define this object we first need to say what a *literal* sentence is.

Definition. A sentence of the form φ or $\neg(\varphi)$ where φ is atomic is called a **literal**.

Definition. The **diagram of** $\mathfrak{A}$, written $\mathrm{diag}(\mathfrak{A})$, is the set of literals that are true in $\mathfrak{A}_A$, i.e.

$$\mathrm{diag}(\mathfrak{A}) = \{\varphi \in \mathcal{L}(\tau_A) : \varphi \text{ is a literal and } \mathfrak{A}_A \models \varphi\}.$$

The **positive diagram of** $\mathfrak{A}$, written $\mathrm{diag}^+(\mathfrak{A})$ is the set of atomic formulae true in $\mathfrak{A}_A$, i.e.

$$\mathrm{diag}^+(\mathfrak{A}) = \{\varphi \in \mathcal{L}(\tau_A) : \varphi \text{ is a atomic and } \mathfrak{A}_A \models \varphi\}.$$

The diagram should be thought of as the "multiplication table" of the structure in analogy with the multiplication table of a group – even though the diagram is not *strictly* a generalization of the multiplication table in the case where $\mathfrak{A}$ is a group. But it contains the same information. For example, suppose G is a structure in the language of groups and we have $a \cdot b = c$ in G. Then the sentence $a \cdot^G b = c$ will be in the diagram. Or suppose $\mathfrak{A}$ is in the structure of partial orders. Then the diagram will include information like $a < b$ and $a \not< b$ in the partial order. In this case the *positive* diagram will contain different information than the diagram. To see this, consider two elements of the partial order which are not related, i.e. $a \not< b$. So neither $a < b$ nor $b < a$ will be in the positive diagram, but we could have some other structure $\mathfrak{B}$ such that $\mathfrak{B}$ satisfies all sentences of $\mathrm{diag}^+(\mathfrak{A})$ *and* $a < b$. In which case the positive diagram of $\mathfrak{B}$ will strictly contain $\mathrm{diag}^+(\mathfrak{A})$.

Of course we would like to say that $\mathfrak{B}$ looks more or less the same as $\mathfrak{A}$ if $\mathfrak{B} \models \mathrm{diag}(\mathfrak{A})$. Indeed there is a relation. To demonstrate it we first need a lemma, which states that homomorphisms commute with arbitrary terms.

Lemma. *Let* $\rho : \mathfrak{A} \to \mathfrak{B}$ *be a homomorphism of* τ*-structures, and* $t(x_0, \ldots, x_{n-1}) \in \mathcal{L}(\tau_{\{x_1,\ldots,x_n\}})$ *and* $a_0, \ldots, a_{n-1} \in A$. *Then* $t(\rho(a_0), \ldots, \rho(a_{n-1})) = \rho(t(a_0, \ldots, a_{n-1}))$.

Proof. We proceed by induction on the complexity of t.

- if t is $c \in \mathcal{C}_\tau$ then $\rho(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$ since ρ is a homomorphism.
- if t is x_i $(0 \leqslant i \leqslant n-1)$ then $\rho(t(a_0, \ldots, a_{n-1})) = \rho(a_i) = t(\rho(a_0), \ldots, \rho(a_{n-1}))$.
- if t is $f(t_0, \ldots, t_{n-1})$ and the statement is true for $t_0, \ldots, t_{n-1}$ then

$$\begin{aligned}
\rho(t^{\mathfrak{A}}(\bar{a})) &= \rho(f^{\mathfrak{A}}(t_0^{\mathfrak{A}}(\bar{a}), \ldots, t_{n-1}^{\mathfrak{A}}(\bar{a}))) \\
&= f^{\mathfrak{B}}(\rho(t_0^{\mathfrak{A}}(\bar{a})), \ldots, \rho(t_{n-1}^{\mathfrak{A}}(\bar{a}))) \\
&= f^{\mathfrak{B}}(t_0^{\mathfrak{B}}(\rho\bar{a})), \ldots, t_{n-1}^{\mathfrak{B}}(\rho\bar{a})) \\
&= t^{\mathfrak{B}}(\rho\bar{a})
\end{aligned}$$

□

With the lemma at hand we can now state the relationship between structures and satisfying the (positive) diagram.

Proposition. *Let* $\mathfrak{A}$ *be a* τ*-structure. The following are equivalent for a* τ*-structure* $\mathfrak{B}$*:*

- *The exists an expansion of* $\mathfrak{B}$ *to* $\mathfrak{B}'$ *in* τ_A *such that* $\mathfrak{B}' \models \mathrm{diag}^+(\mathfrak{A})$.
- *There exists a homomorphism* $\rho : \mathfrak{A} \to \mathfrak{B}$.

Proof. "$\Rightarrow$" Let $\mathfrak{B}'$ be an expansion to τ_A such that $\mathfrak{B}' \models \mathrm{diag}^+(\mathfrak{A})$. Define $\rho : A \to B$ by $\rho(a) = a^{\mathfrak{B}'}$. We check that ρ is a homomorphism.

- For $c \in \mathcal{C}_\tau$. We must show that $\rho(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$. By definition $\rho(c^{\mathfrak{A}}) = (c^{\mathfrak{A}})^{\mathfrak{B}'}$. Now consider the sentence φ which is $c^{\mathfrak{A}} = c$ (this is an atomic formula in $\mathcal{L}(\tau_A)$. Now $\mathfrak{A}_A \models \varphi$ since $(c^{\mathfrak{A}})^{\mathfrak{A}} = c^{\mathfrak{A}}$. So $\varphi \in \mathrm{diag}^+(\mathfrak{A})$ which, by assumption, implies $\mathfrak{B}' \models \varphi$.
- For $f \in \mathcal{F}_\tau$ with $\mathrm{arity}(f) = n$ and $a_1, \ldots, a_n \in A$ then the formula $\psi : f(a_1, \ldots, a_n) = b$ with $b := f^{\mathfrak{A}}(a_1, \ldots, a_n)$. Now $\mathfrak{A}_A \models \psi$ i.e. $\psi \in \mathrm{diag}^+(\mathfrak{A})$. But then by assumption $\mathfrak{B}' \models \psi$ i.e. $f^{\mathfrak{B}'}(a_1^{\mathfrak{B}'}, \ldots, a_n^{\mathfrak{B}'}) = b^{\mathfrak{B}'}$, thus
$$f^{\mathfrak{B}}(\rho(a_1), \ldots, \rho(a_n)) = \rho(b) = \rho(f^{\mathfrak{A}}(a_1, \ldots, a_n)).$$
- For $R \in \mathcal{R}_\tau$ with $\mathrm{arity}(R) = n$ and $a_1, \ldots, a_n \in A$ such that $R^{\mathfrak{A}}(a_1, \ldots, a_n)$. Now let $\theta : R(a_1, \ldots, a_n) \in \mathcal{L}(\tau_A)$. Then $\theta \in \mathrm{diag}^+(\mathfrak{A})$ and so $\mathfrak{B}' \models \theta$ so $R^{\mathfrak{B}'}(a_1^{\mathfrak{B}'}, \ldots, a_n^{\mathfrak{B}'})$ which by definition is $R^{\mathfrak{B}'}(\rho(a_1), \ldots, \rho(a_n))$ so $R^{\mathfrak{B}}(\rho(a_1), \ldots, \rho(a_n))$.

Thus ρ is indeed a homomorphism.

"$\Leftarrow$" Let $\rho : \mathfrak{A} \to \mathfrak{B}$ be a homomorphism. We expand $\mathfrak{B}$ to $\mathfrak{B}'$ in τ_A by setting $a^{\mathfrak{B}'} = \rho(a)$. This is clearly an expansion of $\mathfrak{B}$ to τ_A. We must show that $\mathfrak{B}' \models \mathrm{diag}^+(\mathfrak{A})$. We do this by induction on atomic formulae.

- Suppose $t, s \in \mathcal{T}(\tau_A)$ and $\mathfrak{A}_A \models s{=}t$. Now both s and t are closed terms in τ_A so there are $\tilde{s}$ and $\tilde{t}$ in $\mathcal{T}(\tau_{\{x_i : i \in \omega\}})$ such that $s^{\mathfrak{A}_A} = \tilde{s}^{\mathfrak{A}}(\bar{a})$ and $t^{\mathfrak{A}_A} = \tilde{t}^{\mathfrak{A}}(\bar{a})$ for some $\bar{a} \in A^n$. By assumption $\tilde{s}^{\mathfrak{A}}(\bar{a}) = \tilde{t}^{\mathfrak{A}}(\bar{a})$. Now by the lemma preceding this proposition, ρ commutes with terms and so
$$\rho(\tilde{s}^{\mathfrak{A}}(\bar{a})) = \tilde{s}^{\mathfrak{B}}(\rho\bar{a}) = \tilde{s}^{\mathfrak{B}}(\bar{a}^{\mathfrak{B}'}).$$
but this is the same as $\tilde{s}^{\mathfrak{B}'}(\bar{a}^{\mathfrak{B}'})$ since the interpretation of function symbols in $\mathfrak{B}$ doesn't change under the extension to $\mathfrak{B}'$. Likewise $\rho(\tilde{t}^{\mathfrak{A}}(\bar{a})) = \tilde{t}^{\mathfrak{B}'}(\bar{a}^{\mathfrak{B}'})$. Now since ρ is a function and $\tilde{s}^{\mathfrak{A}}(\bar{a}) = \tilde{t}^{\mathfrak{A}}(\bar{a})$ we have $\tilde{s}^{\mathfrak{B}'}(\bar{a}^{\mathfrak{B}'}) = \tilde{t}^{\mathfrak{B}'}(\bar{a}^{\mathfrak{B}'})$, i.e. $s^{\mathfrak{B}'} = t^{\mathfrak{B}'}$, so $\mathfrak{B}' \models s{=}t$ as well.
- Similar reasoning applies to atomic formulae given by relation symbols and so by induction $\mathfrak{B}' \models \mathrm{diag}^+(\mathfrak{A})$.

□

We have a similar but stronger relationship for the diagram. First we define the notion of *embedding* to be a morphism which respects negations of relations.

Definition. A homomorphism $\iota : \mathfrak{A} \to \mathfrak{B}$ is an **embedding** if it is injective and if $(a_1, \ldots, a_n) \in R^{\mathfrak{A}}$ if and only if $(\iota(a_1), \ldots, \iota(a_n)) \in R^{\mathfrak{B}}$.

Proposition. *Let $\mathfrak{A}$ be a τ-structure. The following are equivalent for a τ-structure $\mathfrak{B}$.*

1. *There exists an expansion $\mathfrak{B}'$ of $\mathfrak{B}$ to τ_A such that $\mathfrak{B}' \models \mathrm{diag}(\mathfrak{A})$.*

2. *There exists an embedding* $\iota : \mathfrak{A} \to \mathfrak{B}$.
3. *There exists a substructure* $\mathfrak{A}' \subseteq \mathfrak{B}$ *such that* $\mathfrak{A} \cong \mathfrak{A}'$.

Proof. First note that $2 \Leftrightarrow 3$ by definition. Most of the rest of the proof is done exactly as before when looking at the positive diagram. In going from 1 to 2 we take $\mathfrak{B}' \models \mathrm{diag}(\mathfrak{A})$ and define $\rho : \mathfrak{A} \to \mathfrak{B}$ by $a \mapsto a^{\mathfrak{B}'}$ and check (like above) that this is a homomorphism. Now ρ will be injective since, if $a \neq b$ in A then it follows that $\mathfrak{A}_A \models \neg(a = b)$, i.e. $\neg(a = b) \in \mathrm{diag}(\mathfrak{A})$. Thus, $\mathfrak{B}' \models \neg(a = b)$ and so $\rho(a) \neq \rho(b)$. The rest of the proof is much the same as before. □

These propositions show the first steps of how model theory works by going between syntax and semantics. We can convert properties which are purely structural into statements about satisfying certain sorts of formulae and sentences.

What we have called the diagram might also be called the **quantifier-free diagram**, since we only include sentences without quantifiers. If we want more information about the structure we can also look at the **elementary diagram**, $\mathrm{eldiag}(\mathfrak{A})$ which by definition is $\mathrm{Th}(\mathfrak{A}_A)$.

3.2 Canonical Models

We now prove a result about the existence of a model of a theory in much the same way as with the term algebra. We take a theory where we would like to find a model and basically just letting the language serve this goal.

Definition. A set T of $\mathcal{L}(\tau)$-sentences is **=-closed** if

- for all closed terms t, s in $\mathcal{T}(\tau)$ and for all formulae φ with one free variable x, if $\varphi(t) \in T$ and if $t = s \in T$ then $\varphi(s) \in T$,
- for all closed terms t we have $t = t \in T$.

Remark. If S is any set of $\mathcal{L}(\tau)$-sentences then there is a smallest =-closed set $\tilde{S}$ containing S.

So the following proposition could be applied to any set of $\mathcal{L}(\tau)$-sentences by passing to the =-closure of the given set first.

Proposition. *If* T *is an =-closed set of atomic sentences then there exists a structure* $\mathfrak{A}$ *such that* $\mathfrak{A} \models T$ *and such that for any* $\mathfrak{B}$ *with* $\mathfrak{B} \models T$ *there is a unique homomorphism* $\mathfrak{A} \longrightarrow \mathfrak{B}$.

Remark. If $T = \varnothing$ then $\mathfrak{A}$ will be the term algebra.

Proof. The domain of $\mathfrak{A}$ will be $\mathcal{T}(\tau)$ modulo the equivalence relation given by $s \sim t$ if and only if $s = t \in T$. Let us show that this is indeed an equivalence relation.

- *Reflexivity*: For all t we have $t \sim t$ since by assumption $t = t \in T$.
- *Symmetry*: Suppose $s \sim t$ so that $s = t \in T$. Consider the formula $\varphi(x)$ given by $x = s$. φ is an atomic formula with one free variable, x. Now $\varphi(s)$ is in T and so by $=$-closedness of T we have $\varphi(t)$ in T, i.e. $t = s \in T$ and so $t \sim s$.
- *Transitivity*: Suppose $s \sim t$ and $t \sim r$. Let $\varphi(x)$ be $x = r$. Then $\varphi(t) \in T$ and since $t \sim s$ by symmetry we have that $\varphi(s) \in T$ so $s \sim r$.

Thus, $\sim$ is an equivalence relation. We now let the domain of $\mathfrak{A}$ be $A := \mathcal{T}(\tau)/\sim$, and denote the equivalence class containing t by $[t]_\sim$. To define the τ-structure on $\mathfrak{A}$ we set

- for $c \in \mathcal{C}_\tau$, $c^{\mathfrak{A}} = [c]_\sim$
- for $f \in \mathcal{F}_\tau$ of arity n we define $f^{\mathfrak{A}}([t_0]_\sim, \dots, [t_{n-1}]_\sim) = [f(t_0, \dots, t_{n-1}]_\sim$
- for $R \in \mathcal{R}_\tau$ of arity n then $([t_0]_\sim, \dots, [t_{n-1}]_\sim) \in R^{\mathfrak{A}}$ if and only if $R(t_0, \dots, t_{n-1}) \in T$.

We must show that these definitions are well-defined and that $\mathfrak{A}$ has the desired properties. For constants there is no problem. But for an n-ary function symbol $f \in \mathcal{F}_\tau$ we must show that the value af $f^{\mathfrak{A}}$ does not depend on the choice of representatives. The same goes for relation symbols. Suppose $(t_0, \dots, t_{n-1})$ is a sequence of terms which are equivalent, coordinate-wise, to $(s_0, \dots, s_{n-1})$. Then we must show that $f^{\mathfrak{A}}([t_0]_\sim, \dots, [t_{n-1}]_\sim) = f^{\mathfrak{A}}([s_0]_\sim, \dots, [s_{n-1}]_\sim)$. To see this we use that

$$f(s_0, \dots, s_{n-1}) = f(s_0, \dots, s_{n-1}) \in T$$

and that T is $=$-closed so since $s_0 \sim t_0$ we have

$$f(t_0, s_1, \dots, s_{n-1}) = f(s_0, \dots, s_{n-1}) \in T$$

and applying this n-times we get

$$f(t_0, \dots, t_{n-1}) = f(s_0, \dots, s_{n-1}) \in T.$$

Similarly, suppose $R \in \mathcal{R}_\tau$ is an n-ary relation symbol, and suppose $R^{\mathfrak{A}}([s_0]_\sim, \dots, [s_{n-1}]_\sim)$ then by succesively substituting t_i's for s_i's we will see that $R^{\mathfrak{A}}([t_0]_\sim, \dots, [t_{n-1}]_\sim)$ also holds. So $\mathfrak{A}$ is now an $\mathcal{L}$-structure.

To show that T is exactly the set of atomic sentences that are satisfied by $\mathfrak{A}$ we use induction on the complexity of atomic sentences. For the case $t = s$ we have

$$\begin{aligned} \mathfrak{A} \models s = t \quad & \text{iff} \quad s^{\mathfrak{A}} = t^{\mathfrak{A}} \\ & \text{iff} \quad [s]_\sim = [t]_\sim \\ & \text{iff} \quad s = t \in T \end{aligned}$$

and similarly, for the case $R(t_0, \ldots, t_{n-1})$ we have

$$\begin{aligned} \mathfrak{A} \models R(t_0, \ldots, t_{n-1}) \quad &\text{iff} \quad R^{\mathfrak{A}}(t_0^{\mathfrak{A}}, \ldots, t_{n-1}^{\mathfrak{A}}) \\ &\text{iff} \quad R(t_0, \ldots, t_{n-1}) \in T. \end{aligned}$$

Now for the final claim, that all elements of $\mathfrak{A}$ have the form $t^{\mathfrak{A}}$ one uses induction on the complexity of terms to show that $[t]_{\sim} = t^{\mathfrak{A}}$. This is clear from the above construction. □

Proposition. *Let* T *be an arbitrary set of atomic sentences. Then there is a structure* $\mathfrak{A}$ *such that*

1.) $\mathfrak{A} \models T$
2.) *Every* $x \in \mathrm{dom}(\mathfrak{A})$ *is of the form* $t^{\mathfrak{A}}$ *for some* $\mathcal{L}$*-term.*
3.) *If* $\mathfrak{B} \models T$ *then there is a unique homomorphism* $f : \mathfrak{A} \to \mathfrak{B}$.

Proof. For *(1.)* and *(2.)* take the =-closure of T and apply the above lemma. *(3.)* follows from the diagram lemma proved last time. □

By *(3.)* of the proposition, $\mathfrak{A}$ is an initial object in the category of models of T.

Example. If $T = \varnothing$ and $\tau = \{f\}$ is a binary function, then we cannot form any closed terms and so cannot form any sentences.

Example. If F is a field and $p(x) \in F[x]$ is an irreducible polynomial over F then considering F[x] as an $\mathcal{L}(\tau_{rings} \cup \{c_i\}_{i \in F[x]})$-structure take

$$T = \{\text{ equations true in } F[x]\}.$$

Then F[x] is the initial structure in the category of T-models. Now consider the enlarged collection $T \cup \{p(x) = 0\}$ and take the =-closure. The initial model for this collection will yield the ring $\mathfrak{A} = F[x]/(p(x))$ where we have added a root to $p(x)$. Moreover we get the quotient map $F[x] \to F[x]/(p(x))$.

Lecture 4

4.1 Relations defined by atomic formulae

Given an $\mathcal{L}$-structure $\mathfrak{A}$ with domain A, and $\varphi(x_0, \ldots, x_{n-1})$ an atomic $\mathcal{L}$-formula we define $\varphi(A^n)$ to be $\{\bar{a} \in A^n : \mathfrak{A} \models \varphi(\bar{a})\}$. We can also allow parameters; if $\psi(\bar{x}, \bar{y})$ is an atomic formula and $\bar{b} \in A^m$ then

$$\psi(A^n, \bar{b}) = \{\bar{a} \in A^n : \mathfrak{A} \models \psi(\bar{a}, \bar{b})\}.$$

4.2 Infinitary languages

Given a signature τ we now define the *infinitary* language $\mathcal{L}_{\infty\omega}$ associated to τ. Roughly speaking the two subscripts describe how many conjunction/disjuntions we are allowed to use and how many quantifications we are allow. The first subscript '∞' indicates that we will allow infinitely many conjunctions and disjunctions. The second subscript 'ω' indicates that we will allow only finitely many quantifiers in a row.

The symbols of $\mathcal{L}_{\infty\omega}$ are all symbols from the signature τ together with the usual logical symbols:

$$=, \neg, \bigwedge, \bigvee, \forall, \exists.$$

The terms, atomic formulae, and literals are defined in the same way as before (i.e. for first-order logic).

Definition. $\mathcal{L}_{\infty\omega}$ is the smallest class such that

- all atomic formulae are in $\mathcal{L}_{\infty\omega}$
- if $\varphi \in \mathcal{L}_{\infty\omega}$ then $\neg\varphi \in \mathcal{L}_{\infty\omega}$
- if $\Phi \subseteq \mathcal{L}_{\infty\omega}$ then $\bigvee \Phi$ and $\bigwedge \Phi$ are in $\mathcal{L}_{\infty\omega}$
- if $\varphi \in \mathcal{L}_{\infty\omega}$ then $\forall x\varphi$ and $\exists x\varphi$ are in $\mathcal{L}_{\infty\omega}$

Remark. We are allowing $\Phi \subseteq \mathcal{L}_{\infty\omega}$ to be an *arbitrary* subset, so we are allowing arbitrary conjunctions and disjunctions, contrary to the case for the usual first-order logic.

Given an $\mathcal{L}$-structure $\mathfrak{A}$ (with domain A) we can now extend the notion of satisfaction "$\models$" to arbitrary formulae of $\mathcal{L}_{\infty\omega}$;

- For atomic formulae the $\models$ relation is the same as before.
- Given $\varphi(\bar{x}) \in \mathcal{L}_{\infty\omega}$ then $\mathfrak{A} \models \neg\varphi(\bar{a})$ if and only if it is not the case that $\mathfrak{A} \models \varphi(\bar{a})$.
- Given $\Phi(\bar{x}) \subseteq \mathcal{L}_{\infty\omega}$ then $\mathfrak{A} \models \bigwedge \Phi(\bar{a})$ if and only if, for all $\varphi(\bar{x}) \in \Phi(\bar{x})$ $\mathfrak{A} \models \varphi(\bar{a})$.
- Given $\Phi(\bar{x}) \subseteq \mathcal{L}_{\infty\omega}$ then $\mathfrak{A} \models \bigvee \Phi(\bar{a})$ if and only if, for at least one of $\varphi(\bar{x}) \in \Phi(\bar{x})$ we have $\mathfrak{A} \models \varphi(\bar{a})$.
- Given $\varphi(y, \bar{x}) \in \mathcal{L}_{\infty\omega}$, then $\mathfrak{A} \models \forall y \varphi(y, \bar{a})$ if and only if for all $b \in A$ we have $\mathfrak{A} \models \varphi(b, \bar{a})$.
- Given $\varphi(y, \bar{x}) \in \mathcal{L}_{\infty\omega}$, then $\mathfrak{A} \models \exists y \varphi(y, \bar{a})$ if and only if for at least one $b \in A$ we have $\mathfrak{A} \models \varphi(b, \bar{a})$.

Now we say that **first-order logic** is the language $\mathcal{L}_{\omega\omega}$ where we allow only finite subsets Φ (in other words we have only finite conjunctions and disjunctions), and only finitely many quantifiers. In general for some cardinal κ we get a language $\mathcal{L}_{\kappa\omega}$ where we allow the subsets $\Phi \subseteq \mathcal{L}_{\kappa\omega}$ to have size $< \kappa$.

In model theory we most often either work within $\mathcal{L}_{\omega\omega}$ or with $\mathcal{L}_{\omega_1\omega}$. The latter language allows *countably* many conjunctions and disjunctions. There are however several properties of first-order logic that the infinitary logics fail to have. Some of these are demonstrated by the following suggested exercises.

Exercise. Give an example of an $\mathcal{L}_{\omega_1\omega}$ sentence Φ such that every finite subsentence og Φ is satisfiable, but Φ is not. (So compactness fails).

Exercise. Axiomatize the following classes of structures with some single sentence in some language using $\mathcal{L}_{\omega_1\omega}$:

- Torsion-free abelian groups.
- Finitely generated fields.
- Linear orders isomorphic to $(\mathbb{Z}, <)$.
- Connected graphs.
- Finite valence graphs.
- Cycle-free graphs.

Exercise. Give an example of a countable language $\mathcal{L}$ and an $\mathcal{L}_{\omega_1\omega}$ sentence Φ such that every model of Φ has cardinality at least $2^{\aleph_0}$. (So Downward Löwenheim-Skolem fails).

4.3 Axiomatization

Definition. A class of $\mathcal{L}$-structures $\mathcal{K}$ is **axiomatizable** if there is some $\mathcal{L}$-theory T such that the class of $\mathcal{L}$-structures satisfying T is $\mathcal{K}$. $\mathcal{K}$ is **$\mathcal{L}$-definable** if we can take $T = \{\varphi\}$ for some $\mathcal{L}$-sentence φ.

The following lemma is important.

Lemma. *Let $\mathfrak{A}$ be an $\mathcal{L}$-structure and $X \subseteq \mathrm{dom}(\mathfrak{A})$ and Y some relation defined by a formula with parameters from X. Then if $f \in \mathrm{Aut}(\mathfrak{A})$ (the group of $\mathcal{L}$-structure automorphisms of $\mathfrak{A}$) fixes X point-wise then f fixes Y set-wise (i.e. $f(Y) = Y$).*

In other words definable sets are invariant under those automorphisms which fix the parameter space. For instance if a set Y is definable without parameters then $Y = f(Y)$ for *every* automorphism. This puts restrictions on the definable sets.

The Arithmetical Hierarchy

The theory of arithmetic is the theory of the structure $\mathbb{N} = (\omega, 0, 1, +, \cdot, <)$.

Definition. Let $\exists x < y\ \varphi$ and $\forall x < y\ \varphi$ be abbreviations of the formulae $\exists x(x < y \wedge \varphi)$ and $\forall x(x < y \to \varphi)$ respectively. These are called **bounded quantifiers**.

Definition. The arithmetic hierarchy is the following hierarchy of subsets of ω.

- φ is Σ^0_0 and Π^0_0 if all quantifiers are bounded.
- φ is in Σ^0_{n+1} if $\varphi = \exists\bar{x}\psi$ for some ψ in Π^0_n.
- φ is in Π^0_{n+1} if $\varphi = \forall\bar{x}\psi$ where ψ is in Σ^0_n.

So the subscript of Σ^0_n and Π^0_n is the number of alterations of (unbounded) quantifiers appearing in the formulae. It can in fact be shown that this hierarchy is proper, i.e. the inclusions $\Sigma^0_n \subseteq \Sigma^0_{n+1}$ are proper for all $n \in \omega$.

Lecture 5

We start with some examples of first-order theories.

Example (Peano Arithmetic). Peano arithmetic is intended as a formalization of the laws of arithmetic on the natural numbers. Let $\tau = \{0, 1, +, \cdot, <\}$ where 0 and 1 are constants, $+$ and $\cdot$ are binary function symbols, and $<$ is a binary relation. The axioms of this theory can be presented in many different ways. There are basically two sorts of axioms: the ones describing the algebraic properties of the natural numbers and the ones describing induction. The algebraic axioms basically state (in the formal language $\mathcal{L}(\tau)$) that $(\mathbb{N}, 0, 1, +, \cdot, <)$ is a discretely ordered semi-ring. They may be stated as follows.

- $\forall x\ x + 0 = x$
- $\forall x \forall y\ (x + y) + 1 = x + (y + 1)$
- $\forall x \forall y\ [(x + 1 = y + 1) \to x = y]$
- $\forall x\ \neg(x + 1 = 0)$
- $\forall x\ x \cdot 0 = 0$
- $\forall x \forall y\ x \cdot (y + 1) = x \cdot y + x$
- $\forall x\ \neg(x < x)$
- $\forall x \forall y\ (x < y) \vee (x = y) \vee (y < x)$
- $\forall x \forall y \forall z\ (x < y \wedge y < z \to x < y)$
- $\forall x \forall y \forall z\ x < y \to x + z < y + z$
- $\forall x \forall y\ (x < y + 1) \to (x < y \vee x = y)$

This takes care of the algebraic axioms. Note that the above list is a finite list of sentences, and so we could take the conjunction over all of them and write them as a single sentence in $\mathcal{L}(\tau)$. This is not the case for the induction axioms. Induction is given by a *schema* of axioms. For each formula $\varphi(x, \bar{y})$ in $\mathcal{L}(\tau)$ we have the axiom, $I(\varphi)$:

$$\forall \bar{y}\ \ ([\varphi(0, \bar{y}) \wedge \forall z\ (\varphi(z, \bar{y}) \to \varphi(z + 1, \bar{y}))] \to \forall x\ \varphi(x, \bar{y}))$$

i.e. if $\varphi(x, \bar{y})$ is true of 0 and if each time $\varphi(z, \bar{y})$ is true then so is $\varphi(z + 1, \bar{y})$, then $\varphi(x, \bar{y})$ is true for all x.

The collection of all these infinitely many axioms (both the algebraic and the inductive) is called the theory of **Peano arithmetic**, or PA. It is of course meant to axiomatize $\mathbb{N}$. There are however many other *non-standard* models of Peano arithmetic. A theorem of Tenenbaum however says that no other models of Peano arithmetic has a recursive presentation. In a sense this means that one will never "see" any of the non-standard models of PA.

Example (Orders). Let $\tau = \{<\}$ be the signature of a single binary relation symbol. The **partial orders** are τ-structures satisfying the axioms

- $\forall x\, \neg(x < x)$
- $\forall x \forall y \forall z\, (x < y \wedge y < z) \to (x < z)$

The theory of **linear orders** is a sub theory of the theory of partial orders, i.e. it contains the two axioms above and the extra axiom

$$\forall x \forall y (x < y \vee x = y \vee y < x).$$

The theory of **well-orders** is the theory of linear orders together with the statement that

for every nonempty subset X there exists a least element $a \in X$.

This last statement is not a first-order statement since we are quantifying over both subsets and elements of the subsets. So this is a second-order statement and can be made rigorous in second-order logic.

Exercise. Show that the class of well-orders is axiomatizable in $\mathcal{L}_{\omega_1 \omega_1}$. [Hint: $(A, <)$ is well-ordered iff there are no strict descending chains.]

Example (ZFC). **Zermelo-Fraenkel set theory with choice** (ZFC) is an axiomatization system for doing set theory. As with PA, ZFC is usually given by an axiom schema. In fact ZFC is not finitely axiomatizable.

Example (ACF). **Algebraically Closed Fields** (ACF). This is again given by a schema of axioms which express that every monic polynomial, of degree m (for each $m \in \omega$) with coefficients in the field, has a solution. ACF is in fact not finitely axiomatizable, but there is an open question which asks whether ACF is finite-variable axiomatizable. I.e. if we allow only finitely many variables in the construction of the language, can we axiomatize ACF?

5.1 Preservation of Formulae

Fix a signature τ. We work in $\mathcal{L}_{\omega\omega}$ although some of the following makes sense in higher-order logics. We start with the $\forall_n$ (read: "A n") and $\exists_n$ (read: "E

n") hierarchy of formulae. The subscripts refer to the number of alterations of quantifiers there are in a given formula.

Definition. The class of $\forall_0$ formulae is the same as the class of $\exists_0$ formulae and they are the quantifier-free formulae. A formula is $\forall_{n+1}$ if it has the form

$$\forall \bar{y} \left(\bigvee \bigwedge \Phi \right)$$

where every $\varphi \in \Phi$ is an $\exists_n$-formula.
A formula is $\exists_{n+1}$-formula if it has the form

$$\exists \bar{y} \left(\bigvee \bigwedge \Phi \right)$$

where each $\varphi \in \Phi$ is a $\forall_n$-formula.

Note that $\bigcup_n \forall_n = \bigcup_n \exists_n$.

Remark. Often one says that a formula which is "equivalent" to a formula in a given class is in that class. I.e. if φ is "equivalent" to a formula ψ which is $\forall_n$ then we may say that φ is $\forall_n$ as well. Here "equivalent" means equivalence modulo some implicit background theory T. I.e. φ and ψ are **equivalent modulo** T if for all $\mathfrak{A}$ such that $\mathfrak{A} \models T$ then $\mathfrak{A} \models \varphi \leftrightarrow \psi$.

Remark. The classification given above is similar to the arithmetical hierarchy in recursion theory (see lecture 4), where statements are divided into Σ^0_n and Π^0_n classes. However, these hierarchies are different, namely in the lowest level $\Sigma^0_0 = \Pi^0_0$ allows the use of *bounded quantifiers*. Bounded quantifiers are *not* allowed in the $\forall_0$ and $\exists_0$ formulae.

A formula is **prenex** if it consists of a (possibly empty) string of quantifiers followed by a quantifier-free formula.

Proposition. *For every $\psi \in \mathcal{L}(\tau)$ there exists $\theta \in \bigcup_n \forall_n = \bigcup \exists_n$ such that $\psi \leftrightarrow \theta$. In words; every formula is equivalent to a formula in prenex normal form.*

Proof. We work by induction on the complexity of ψ.
If ψ is atomic then it is $\forall_0$ (and $\exists_0$) already. If $\psi = \neg\theta$ then by induction θ is in, say, $\exists_n$ so θ has the form $\exists \bar{y} \tilde{\theta}$ with $\tilde{\theta}$ in $\forall_{n-1}$, so $\neg\theta$ is equivalent to $\forall \bar{y} \neg \tilde{\theta}$.
If $\psi = (\varphi \vee \theta)$ then we may assume $\varphi \in \forall_n$ and $\theta \in \forall_n$ then by definition ψ is in $\forall_n \subseteq \exists_{n+1}$ (since we can always put irrelevant quantifiers in from of a formula).
Likewise for $\psi = \exists\theta$, we may assume $\theta \in \forall_n$ then $\psi \in \exists_{n+1}$. By induction we are done. $\square$

We shall see later that this hierarchy is in fact proper, in the sense that the inclusions $\forall_n \subseteq \exists_n$ and $\exists_n \subseteq \forall_{n+1}$ are proper.

Remark. This is related to Hilbert's 10th problem; Find an algorithm to decide for $p(x) \in \mathbb{Z}[x]$ whether there exists $\bar{a} \in \mathbb{Z}^n$ such that $p(\bar{a}) = 0$. The Matiyasevich-Davis-Putnam-Robinson (MDPR) theorem states that no such algorithm exists. This problem can then be asked for polynomials over the rational. This is an open problem. But, a recent theorem of Jochen Koenigsmann states that there is a *universal definition* of the integers inside the rationals. I.e. a $\forall_1$-formula $\theta(x)$ such that $\mathbb{Q} \models \theta(a)$ if and only if $a \in \mathbb{Z}$. If there were an $\exists_1$-definition then the MDPR-theorem would imply that there is no algorithm to decide over $\mathbb{Q}$ either.
Later in the course we will prove the following fact: There exist two models $\mathfrak{A}$ and $\mathfrak{B}$ of the theory of the rational, $(\mathbb{Q}, +, \cdot, 0, 1)$, such that $\mathfrak{A} \subseteq \mathfrak{B}$ and there is some polynomial over $\mathfrak{A}$ which has no solutions over $\mathfrak{A}$ but does have solutions over $\mathfrak{B}$.

We shall now look at what kinds of formulae are preserved by certain types of maps.

Proposition (Going-up)**.** *If $\iota : \mathfrak{A} \to \mathfrak{B}$ is an embedding and $\varphi(\bar{x})$ is $\exists_1$, then $\mathfrak{A} \models \varphi(\bar{a})$ implies that $\mathfrak{B} \models \varphi(\iota\bar{a})$. Equivalently if $\mathfrak{A} \subseteq \mathfrak{B}$ then $\mathfrak{B} \models (\iota\bar{a})$.*

Notation. If $\varphi(\bar{x})$ is a formula and $\bar{a}$ is a tuple of the same length as $\bar{x}$ then we write $\varphi(\bar{a}/\bar{x})$ for the formula where we have substituted $\bar{a}$ for $\bar{x}$.

Proof. Write φ as $\exists\bar{y}(\bigvee\bigwedge\Psi)$ where Ψ is a set of $\forall_0$ (i.e. quantifier-free) formulae. Then if $\mathfrak{A}_{\bar{a}} \models \varphi(\bar{a})$ then there exist $\bar{b}$ from $\mathfrak{A}$ such that $\mathfrak{A}_{\bar{a},\bar{b}} \models \bigvee\bigwedge\Psi(\bar{a}/\bar{x}, \bar{b}/\bar{y})$. We have already shown that if $\theta(\bar{z})$ is quantifier-free and $\mathfrak{A} \subseteq \mathfrak{B}$ then for all $\bar{c}$ in $\mathfrak{A}$ we have $\mathfrak{A} \models \theta(\bar{c})$ if and only if $\mathfrak{B} \models \theta(\bar{c})$. So $\mathfrak{B} \models \bigvee\bigwedge\Psi(\bar{a}/\bar{x}, \bar{b}/\bar{y})$, i.e. $\mathfrak{B} \models \exists\bar{y}\bigvee\bigwedge\Psi(\bar{a}/\bar{x}, \bar{y})$. □

If we weaken the hypothesis and assume only that there is a homomorphism between $\mathfrak{A}$ and $\mathfrak{B}$ we can still get a result. We call a formula $\exists_1^+$ if no negations are involved, i.e. if it has the form $\exists\bar{y}(\bigvee\bigwedge\Phi)$ where all elements of Φ are atomic.

Proposition. *If $\rho : \mathfrak{A} \to \mathfrak{B}$ is a τ-homomorphism and φ is $\exists_1^+$, then $\mathfrak{A}_{\bar{a}} \models \varphi(\bar{a})$ implies $\mathfrak{B}_{\rho\bar{a}} \models \varphi(\rho\bar{a})$.*

Proof. Immediate from the definition of homomorphism. □

The "Going-up" proposition has a dual "Going-down" proposition.

Proposition (Going-down)**.** *If $\mathfrak{A} \subseteq \mathfrak{B}$ and φ is $\forall_1$ then $\mathfrak{B} \models \varphi(\bar{a})$ implies that $\mathfrak{A} \models \varphi(\bar{a})$.*

Proof. We note that φ is equivalent to a formula of form $\neg\exists\neg$. Then apply the "going-up" proposition. $\square$

As a nice consequence of this propositions, suppose T is a theory where all axioms of T are $\forall_1$. Then the class of models is closed under formation of substructures. I.e. if $\mathfrak{A} \models T$ then all substructures of $\mathfrak{A}$ also model T. Dually, if T is a theory all of whose sentences are $\exists_1$ then, by the going-up proposition, the class of models of T is closed under formation of superstructures. I.e. if $\mathfrak{A} \models T$ and $\mathfrak{B}$ is some superstructure of $\mathfrak{A}$ then $\mathfrak{B} \models T$ as well.
We shall in fact see that these characterizations of universal and existential theories have converses. That is, if a theory T has the property that whenever $\mathfrak{A} \models T$ then for all substructures $\mathfrak{B} \subseteq \mathfrak{A}$, $\mathfrak{B} \models T$, then T is universal. Similarly if T has the property that whenever $\mathfrak{A} \models T$ and $\mathfrak{B} \supseteq \mathfrak{A}$ then $\mathfrak{B} \models T$, then T is existential.

We now turn to situations where we can preserve $\forall_2$-sentences.

Definition. A **chain** of models is a sequence $(\mathfrak{A}_i)_{i\in I}$ of τ-structures such that $(I, <)$ is totally ordered and such that

$$i < j \Rightarrow \mathfrak{A}_i \subseteq \mathfrak{A}_j.$$

Given a chain $(\mathfrak{A}_i)_{i\in I}$ of τ-structures we can form the **direct limit**[1]

$$\tilde{\mathfrak{A}} = \bigcup_{i\in I} \mathfrak{A}_i.$$

The domain of $\tilde{\mathfrak{A}}$ will be the union $\bigcup_{i\in I} \mathrm{dom}(\mathfrak{A}_i)$. The interpretations of the symbols will be as follows.

- for $c \in \mathcal{C}_\tau$ let $c^{\tilde{\mathfrak{A}}} = c^{\mathfrak{A}_i}$ for any choice of $\mathfrak{A}_i$
- for $f \in \mathcal{C}_\tau$ then $f^{\tilde{\mathfrak{A}}} = \bigcup_{i\in I} f^{\mathfrak{A}_i}$
- for $R \in \mathcal{R}_\tau$ then $R^{\tilde{\mathfrak{A}}} = \bigcup_{i\in I} f^{\mathfrak{A}_i}$.

All these choices are well-defined and since $\mathfrak{A}_i \subseteq \mathfrak{A}_j$ whenever $i < j$ we get that $\mathfrak{A}_i \subseteq \tilde{\mathfrak{A}}$ for all $i \in I$.

$\forall_2$ sentences "go up" in chains.

Proposition. *If* φ *is* $\forall_2$ *and if for all* $i \in I$ $\mathfrak{A}_i \models \varphi$ *then* $\tilde{\mathfrak{A}} \models \varphi$.

[1] It is in fact a direct limit in the category theoretic sense.

Proof. We can write φ as $\forall\bar{x}\exists\bar{y}\ \theta$ with θ quantifier-free. Let $\bar{a}$ be a sequence from $\tilde{\mathfrak{A}}$. Since $\bar{a}$ is finite there exists $i \in I$ such that $\bar{a}$ comes from $\mathrm{dom}(\mathfrak{A}_i)$. Now since $\mathfrak{A}_i \models \varphi$ we have

$$\mathfrak{A}_{i,\bar{a}} \models \exists\bar{y}\ \theta(\bar{a}/x)$$

and so by the going-up for $\exists_1$,

$$\tilde{\mathfrak{A}}_{\bar{a}} \models \exists\bar{y}\ \theta(\bar{a}/x)$$

and since this was true for any choice of $\bar{a}$ it follows that

$$\tilde{\mathfrak{A}} \models \varphi.$$

$\square$

The converse is also true, i.e. a theory T admits an $\forall_2$-axiomatization if and only if it is preserved under unions of chains.

Lecture 6

We discuss the basic idea of comparing different structures and ways of regarding the same structure in different languages. At one level this allows us to completely forget about syntax and focus on the definable sets. On another level it brings the syntax back to the fore because we will have specific ways of referring, to specific sets which might appear, as though they are actually part of the language. Hodges calls this *atomisation* although most people call it *Morleyisation*. As Hodges points out Skolem introduced the method before Morley.

Last Time

Let us first recall briefly the chain construction. We have a chain $(\mathfrak{A}_i)_{i \in I}$ of τ-structures indexed by a totally ordered set $(I, <)$. This is a functor from the category $(I, <)$ to the category $\mathrm{Str}(\tau)$ of τ-structures. The content of this rephrasing is just that and arrow $i < j$ is mapped to an arrow $\mathfrak{A}_i \subseteq \mathfrak{A}_j$. Given this chain we may form the union $\bigcup \mathfrak{A}_i$ which as its domain is the union of the domains of $\mathfrak{A}_i$ and which is given the natural τ-structure.

We proved last time that if φ is an $\forall_2$ sentence in $\mathcal{L}(\tau)$ and if for all i we have $\mathfrak{A}_i \models \varphi$ then $\bigcup \mathfrak{A}_i \models \varphi$. This proposition is a slight elaboration on the proposition that $\exists_1$ sentence "go up".

6.1 Theories and Models

Definition. If $\mathcal{K}$ is a class of τ-structures, then $\mathrm{Th}(\mathcal{K})$ is the set of all τ sentences φ such that for all $\mathfrak{A} \in \mathcal{K}$ we have $\mathfrak{A} \models \varphi$. I.e.

$$\mathrm{Th}(\mathcal{K}) := \{\varphi \in \mathcal{L}(\tau) : \varphi \text{ is a sentence } \forall \mathfrak{A} \in \mathcal{K}\ \mathfrak{A} \models \varphi\}$$

Definition. If T is a set of τ-sentences, then $\mathrm{Mod}(T)$ is the class of τ-structures $\mathfrak{A}$ such that $\mathfrak{A} \models T$. I.e.

$$\mathrm{Mod}(T) := \{\mathfrak{A} \in \mathrm{Str}(\tau) : \mathfrak{A} \models T\}$$

$$\begin{tikzcd}
\text{classes of } \tau\text{-structures} \arrow[rr, bend left, "\mathrm{Th}(-)"] & & \text{sets of } \tau\text{-sentences} \arrow[ll, bend left, "\mathrm{Mod}(-)"]
\end{tikzcd}$$

One immediately asks whether $\mathrm{Th}(-)$ and $\mathrm{Mod}(-)$ are each others inverses? They are not. But they are connected[1]. Indeed we have, by definition, that

$$\mathrm{Th}(\mathrm{Mod}(T)) \supseteq T$$

and

$$\mathrm{Mod}(\mathrm{Th}(\mathcal{K})) \supseteq \mathcal{K}.$$

Both inclusions may be strict.

Definition. Given a theory T and a sentence φ we say that T **semantically implies** φ, written $T \models \varphi$ iff any model of T is also a model of φ.

Given a theory T the set $\mathrm{Th}(\mathrm{Mod}(T))$ is the set of **semantic consequences of** T. It is the set of sentences that are satisfied by all models of T, i.e.

$$\mathrm{Th}(\mathrm{Mod}(T)) = \{\varphi : T \models \varphi\}.$$

Similarly $\mathrm{Mod}(\mathrm{Th}(\mathcal{K}))$ is the smallest definable class of structures containing $\mathcal{K}$.

Notation. If $\mathcal{K}$ is the singleton class $\{\mathfrak{A}\}$, then we write $\mathrm{Th}(\mathfrak{A})$ instead of $\mathrm{Th}(\{A\})$.

Definition. We say two structures $\mathfrak{A}$ and $\mathfrak{B}$ are **elementarily equivalent**, written $\mathfrak{A} \equiv \mathfrak{B}$, if $\mathrm{Th}(\mathfrak{A}) = \mathrm{Th}(\mathfrak{B})$.

Definition. If $\mathfrak{A} \subseteq \mathfrak{B}$ then $\mathfrak{A}$ is an **elementary substructure** of $\mathfrak{B}$, written $\mathfrak{A} \preccurlyeq \mathfrak{B}$ if the inclusion map preserves all formulae of $\mathcal{L}(\tau)$. Such an inclusion map is called an **elementary inclusion**.

See lecture 7 for an example where $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{A} \equiv \mathfrak{B}$ but $\mathfrak{A} \not\preccurlyeq \mathfrak{B}$.

Definition. We say two theories S and T are **equivalent** if $\mathrm{Mod}(T) = \mathrm{Mod}(S)$.

[1] They form a Galois connection between the "posets" $(\mathcal{P}(\mathrm{Str}(\tau)), \supseteq)$ and $(\mathcal{P}(\mathcal{L}(\tau)), \subseteq)$, except that $\mathrm{Str}(\tau)$ is not a set and so neither is $\mathcal{P}(\mathrm{Str}(\tau))$!

We may often find that $\mathrm{Mod}(\mathrm{Th}(\mathcal{K}))$ strictly contains $\mathcal{K}$. For instance if $\mathcal{K}$ is a singleton $\{\mathfrak{A}\}$ where $\mathfrak{A}$ is some τ-structure, then $\mathrm{Mod}(\mathrm{Th}(\mathcal{K}))$ will contain all those τ-structures which are elementarily equivalent to $\mathfrak{A}$. So if there exist models, $\mathfrak{B}$, that are elementarily equivalent to $\mathfrak{A}$, i.e. $\mathrm{Th}(\mathfrak{A}) = \mathrm{Th}(\mathfrak{B})$ but such that $\mathfrak{A} \not\cong \mathfrak{B}$ then $\mathrm{Mod}(\mathrm{Th}(\mathfrak{A}))$ will properly contain $\{\mathfrak{A}\}$. In fact, given τ, there *always* exist $\mathfrak{A}$ and $\mathfrak{B}$ that are not isomorphic yet elementarily equivalent.

Proposition. *There exist $\mathfrak{A}$ and $\mathfrak{B}$ two τ-structures such that $\mathfrak{A} \not\cong \mathfrak{B}$ but $\mathfrak{A} \equiv \mathfrak{B}$.*

Proof. Consider the restricted functor $\mathrm{Th} : \mathrm{Str}(\tau) \longrightarrow \mathcal{P}(\mathcal{L}(\tau))$. Now $\mathrm{Str}(\tau)$ is a class (there are as many τ-structures as there are sets) and $\mathcal{P}(\mathcal{L}(\tau))$ is a set of cardinality at most $2^{|\mathcal{L}(\tau)|}$. By the Pigeon-hole-principle this is not injective. Even considering $\mathrm{Str}(\tau)$ up to isomorphism it is still a class since isomorphism preserves cardinality. Thus there are $\mathfrak{A}$ and $\mathfrak{B}$ such that $\mathfrak{A} \not\cong \mathfrak{B}$ and yet $\mathrm{Th}(\mathfrak{A}) = \mathrm{Th}(\mathfrak{B})$. $\square$

We shall prove much stronger results than this later in the course.

The kinds of classes of structures that we will be most interested in will be those that appear as the classes of models of some theory T.

Definition. A class of τ-structures, $\mathcal{K}$, is an **elementary class** if $\mathcal{K} = \mathrm{Mod}(T)$ for some T. In this case we say that T **axiomatizes** $\mathcal{K}$.

Definition. Let $\mathcal{K}$ be a class of τ-structures. Then $\mathfrak{A} \in \mathcal{K}$ is **existentially closed in** $\mathcal{K}$ (or "e.c. in $\mathcal{K}$") if; given any $\mathfrak{B} \in \mathcal{K}$ with $\mathfrak{A} \subseteq \mathfrak{B}$ then, for every $\exists_1$ sentence ψ in $\mathcal{L}(\tau_A)$, if $\mathfrak{B} \models \psi$ then $\mathfrak{A} \models \psi$.

So a structure is existentially closed if you have already put in all the witnesses.

Theorem 6.1. *If $\mathcal{K} = \mathrm{Mod}(T)$ where T is $\forall_2$-axiomatizable, then for all $\mathfrak{A} \in \mathcal{K}$ there exists some $\mathfrak{B} \in \mathcal{K}$ such that $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{B}$ is existentially closed.*

Proof. Given $\mathfrak{A} \in \mathcal{K}$ we build a chain of models $(\mathfrak{A}_n)_{n\in\omega}$ in $\mathcal{K}$ and then take the union. Let $\mathfrak{A}_0 := \mathfrak{A}$. We construct $\mathfrak{A}_1$ as follows. Let $\{\varphi_i \in \mathcal{L}(\tau_{A_0}) : \varphi_i \text{ is } \exists_1\text{-sentence}\}$ be an enumeration of the existential sentences with parameters from $A_0 = A$. Now let $\mathfrak{A}_{1,0} := \mathfrak{A}_1$ and for φ_0 we ask whether there exists any $\mathfrak{B} \in \mathcal{K}$ with $\mathfrak{B} \models \varphi_0(\bar{a})$ and $\mathfrak{A}_{1,0} \not\models \varphi_0$, if so then let $\mathfrak{A}_{1,1} := \mathfrak{B}$. Now at stage i we ask the same question for φ_i and if $\mathfrak{A}_{1,i}$ is not existentially closed with respect to φ_i then pick some $\mathfrak{B} \in \mathcal{K}$ that witnesses this and let $\mathfrak{A}_{1,i+1} := \mathfrak{B}$. Thus we get a chain of order type equal to the order type of $\mathcal{L}(\tau_{A_0})$. We take the union of this chain. This union is $\mathfrak{A}_1$. Now $\mathfrak{A}_1 \in \mathcal{K}$ since $\forall_2$ sentences are preserved in unions of chains. Likewise at stage n we construct $\mathfrak{A}_{n+1}$ by going through all sentences

$\varphi_i(\bar{a})$ with parameters from A_n. At each stage we have $\mathfrak{A}_n \in \mathcal{K}$ and $\mathfrak{A}_n \subseteq \mathfrak{A}_{n+1}$ by construction. Now we take the union of the ω-chain $(\mathfrak{A}_n)_{n\in\omega}$

$$\mathfrak{B} := \bigcup_{n\in\omega} \mathfrak{A}_n$$

Then $\mathfrak{B}$ is again in $\mathcal{K}$ by the preservation of $\forall_2$-sentences in chains. Also $\mathfrak{B}$ is existentially closed since given any $\exists_1$ sentence φ with parameters from B then since there are only finitely many of these parameters occurring in φ we have that $\varphi = \varphi_i$ for some $\varphi_i \in \mathcal{L}(\mathfrak{A}_n)$ for some n. At stage n we ensured that $\mathfrak{A}_{n+1} \subseteq \mathfrak{B}$ is existentially closed with respect to φ. This finishes the proof. □

Example (Linear orders). Let τ be the signature $\mathcal{C}_\tau = \mathcal{F}_\tau = \varnothing$ and $\mathcal{R}_\tau = \{<\}$. Let T be the theory of linear orders.

- Let $\mathfrak{A} = (\omega, <)$ is *not* existentially closed. To see this let ψ be $\exists x(0 < x < 1)$ then take the natural extension of $\mathfrak{A}$ by adding $\frac{1}{2}$ to the set. Call this τ-structure $\mathfrak{B}$ then $\mathfrak{B} \models \psi$, and $\mathfrak{A} \subseteq \mathfrak{B}$ but $\mathfrak{A} \not\models \psi$.
- Let $\tilde{\mathfrak{A}}$ have domain $\{\frac{a}{2^n} : a \in \mathbb{N}, n \in \mathbb{N}\}$, with the natural order. Then $\tilde{\mathfrak{A}}$ is *not* existentially closed. For instance $(\mathbb{R}, <) \models \exists x\ x < 0$ and $\tilde{\mathfrak{A}} \not\models \exists x\ x < 0$.
- $(\mathbb{R}, <)$ *is* existentially closed. This requires a bit of work to show.

Example (Fields). Let τ be the signature of fields. An existentially closed field is E.C. if and only if it is algebraically closed

Example (Groups). It is difficult to describe explicitly the E.C. groups. Of course one can give examples of equations that are necessarily true in E.C. groups, for instance $\forall x \exists y\ y^n = x$.

In fact the class of E.C. groups cannot be axiomatized. We can however axiomatize the class of E.C. fields and the class of E.C. linear orders.

6.2 Unnested formulae

Definition. An **unnested atomic formula** is one of the form

- $x = c$, for $c \in \mathcal{C}_\tau$ and x a variable.
- $F\bar{x} = y$ where $F \in \mathcal{F}_\tau$ and $\bar{x}, y$ are variables.
- $R\bar{x}$, where $R \in \mathcal{R}_\tau$ and $\bar{x}$ are variables.
- $x = y$, where x and y are variables.

An **unnested formula** is built from the unnested atomic formulae by the usual rules.

Lemma. *Every formula $\varphi \in \mathcal{L}(\tau)$ is equivalent to some unnested formula $\tilde{\varphi}$. In fact if φ is atomic then $\tilde{\varphi}$ may be take to be either $\exists_1$ or $\forall_1$, and if φ is $\forall_n$ or $\exists_n$ then $\tilde{\varphi}$ may be taken to have the same quantifier complexity.*

Proof. Whenever some term contains a function symbol applied to something unnested we will strip of the function symbol and replace it by a new variable.
Let us start with φ an atomic formula. We will show that there is an equivalent existential formula where each of the sub-formulae have terms that are no more complicated than the one we had before (and at least one has complexity strictly less than before). Suppose for instance that φ is

$$R(t_0, \ldots, t_{n-1})$$

and suppose $t_0 = F(s_0, \ldots, s_k)$ where s_i are simpler terms. Then φ is equivalent to

$$\exists x_0, \ldots, x_{n-1}, y \ (y{=}F(\bar{x}) \wedge \bigwedge(x_i = s_i) \wedge R(y, t_2, \ldots, t_{n-1})).$$

φ is also equivalent to the formula

$$\forall x_0, \ldots, x_{n-1}, y \ (y{=}F(\bar{x}) \wedge \bigwedge(x_i = s_i) \longrightarrow R(y, t_2, \ldots, t_{n-1})).$$

then we complete the proof by induction. Of course one needs to do a similar reduction in the case that φ is an equality of terms, or a more general formula. $\square$

Remark. The above procedure (as described in the proof of the lemma) is analogous to a procedure in the theory of differential equations. Here one can turn an order n differential equation in one variable into an equivalent first-order differential equation in n variables. For instance given the equation

$$\sum_{i=0}^{n} a_i \frac{d^i}{dt^i} f = 0.$$

Then by defining $\bar{y} = (f, \frac{d}{dt}f, \frac{d^2}{dt^2}f, \ldots, \frac{d^{n-1}}{dt^{n-1}}f)$ we get an equivalent system of first-order differential equations

$$\sum_{i=0}^{n-1} a_i y_i + a_n \frac{d}{dt} y_{n-1} = 0$$

where

$$y_{i+1} = \frac{d}{dt} y_i.$$

Unnested formulae are useful when dealing with an interpretation of one language in another language where they allow us to deal with just the basic structure.

6.3 Definitional expansions

There are cases where, when extending the language in some sense gives no further structure, i.e. whatever new structure we get in the new language, was already there in the old language. An example will make this clear.

Example. Let $(\mathbb{R}, 0, 1, +, \cdot)$ be $\mathbb{R}$ in the signature of rings, $\tau = \{0, 1, +, \cdot\}$. In this structure the ordinary relation $x \leqslant y$ on $\mathbb{R}$ is already definable! For instance we could set

$$x \leqslant y \qquad \text{iff} \qquad \exists z\ (x + z^2 = y).$$

Now $\leqslant$ is not in the signature, but the *set* (in $\mathbb{R}^2$) given by the relation $x \leqslant y$ is definable in $(\mathbb{R}, 0, 1, +, \cdot)$. Thus extending the signature to the signature of ordered rings, i.e. $\tau^+ = \tau \cup \{\leqslant\}$ seems not to give us any new definable sets[2].

Definition. If $\tau \subseteq \tau'$ is an extension of signatures, and $\mathfrak{A}'$ is a τ'-structure and $\mathfrak{A} := \mathfrak{A}'|_\tau$ we say that $\mathfrak{A}'$ is a **definitional expansion of** $\mathfrak{A}$ if every τ'-definable set is already τ-definable.

This definition requires us to look at all τ'-definable sets. The following equivalent criterion allows us to focus on the definitions of the symbols of τ' in terms of the simpler language $\mathcal{L}(\tau)$.

Theorem 6.2. *Let $\tau \subseteq \tau'$ be an extension of signatures and let $\mathfrak{A} = \mathfrak{A}'|_\tau$. If*

- *for each $c \in \mathcal{C}_{\tau'}$ there is some $\theta_c(x) \in \mathcal{L}(\tau)$ such that*

$$\mathfrak{A} \models \theta_c(a) \qquad \textit{iff} \qquad a = c^{\mathfrak{A}'}$$

- *for each $f \in \mathcal{F}_{\tau'}$ of arity n there is some $\theta_f(\bar{x}, y)$ such that*

$$\mathfrak{A} \models \theta_f(\bar{a}, b) \qquad \textit{iff} \qquad f^{\mathfrak{A}'}(\bar{a}) = b$$

- *and for each $R \in \mathcal{R}_{\tau'}$ of arity n there is some $\theta_R(\bar{x})$ such that*

$$\mathfrak{A} \models \theta_R(\bar{a}) \qquad \textit{iff} \qquad \bar{a} \in R^{\mathfrak{A}'}.$$

Then $\mathfrak{A}'$ is a definitional expansion of $\mathfrak{A}$.

Proof. Let $\chi(\bar{x})$ be a formula in $\mathcal{L}(\tau')$. We want a formula $\tilde{\chi}(\bar{x})$ in $\mathcal{L}(\tau)$ such that $\tilde{\chi}(\mathfrak{A}') = \chi(\mathfrak{A}')$. Since every formula is equivalent to an unnested formula we may assume that χ is unnested. Now we work by induction on the complexity of χ.

[2]This is of course not a rigorous statement since there could *a priori* be some strange new definable sets when we introduce $\leqslant$ into the signature.

- The case where χ is atomic is covered immediately by the assumptions of the theorem.
- If χ is $\chi_1 \wedge \chi_2$ then by induction hypothesis χ_1 and χ_2 have equivalent forms, and so $\tilde{\chi} = \tilde{\chi}_1 \wedge \tilde{\chi}_2$. Similarly for disjunctions and negations.
- If χ is $\exists x\ \xi$ then $\tilde{\chi}$ is just $\exists x\ \tilde{\xi}$. In this last step we must be careful not to reuse variables.

By induction we are done. $\square$

Remark. The fact that we can assume that χ is unnested makes the above proof much easier since we do not need to look carefully at all possible nested terms that occur in the formulae.

Example. From the theorem it is now clear that $(\mathbb{R}, 0, 1, +, \cdot, \leqslant)$ is in fact a definable expansion of $(\mathbb{R}, 0, 1, +, \cdot)$.

Lecture 7

7.1 Definitional expansions continued

Let τ' be an extension of the signature τ. Let T be an τ-theory and T' be an τ'-theory, such that every model of T' has a reduct back to a model of T. Furthermore assume that for each new symbol in τ' we have a definition of that symbol in terms of $\mathcal{L}(\tau)$. I.e.

- For each $c \in \mathcal{C}_{\tau'} \setminus \mathcal{C}_{\tau}$ we have $\theta_c(x) \in \mathcal{L}(\tau_x)$,
- For each $f \in \mathcal{F}_{\tau'} \setminus \mathcal{F}_{\tau}$ we have $\gamma_f(\bar{x}, y)$,
- For each $R \in \mathcal{R}_{\tau'} \setminus \mathcal{R}_{\tau}$ we have $\psi_R(\bar{x})$,

and such that T says

- $\exists^{=1} x\ \theta_c(x)$ for each $c \in \mathcal{C}_{\tau'} \setminus \mathcal{C}_{\tau}$.
- $\forall \bar{x} \exists^{=1} y\ \gamma_f(\bar{x}, y)$ for each $f \in \mathcal{F}_{\tau'} \setminus \mathcal{F}_{\tau}$,

and such that T' says that these formulae formally define the constants, functions and relations, i.e.

- $\forall x \theta_c(x) \leftrightarrow x = c$ for each $c \in \mathcal{C}_{\tau'} \setminus \mathcal{C}_{\tau}$
- $\forall \bar{x} \forall y\ [\gamma_f(\bar{x}, y) \leftrightarrow f(\bar{x}) = y]$ for each $f \in \mathcal{F}_{\tau'} \setminus \mathcal{F}_{\tau}$,
- $\forall \bar{x}\ [\psi_R(\bar{x}) \leftrightarrow R(\bar{x})]$ for each $R \in \mathcal{R}_{\tau'} \setminus \mathcal{R}_{\tau}$.

Given $\tau \subseteq \tau'$, T and T' as above, then we have the restriction map

$$\mathrm{Res}_{\tau} : \mathrm{Str}(\tau') \longrightarrow \mathrm{Str}(\tau).$$

Now by assumption we can restrict the restriction map to $\mathrm{Mod}(T')$. The induced map on $\mathrm{Mod}(T')$ has range inside $\mathrm{Mod}(T)$ by assumption. With this setup we state the following proposition.

$$\begin{array}{ccc} \mathrm{Str}(\tau') & \xrightarrow{\mathrm{Res}_\tau} & \mathrm{Str}(\tau) \\ \big\uparrow \iota & & \big\uparrow \iota \\ \mathrm{Mod}(T') & \xrightarrow[\mathrm{Res}_\tau]{} & \mathrm{Mod}(T) \end{array}$$

Proposition. *Given* $\tau \subseteq \tau'$, T *and* T' *as above, the induced map* $\mathrm{Mod}(T') \longrightarrow \mathrm{Mod}(T)$ *is a bijection (of classes).*

Proof. Suppose $\mathfrak{A} \models T'$. We check that $\mathfrak{A}|_\tau \models T$. Let $\varphi \in T$. There are two cases.

- φ is $\exists^{=1} x\ \theta_c(x)$ for some constant c. For any $a \in \mathrm{dom}(\mathfrak{A})$ then $\mathfrak{A}_a \models \theta_c(a)$ if and only if $a = c^{\mathfrak{A}}$ thus $\mathfrak{A} \models \exists^{=1} x\ \theta_c(x)$ so $\mathfrak{A}|_\tau \models \varphi$ as well.
- φ is $\forall \bar{x} \exists^{=1} y\ \gamma_f(\bar{x}, y)$ for some function symbol f. Now $\mathfrak{A} \models \forall \bar{x} \exists^{=1} y\ f(\bar{x}) = y$ so

  $$\mathfrak{A} \models \forall \bar{x} \forall y\ [f(\bar{x}) = y \leftrightarrow \gamma_f(\bar{x}, y)]$$

  so $\mathfrak{A} \models \forall \bar{x} \exists^{=1} y\ \gamma_f(\bar{x}, y)$ which implies that $\mathfrak{A}|_\tau \models \varphi$.

Thus $\mathfrak{A}|_\tau \models T$.
Conversely, suppose $\mathfrak{A} \models T$. We want to expand $\mathfrak{A}$ to some τ'-structure $\mathfrak{A}'$ which is a model of T'.

- For $c \in \mathcal{C}'_\tau \setminus \mathcal{C}_\tau$ define $c^{\mathfrak{A}'}$ to be the unique $a \in \mathrm{dom}(\mathfrak{A})$ such that $\mathfrak{A}_a \models \theta_c(a)$.
- For $f \in \mathcal{F}'_\tau \setminus \mathcal{F}_\tau$ then we define $f^{\mathfrak{A}'}$ by

  $$f^{\mathfrak{A}'}(\bar{a}) = b \Leftrightarrow \mathfrak{A} \models \gamma_f(\bar{a}, b).$$

  Note that this actually defines a function because of what γ_f says.
- For $R \in \mathcal{R}'_\tau \setminus \mathcal{R}_\tau$ we let

  $$R^{\mathfrak{A}'} := \{\bar{a} \in \mathrm{dom}(\mathfrak{A})^{\mathrm{arity}(R)} : \mathfrak{A}_{\bar{a}} \models \psi_R(\bar{a})\}.$$

This makes $\mathfrak{A}'$ into an τ'-structure which is a model of T' and $\mathfrak{A}'|_\tau = \mathfrak{A}$. This completes the proof. $\square$

7.2 Atomisation/Morleyisation

The following construction is usually called *Morleyisation*. Hodges however, calls it *Atomisation*. He points out that Thoralf Skolem used this construction before Morley did. Since the term "Skolemisation" has a different meaning, Hodges decides that "atomisation" is both more correct and more descriptive.

Given a signature τ we build a new signature τ' (which will *not* be an expansion). Let $\mathcal{C}_{\tau'} = \mathcal{F}_{\tau'} = \varnothing$ and

$$\mathcal{R}_{\tau'} := \{R_{(\varphi,n)} : \varphi \in \mathcal{L}(\tau_{\{x_i:i<n\}})\}$$

with $\text{arity}(R_{(\varphi,n)}) = n$[1].

We now make a definitional expansion from a theory in τ to a theory in $\tau' \cup \tau$. Consider $T = \varnothing$ the empty theory in $\mathcal{L}(\tau)$, and T' a theory in $\mathcal{L}(\tau \cup \tau')$ given by

$$T' := \{\forall \bar{x}[R_{(\varphi,n)}(\bar{x}) \leftrightarrow \varphi(\bar{x})] : \varphi \in \mathcal{L}(\tau_{\{x_i:i<n\}})\}.$$

Then T and T' trivially satisfy the conditions for the definitional expansions as in the above section, since there are no new constant symbols and no new function symbols. By the proposition we proved for definitional expansions, each τ-structure $\mathfrak{A}$ (i.e. any model of $T = \varnothing$) admits a unique definitional expansion to a $\tau \cup \tau'$-structure $\mathfrak{A}'$ such that $\mathfrak{A}' \models T'$.

Definition. With the setup as describe above, the **atomisation** of $\mathfrak{A}$ is the reduct of $\mathfrak{A}'$ down to τ', i.e. $\mathfrak{A}^{\mathrm{Atom}} := \mathfrak{A}'|_{\tau'}$.

Proposition. *Let $\mathfrak{A}^{\mathrm{Atom}}$ be the atomisation of $\mathfrak{A}$. Then every definable set in $\mathfrak{A}^{\mathrm{Atom}}$ is defined by an atomic τ'-formula*

Proof. This is true by definition of definitional expansions.

Any subset $X \subseteq \text{dom}(\mathfrak{A}^{\mathrm{Atom}})^n = A^n$ is $\mathcal{L}(\tau')$-definable if and only if it is $\mathcal{L}(\tau)$-definable. So by construction X is definable if and only if it is definable by an atomic formula. $\square$

Remark. Depending on the definition one has of atomic formula we may need to assume that the definable sets in the propostion are defined in at least one variable. This is a necessary assumption if one does not count true ($\top$) and false ($\perp$) as atomic sentences.

Corollary. *With T' as above, if $\mathfrak{A}' \subseteq \mathfrak{B}'$ and $\mathfrak{A}', \mathfrak{B}' \models T'$ then $\mathfrak{A}' \preccurlyeq \mathfrak{B}'$ and $\mathfrak{A} \preccurlyeq \mathfrak{B}$ where $\mathfrak{A} := \mathfrak{A}'|_\tau$ and $\mathfrak{B} := \mathfrak{B}|_\tau$.*

Proof. We use the Tarski-Vaught Criterion, namely that $\mathfrak{A}' \preccurlyeq \mathfrak{B}'$ if and only if, for any formula $\theta(x) \in \mathcal{L}(\tau_{A'})$ we have $\mathfrak{A}' \models \exists x\theta(x)$ iff $\mathfrak{B}' \models \exists x\theta(x)$. The forward direction is immediate. The backwards direction we may prove by induction on the complexity of our formula. Suppose φ is the τ-formula. We must show that $\mathfrak{A}_{\bar{a}} \models \varphi(\bar{a})$ iff $\mathfrak{B}_{\bar{a}} \models \varphi(\bar{a})$.

[1]Many authors do not include the subscript n. They simply write R_φ without specifying the arity.

- For φ atomic it follows from $\mathfrak{A} \subseteq \mathfrak{B}$
- For $\varphi \equiv \neg\psi$ note that negation preserves the biimplication.
- The conjunction and disjunction are immidiate.
- For $\varphi \equiv \exists x\psi$ this is exactly our hypothesis.

$\square$

Remark. The above corollary only holds if we allow $\top$ and $\perp$ as atomic formulae. Otherwise we must assume that $\mathfrak{A} \equiv_{\mathcal{L}(\tau)} \mathfrak{B}$.

The atomisation process is useful for simplifying some arguments. Furthermore if one is only interested in the class of definable sets of a given structure then the atomisation is also useful since it has the same class, only this time each set is defined by atomic formulae. But if one actually wants to determine what these definable sets are, then the atomisation is completely useless.

We now give some examples and non-examples of elementary substructures.

Example. Let $\tau = \{<\}$. Let $\mathfrak{B} = (\mathbb{Q}, <)$ and $\mathfrak{A} = (\mathbb{Z}[\frac{1}{2}], <)$. Then $\mathfrak{A} \preccurlyeq \mathfrak{B}$. We will not prove this now.

Example. Let $\tau = \{<\}$. Let $2\mathbb{Z} = (2\mathbb{Z}, <)$ and $\mathbb{Z} = (\mathbb{Z}, <)$. Then $2\mathbb{Z} \subseteq \mathbb{Z}$ as τ-structures. Furthermore $2\mathbb{Z} \models \forall x \neg(x < 4 \wedge 2 < x)$ and $\mathbb{Z} \models \exists x(x < 4 \wedge 2 < x)$ so $2\mathbb{Z} \not\preccurlyeq \mathbb{Z}$. However, since $2\mathbb{Z}$ and $\mathbb{Z}$ are isomorphic as $\mathcal{L}(\tau)$-structures we do have that $2\mathbb{Z} \equiv_{\mathcal{L}(\tau)} \mathbb{Z}$.

Question. Does there exist $\mathfrak{B}$ a τ-structure and $\mathfrak{A} \subseteq \mathfrak{B}$ and $f : \mathfrak{B} \longrightarrow \mathfrak{A}$ a definable isomorphism (i.e. the graph of f is a definable set) such that $\mathfrak{A} \not\preccurlyeq \mathfrak{B}$? [Hint: $S : \omega \longrightarrow \mathbb{Z}_+$ the succesor map.]

Question. Is it true that given $\mathfrak{B}$ and $\mathfrak{A} \subseteq \mathfrak{B}$ and a definable isomorphism $f : \mathfrak{B} \longrightarrow \mathfrak{A}$ such that $\mathfrak{A} \preccurlyeq \mathfrak{B}$ then we must have $\mathfrak{A} = \mathfrak{B}$?

In some sense the notion of an extension of a structure is the right notion from the level of the atomic formulae. But if one is interested in the definable sets and how to interpret the formulae from one structure to another, then *elementary* extension is the right notion.

Lecture 8

8.1 Quantifier Elimination

In practice one does not use the atomisation procedure to gain information about the definable sets. To actually gain information one can hope to find a reasonable class of formulae from which every definable set can be defined.

Definition. Given a signature τ and a class $\mathcal{K}$ of τ-structures. A set $\Phi \subseteq \mathcal{L}(\tau)$ of τ-formulae is an **elimination set for** $\mathcal{K}$ if

- for every formula $\psi(\bar{x}) \in \mathcal{L}(\tau)$ (with at least one free variable) there exists a boolean combination $\varphi(\bar{x})$ of formulae from Φ such that for all $\mathfrak{A} \in \mathcal{K}$ we have $\mathfrak{A} \models \forall \bar{x}(\varphi \leftrightarrow \psi)$.

We say that $\mathcal{K}$ has **quantifier elimination** if we can take the elimination set Φ to be the collection of atomic formulae.

Remark. For any τ and any $\mathcal{K}$ there always exists an elimination set, namely $\mathcal{L}(\tau)$ itself.

Remark. Relativising the atomisation construction to $\Phi \subseteq \mathcal{L}(\tau)$, then Φ is an elimination set for $\mathcal{K}$ if and only if each definable expansion of structures of $\mathcal{K}$ to the relative atomisation has quantifier elimination.

Definition. For a structure $\mathfrak{A}$, we say that $\mathfrak{A}$ **eliminates quantifiers** if $\{\mathfrak{A}\}$ does so. Similarly if $\mathcal{K} = \mathrm{Mod}(T)$ then we say T **eliminates quantifiers**.

8.2 Quantifier Elimination for $(\mathbb{N}, <)$

We will prove quantifier elimination for $\mathrm{Th}(\mathbb{N}, <)$. The proof will be very syntactic and quite heavy-handed. This should in part serve as motivation for the more structural methods of quantifier elimination that will be developed later in the course.

Theorem 8.1. *In* $\mathcal{L}(<)$ *the set* Φ *consisting of the atomic formulae together with*

$$\text{for every } n \in \mathbb{Z}_+ \qquad B_n(x, y) := \exists \bar{z}\ x < z_1 < \cdots < z_n < y$$

and

$$\text{for every } n \in \mathbb{Z}_+ \qquad L_n(x) := \exists \bar{z}\ x > z_1 > \cdots > z_n,$$

is an elimination set for $\mathrm{Th}(\mathbb{N}, <)$. *I.e. in the signature* τ' *where* $\mathcal{C}_{\tau'} = \mathcal{F}_{\tau'} = \varnothing$ *and*

$$\mathcal{R}_{\tau'} = \{<\} \cup \{B_n : n \in \mathbb{Z}_+\} \cup \{L_n : n \in \mathbb{Z}_+\}.$$

then the expansion $\mathrm{Th}_{\mathcal{L}(\tau')}(\mathbb{N}, <)$ *has quantifier elimination.*

Remark. $B_n(x, y)$ says 'there are n elements between x and y'. $L_n(x)$ says 'there are n elements less than x'.

Proof. By the lemma proved last time it suffices to show that if $\varphi(x, y_0, \ldots, y_{n-1}) \in \mathcal{L}(\tau')$ is quantifier-free with free variables amongst $x, \bar{y}$, then $\exists x \varphi(x, \bar{y})$ is equivalent to a quantifier-free formula in $\mathcal{L}(\tau')$.

Write φ in (the equivalent) disjunctive normal form as $\bigvee_i \bigwedge_j \theta_{i,j}$ where each $\theta_{i,j}$ is a literal. Now since the operator $\exists$ distributes over $\bigvee$ it suffices to eliminate quantifiers from the conjuncts. I.e. it suffices to show that $\exists x \bigwedge \theta_{i,j}$ is equivalent to some quantifier-free formula.

So from now on φ will be renamed to $\bigwedge \theta_{i,j}$.

Now we must figure out what the literals $\theta_{i,j}$ can possibly be. For example they can be of the following form, $L_n(y_i), L_n(x), \neg L_n(y_j), \neg L_n(x), B_n(x, y_i), \neg B_n(x, y_i)$, $x < y_i$, $\neg(x < y_i), \ldots$.

We make another simplifying observation: φ is equivalent to a big disjunction over all possible "order relations" between the elements $y_0, \ldots, y_{n-1}, x$ of the formula describing this given "order relation" conjoined with $\varphi(x, \bar{y})$. Here "order relation" means a possible way that the variables $y_0, \ldots, y_{n-1}$ and x can be related via the $<$ relation symbol. For example one such order relation $\psi(x, \bar{y})$ could be

$$y_0 = y_1 = y_2 < y_3 < \cdots < y_i = \cdots = y_{n-1} < x.$$

There are only finitely many different such order relations. Now the observation is that $\varphi(x, \bar{y})$ is equivalent to

$$\bigvee \{\psi(x, \bar{y}) \wedge \varphi(x, \bar{y}) : \psi(x, \bar{y}) \text{ is an order relation}\}.$$

So we may assume that we have already pinned down completely the order relation of $y_0, \ldots, y_{n-1}, x$. This means that any other order condition contained inside φ (for example one of the $\theta_{i,j}$ could say $x < y_i$ or $x = y_j$) will now either be redundant or explicitly contradictory with the order relation. If one of the literals $\theta_{i,j}$ is an explicit contradiction to the order relation then it is easy to find an equivalent quantifier-free formula, namely any false sentence. If one of the literals $\theta_{i,j}$ is redundant then we need not worry about it.

Now we complete the proof by considering the remaining cases.

- If $\varphi \longrightarrow x{=}y_j$ for some $j \leqslant n-1$ then $\exists x\ \varphi(x, \bar{y})$ is equivalent to $\varphi(y_j, \bar{y})$, which is quantifier-free.
- So we may assume $\varphi \longrightarrow \bigwedge_i x \neq y_i$. Then there will be a single smallest interval where x is. I.e. $\varphi \longrightarrow y_i < x < y_j$ for some unique i, j such that $\varphi \longrightarrow \bigwedge_k \neg(y_i < y_k < y_j)$. Of course x could also be smaller than or greater than all the y_i's. This gives two more cases, but for convenience we shall allow "$y_i = \pm\infty$" and "$y_j = \pm\infty$". Pictorially we now fix ourselves in the following generic situation:

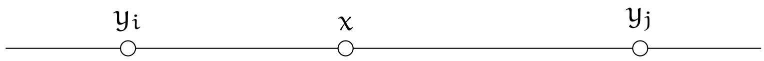

  Now we must consider the other $\theta_{i,j}$'s in φ. I.e. the conditions which make use of the symbols L_n and B_n (with possible negations). Since we have fixed the order relation for φ we know that $\exists x\varphi$ holds if and only if there exists an x which satisfies each of the extra conditions individually. Thus it suffices to eliminate quantifiers from simple formulae of the form $\exists x\theta_{i,j}(x, \bar{y})$ for the cases where $\theta_{i,j}$ is one of the L_n or B_n (or negations thereof). Note that we need only concern ourselves with the $\theta_{i,j}$'s that have instances of x in them. There are four cases:

  - $\exists x L_r(x)$ is equivalent to $L_{r+1}(y_i)$.
  - $\exists x \neg L_r(x)$ is equivalent to $\neg L_{r-1}(y_j)$
  - If $x < y_k$ then $\exists x B_t(x, y_k)$ is equivalent to $B_{t+1}(y_i, y_k)$. If $y_k < x$ then $B_t(x, y_k)$ is equivalent to $B_{t+1}(y_j, y_k)$.
  - For $x < y_k$ then $\exists x \neg B_t(x, y_k)$ is equivalent to $\neg B_{t-1}(y_j, y_k)$. If $y_k < x$ then $\exists x \neg B_t(x, y_k)$ is equivalent to $\neg B_{t+1}(y_i, y_k)$.

This completes the proof. $\square$

Once we have the compactness theorem one can use back-and-forth type arguments to greatly reduce the trouble with proving quantifier elimination results. The above proof demonstrates somewhat the idea behind proofs to come; we tried to "complete" the formula $\exists x\varphi(x, \bar{y})$ as much as possible so that there is only one formula to think about.

8.3 Skolem's theorem

We now move to more structural ideas. First we prove Skolem's theorem, showing that given any model we can find an elementary submodel of size $\leqslant$ the cardinality of the language.

Theorem 8.2. *For any τ-structure $\mathfrak{A}$ there exists $\mathfrak{B} \preccurlyeq \mathfrak{A}$ with $|\mathfrak{B}| \leqslant |\mathcal{L}(\tau)|$.*

Proof. Let $A = \mathrm{dom}(\mathfrak{A})$. We will build an increasing sequence $B_0 \subseteq B_1 \subseteq \cdots$ of subsets of A such that the union $B := \bigcup_i B_i$ will give us the domain of an elementary substructure of $\mathfrak{A}$.

At stage 0 set $B_0 := \varnothing$. At stage $n+1$ list all $\varphi(x)$ in $\mathcal{L}(\tau_{B_n,x})$ (i.e. φ has one free variable x and parameters from B_n). For each φ, if $\mathfrak{A}_{B_n} \models \exists x \varphi(x)$ then let a_φ be a witness. Now set

$$B_{n+1} := B_n \cup \{a_\varphi : \varphi \in \mathcal{L}(\tau_{B_n,x}) \text{ and } \mathfrak{A}_{B_n} \models \exists x \varphi(x)\}$$

First let us check that $|B_n| \leqslant |\mathcal{L}(\tau)|$. For $n = 0$ this is clear. For $n+1$ by have by induction hypothesis that $|B_n| \leqslant |\mathcal{L}(\tau)|$. Then $|\mathcal{L}(\tau_{B_n})| = |\mathcal{L}(\tau)|$, which implies that

$$|B_{n+1}| \leqslant |B_n| + |\mathcal{L}(\tau)| \leqslant |\mathcal{L}(\tau)|.$$

So

$$|B| = |\bigcup_n B_n| \leqslant |\mathcal{L}(\tau)|.^{1}$$

Finally we claim that B is the domain of an elementary substructure of $\mathfrak{A}$. It is the domain of a substructure because, by construction, it contains witnesses to statements of the form $\exists x\ x{=}c$ for each constant symbol c, and $\exists x f(b_1, \ldots, b_n){=}x$ for each function symbol f and $b_1, \ldots b_n \in B$. Finally it is an elementary substructure by the Tarski-Vaught Test. $\square$

As a simply corollary we have.

Corollary. *If $\mathfrak{A}$ is any τ-structure and λ a cardinal such that $|\mathcal{L}(\tau)| \leqslant \lambda \leqslant |\mathfrak{A}|$. Then there exists $\mathfrak{B} \preccurlyeq \mathfrak{A}$ such that $|\mathfrak{B}| = \lambda$.*

Proof. Let $Z \subseteq A$ be a subset with $|Z| = \lambda$. Consider the expansion τ_Z. Then by Skolem's theorem we find $\mathfrak{B}_Z \preccurlyeq \mathfrak{A}_Z$ with $|B_Z| \leqslant |\mathcal{L}(\tau_Z)|$. But we also have $|Z| \leqslant |B_Z|$ and so $B_Z = |Z| = \lambda$. Now let $\mathfrak{B} := \mathfrak{B}_Z|_\tau$ to get $\mathfrak{B} \preccurlyeq \mathfrak{A}$. $\square$

[1] Here we are using the theorem from set theory that a union of an ω-chain of elements having cardinality $\leqslant \lambda$ (for $\lambda \geqslant \omega$) has cardinality $\leqslant \lambda$.

8.4 Skolem functions

Skolem himself proved his theorem by using what we shall call Skolemisation. The process will be used again and again.

Definition. A theory T in a signature τ has **Skolem functions** if for each formula $\varphi(\bar{x}, y) \in \mathcal{L}(\tau)$ there is a (not necessarily unique) function symbol f_φ with $\operatorname{arity}(f_\varphi) = \operatorname{length}(\bar{x})$ such that T contains the formula

$$\forall \bar{x} \, [\exists y \;\varphi(\bar{x}, y) \longleftrightarrow \varphi(\bar{x}, f_\varphi(\bar{x})] \, .$$

So the function f_φ finds witnesses (depending on $\bar{x}$) whenever $\exists y \, \varphi(\bar{x}, y)$ is true.

Remark. Some writers replace function symbols by terms, so that "Skolem functions" are actually terms of the signature. In this way one can better handle the case where $\operatorname{length}(\bar{x}) = 0$. We would need a "0-ary function symbol" which our definition does not allow. In our definition we can simply add dummy variables so that $\operatorname{length}(\bar{x}) > 0$.

Remark. Some people say that T *has built in Skolem functions* if it has a definitional expansion with Skolem functions. Our notion of a theory with Skolem functions is more restrictive.

Example. If $\tau = \varnothing$ then the "theory of equality" is a theory *without* Skolem functions.

Example. The theory $\operatorname{Th}(\mathbb{N}, +, \cdot, 0, 1, <)$ has built in Skolem functions (in the sense of the above remark) but does *not* have Skolem functions in our sense.

Next time we shall show how to add Skolem functions to our theories, a process called *Skolemisation*. A Corollary to this will be another proof of the Löwenheim-Skolem theorem.

Lecture 9

9.1 Skolemisation

Theorem 9.1 (Skolemisation Theorem)**.** *For any signature τ there exists a signature τ' ($= \tau^{\mathrm{skolem}}$) and a τ'-theory T' ($= T^{\mathrm{skolem}}$) such that*

- T' *has skolem functions*
- *Every τ-structure extends to a τ'-structure which models T'. I.e. the restriction map* $\mathrm{Mod}(T') \longrightarrow \mathrm{Str}(\tau)$ *is surjective.*

Proof. The basic idea is just to put the required Skolem functions into the signature. Of course just doing this for τ (say by extending τ to τ^+) doesn't work since there will be new formulae in the language $\mathcal{L}(\tau^+)$ which lack Skolem functions. To remedy this we construct a chain and take a union.

We construct an increasing sequence of signatures $\tau_0 \subseteq \tau_1 \subseteq \cdots$ and theories $T_0 \subseteq T_1 \subseteq \cdots$ (where T_n is an $\mathcal{L}(\tau_n)$-theory). Then let

$$\tau' := \bigcup_n \tau_n \qquad \text{and} \qquad T' := \bigcup_n T_n.$$

The construction is as follows.

- Let $\tau_0 := \tau$ and $T_0 := \varnothing$.
- At stage n define $\mathcal{C}_{\tau_{n+1}} := \mathcal{C}_{\tau_n}$, $\mathcal{R}_{\tau_{n+1}} := \mathcal{R}_{\tau_n}$, and

  $$\mathcal{F}_{\tau_{n+1}} := \mathcal{F}_{\tau_n} \cup \{f_{(\varphi,m)} : \varphi \in \mathcal{L}(\tau_n), \text{ with free variables amongst } x_0, \ldots, x_{m-1}\}.$$

  The theory T_{n+1} will be T_n together with

  $$\{\forall x_0, \ldots, x_{m-1}[\varphi(x_0, \ldots, x_{m-1}, f_{(\varphi,m)}(\bar{x})) \longleftrightarrow \exists y\ \varphi(\bar{x}, y)] \\ : \varphi \in \mathcal{L}(\tau_n) \text{ free variables in } x_0, \ldots, x_{m-1}\}$$

Now we claim that T' has Skolem functions. Indeed if $\varphi(\bar{x}, y) \in \mathcal{L}(\tau')$ with free variables amongst $\bar{x}, y$ then $\varphi(\bar{x}, y) \in \mathcal{L}(\tau_n)$ for some n. Now we constructed T_{n+1} to say that

$$\forall \bar{x}(\varphi(\bar{x}, f_{(\varphi,m)}(\bar{x})) \longleftrightarrow \exists y\ \varphi(\bar{x}, y)).$$

So T' does have a skolem function for φ.

Now we show the second claim, namely that the restriction map $\mathrm{Mod}(T') \longrightarrow \mathrm{Str}(\tau)$ is surjective. We show that if $\mathfrak{A}_n \in \mathrm{Mod}(T_n)$ then there exists $\mathfrak{A}_{n+1} \in \mathrm{Mod}(T_{n+1})$ such that $\mathfrak{A}_n = \mathfrak{A}_{n+1}|_{\tau_n}$; To find $\mathfrak{A}_{n+1}$ we basically need to show how to interpret the Skolem functions that entered at stage n.

Let $f_{\varphi,m} \in \mathcal{F}_{\tau_{n+1}}$ be a new function symbol in τ_{n+1}. To interpret $f_{\varphi,m}$ we shall need the *axiom of choice*:

For all sets X *there exists a map* $f : X \setminus \{\varnothing\} \to \bigcup X$ *such that* $\forall y \in X\ f(y) \in y$.

Let $X = \{\varphi(\bar{a}, \mathfrak{A}) : \bar{a} \in A^m\}$ and let g be a choice function for X as afforded by the axiom of choice. Then define

$$f_{\varphi,m}^{\mathfrak{A}_{n+1}}(\bar{a}) = \begin{cases} g(\varphi(\bar{a}, \mathfrak{A})) & \text{if } \varphi(\bar{a}, \mathfrak{A}) \neq \varnothing \\ a_0 & \text{otherwise} \end{cases}$$

Note that the second clause only happens when $\varphi(\bar{a}, \mathfrak{A})$ is empty, but then whatever $f_{\varphi,m}$ does to $\bar{a}$ doesn't matter. This interpretation makes $\mathfrak{A}_{n+1}$ into a τ_{n+1}-structure which models T_{n+1} and restricts back to $\mathfrak{A}_n$. □

For a model $\mathfrak{B}$ of a theory T with Skolem function the notion of substructure and elementary substructure coincide! This follows as a Corollary to the following proposition.

Proposition. *If* T *is a* τ*-theory with Skolem functions then for every formula* $\theta(\bar{x})$ *with* at least one free variable, *there is a quantifier-free formula* $\varphi(\bar{x})$ *such that* $T \models \forall \bar{x}(\theta(\bar{x}) \leftrightarrow \varphi(\bar{x}))$.

Proof. We work by induction on the complexity of θ. The atomic case is immediate. Boolean combinations are also immediate. For the case $\theta(\bar{x})$ is $\exists y\ \psi(\bar{x}, y)$ then we have

$$T \models \forall \bar{x}(\theta(\bar{x}) \leftrightarrow \psi(\bar{x}, f_\psi))$$

since T has Skolem functions. By induction we may find an equivalent formula for ψ. □

Corollary. *If* T *has Skolem functions and* $\mathfrak{B} \models T$, *then if* $\mathfrak{A} \subseteq \mathfrak{B}$ *then* $\mathfrak{A} \preccurlyeq \mathfrak{B}$ *provided* $\mathfrak{A} \neq \varnothing$.

Proof. Let $\mathfrak{A} \subseteq \mathfrak{B}$ be nonempty. We show $\mathfrak{A} \preccurlyeq \mathfrak{B}$. By the Tarski-Vaught criterion it suffices to check that for any formula $\varphi(\bar{x}, y)$ and $\bar{a}$ a tuple of elements from $\mathfrak{A}$ (which exists since $\mathfrak{A} \neq \varnothing$) then $\mathfrak{A} \models \exists y\ \varphi(\bar{a}, y))$ iff $\mathfrak{B} \models \exists y\ \varphi(\bar{a}, y)$. The forward direction is immediate. Suppose $\mathfrak{B} \models \exists y\ \varphi(\bar{a}, y)$ then since T has Skolem functions, $\mathfrak{B} \models \varphi(\bar{a}, f_\varphi(\bar{a}))$ and since $\varphi(\bar{a}, f_\varphi(\bar{a}))$ only involves parameters from $\mathfrak{A}$ we have $\mathfrak{A} \models \varphi(\bar{a}, f_\varphi(\bar{a}))$ which implies that $\mathfrak{A} \models \exists y\ \varphi(\bar{a}, y)$. □

So skolemisation gives a way of building elementarily equivalent substructures. As a corollary we get the full Downward Löwenheim-Skolem theorem.

Theorem 9.2 (Downward Löwenheim-Skolem theorem). *Let $\mathcal{L}(\tau)$ be a first-order language, $\mathfrak{A}$ a τ-structure, X a set of elements of $A = \text{dom}(\mathfrak{A})$ and λ a cardinal such that $|\mathcal{L}(\tau)| + |X| \leqslant \lambda \leqslant |\mathfrak{A}|$. Then $\mathfrak{A}$ has an elementary substructure $\mathfrak{B}$ of cardinality λ with $X \subseteq \text{dom}(\mathfrak{B})$.*

Proof. We skolemise the empty τ-theory $T = \varnothing$ to get a τ^{skolem}-theory T^{skolem} and an extension of $\mathfrak{A}$ to a model $\mathfrak{A}^{\text{skolem}}$ of T^{skolem}. Let Y be a subset of A with $|Y| = \lambda$ and $X \subseteq Y$. Then Let $\mathfrak{B}'$ be the substructure generated by Y. Finally take the reduct $\mathfrak{B}$ of $\mathfrak{B}'$ to τ. Now $|\mathfrak{B}| \leqslant |Y| + |\mathcal{L}(\tau^{\text{skolem}})| = \lambda + |\mathcal{L}(\tau)| = \lambda = |Y| \leqslant |\mathfrak{B}|$. By the above corollary $\mathfrak{B}' \preccurlyeq \mathfrak{A}^{\text{skolem}}$ hence $\mathfrak{B} \preccurlyeq \mathfrak{A}$. $\square$

9.2 Back-and-Forth and Games

We will now discuss games for testing equivalence of structures. There are many different forms of games. Different forms of games will bring different notions of equivalence which correspond to different logics on the structures.

As a prelude we prove a famous theorem due to Cantor.

Theorem 9.3 (Cantor's Back-and-Forth Theorem). *If $(A, \leqslant)$ and $(B, \leqslant)$ are nonempty countable dense linear orders without endpoints then they are isomorphic.*

Notation. The abbreviation "DLO" is commonly used for the theory of dense linear orders without endpoints.

Proof. Let $A = \{a_n : n \in \omega\}$ and $B = \{b_n : n \in \omega\}$ be some enumerations of A and B respectively. We shall construct an increasing sequence $\{f_n : n \in \omega\}$ of partial isomorphisms (i.e. f_n is an isomorphisms between its domain and codomain thought of as substructures) such that

1. $f_n \subseteq f_{n+1}$,
2. $\text{dom}(f_n) \supseteq \{a_0, \ldots, a_{n-1}\}$ and $\text{range}(f_n) \supseteq \{b_0, \ldots, b_{n-1}\}$,
3. and f_n is finite for all n.

We do this as follows. We let $f_0 := \varnothing$. At stage $n+1$ we want to extend f_n to ensure that $a_n \in A$ is in the domain and that b_n is in the range of f_{n+1}. There are four cases.

- If $a_n \in \text{dom}(f_n)$ then $f_{n+\frac{1}{2}} = f_n$

- If $\forall a \in \mathrm{dom}(f_n)\ a_n < a$ then since $B \models \mathrm{DLO}$ there exists $b' \in B$ such that $\forall b \in \mathrm{range}(f_n)\ b' < b$ (here we are using that $\mathrm{range}(f_n)$ is finite). Then set $f_{n+\frac{1}{2}}(a_n) = b'$.
- If $\forall a \in \mathrm{dom}(f_n)\ a < a_n$ then since $B \models \mathrm{DLO}$ there exists $b' \in B$ such that $\forall b \in \mathrm{range}(f_n)\ b < b'$ (again since $\mathrm{range}(f_n)$ is finite). Set $f_{n+\frac{1}{2}}(a_n) = b'$.
- If there is $a, b \in \mathrm{dom}(f_n)$ with $a < b$ and $(a, b) \cap \mathrm{dom}(f_n) = \emptyset$ and $a < a_n < b$ then $\forall c \in \mathrm{range}(f_n)$ we have $\neg(f_n(a) < c < f_n(b))$ since f_n is an isomorphism. Now since B is dense there is some d such that $f_n(a) < d < f_n(b)$. Pick such a d and define $f_{n+\frac{1}{2}}(a_n) = d$.

This tells us how to map a_n forward. Now dual arguments show how to extend f_n^{-1} to $f_{n+\frac{1}{2}}^{-1}$ so that $f_{n+\frac{1}{2}}^{-1}$ is defined on b_n. Putting both directions together we get the maps f_{n+1} and f_{n+1}^{-1}.
The sequence $(f_n)_{n\in\omega}$ clearly satisfies the requirements 1. 2. and 3.. Now letting

$$f = \bigcup_n f_n$$

we get that f is an isomorphism between $A = \mathrm{dom}(f)$ and $B = \mathrm{range}(f)$. $\square$

Remark. An alternative formulation of the theorem is that DLO is an $\aleph_0$-categorical theory.

9.3 The Ehrenfeucht-Fraïssé game

The proof of Cantor's theorem is an example of the *back and forth method*. We can formalize this argument in terms of a game called the *Ehrenfeucht-Fraïssé game of length* ω.

There are two players; $\forall$ (Abelard) and $\exists$ (Heloise). Let γ be an ordinal. The **Ehrenfeucht-Fraïssé game** of length γ between τ-structures $\mathfrak{A}$ and $\mathfrak{B}$, denoted $\mathrm{EF}_\gamma(\mathfrak{A}, \mathfrak{B})$, has γ moves. At move α, $\forall$ picks an element from either $\mathfrak{A}$ or $\mathfrak{B}$. $\exists$ responds with an element from the other model. A play of the $\mathrm{EF}_\gamma(\mathfrak{A}, \mathfrak{B})$-game is a γ-tuple $(a_\alpha, b_\alpha)_{\alpha<\gamma}$ where $a_\alpha \in A$ and $b_\alpha \in B$.
Player $\exists$ wins the play $(a_\alpha, b_\alpha)_{\alpha<\gamma}$ if the map $a_\alpha \longmapsto b_\alpha$ has the property that for every for atomic formula $\varphi(\bar{x})$ we have

$$\mathfrak{A} \models \varphi(\bar{a}) \qquad \Longleftrightarrow \qquad \mathfrak{B} \models \varphi(\bar{b})$$

where $\bar{a}$ is a tuple from $(a_\alpha)_{\alpha<\gamma}$ and $\bar{b}$ is the image of $\bar{a}$ under the map $a_\alpha \longmapsto b_\alpha$.
A **winning strategy** for $\mathrm{EF}_\gamma(\mathfrak{A}, \mathfrak{B})$ is a function from the set of partial plays to plays up through stage α together with $\forall$'s play are stage α, which returns a play for $\exists$, such that if $\exists$ follows this function then she always wins.

Definition. We say that $\mathfrak{A}$ is γ**-equivalent** to $\mathfrak{B}$, written $\mathfrak{A} \sim_\gamma \mathfrak{B}$, if $\exists$ has a winning strategy for $\mathrm{EF}_\gamma(\mathfrak{A}, \mathfrak{B})$.

Remark. Note that if $\mathfrak{A} \cong \mathfrak{B}$ then $\mathfrak{A} \sim_\gamma \mathfrak{B}$ for any ordinal γ. The winning strategy is given simply by the isomorphism.

Remark. $\mathfrak{A} \sim_0 \mathfrak{B}$ if and only if for all atomic *sentences* φ we have $\mathfrak{A} \models \varphi$ iff $\mathfrak{B} \models \varphi$.

Example. $(\mathbb{Q}, <) \sim_\omega (\mathbb{R}, <)$. This follows from the Back and Forth method demonstrated in the proof of Cantors theorem.

Example. $(\mathbb{Q}, <) \not\sim_{\omega+1} (\mathbb{R}, <)$. To see this we must show that $\forall$ can force a win in the $\mathrm{EF}_{\omega+1}((\mathbb{Q}, <), (\mathbb{R}, <))$-game. To do this $\forall$ may start with an enumeration $(q_n)_{n<\omega}$ of $\mathbb{Q}$. At each stage $n < \omega + 1$ in the game, $\forall$ picks an element of $\mathbb{R}$ corresponding to the rational number q_n sitting inside of $\mathbb{R}$. Then $\exists$ must always pick elements from $\mathbb{Q}$. Now at the ω'th play $\forall$ picks some irrational element of $\mathbb{R}$. $\exists$ must now pick one of its previous choices from $\mathbb{Q}$ and looses the game since the resulting function will not be an isomorphism.

Lecture 10

10.1 The Unnested Ehrenfeucht-Fraïssé Game

We have introduced the Ehrenfeucht-Fraïssé game $\mathrm{EF}_\omega(\mathfrak{A}, \mathfrak{B})$ as a way of testing for similarities between the structures $\mathfrak{A}$ and $\mathfrak{B}$. In particular we saw that if $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic then $\exists$ has a winning strategy (i.e. then $\mathfrak{A} \sim_\omega \mathfrak{B}$). The natural question is then: how similar are $\mathfrak{A}$ and $\mathfrak{B}$ if we know that $\exists$ has a winning strategy for $\mathrm{EF}_\omega(\mathfrak{A}, \mathfrak{B})$? The answer is somewhere between elementary equivalence and isomorphism.

Notation. Recall that $\mathfrak{A} \equiv_{\infty\omega} \mathfrak{B}$ means that $\mathfrak{A}$ and $\mathfrak{B}$ agree on all sentence of the infinitary language $\mathcal{L}_{\infty\omega}(\tau)$. In particular $\mathfrak{A} \equiv_{\infty\omega} \mathfrak{B}$ implies $\mathfrak{A} \equiv \mathfrak{B}$.

Theorem 10.1. *If τ is a countable signature, then $\mathfrak{A} \sim_\omega \mathfrak{B}$ if and only if $\mathfrak{A} \equiv_{\infty\omega} \mathfrak{B}$.*

Proof.
"$\Longrightarrow$" We will show by induction on the complexity of an $\mathcal{L}_{\infty\omega}(\tau)$-sentence Φ, for *any* signature τ, that $\mathfrak{A} \sim_\omega \mathfrak{B}$ implies $\mathfrak{A} \models \Phi \Leftrightarrow \mathfrak{B} \models \Phi$.

- If Φ is atomic then since $\mathfrak{A} \sim_\omega \mathfrak{B}$ implies $\mathfrak{A} \sim_0 \mathfrak{B}$ which implies that $\mathfrak{A} \models \Phi \Leftrightarrow \mathfrak{B} \models \Phi$.
- If Φ is $\neg\Psi$ then $\mathfrak{A} \models \Phi$ iff $\mathfrak{A} \not\models \Psi$ iff (by Induction Hypothesis) $\mathfrak{B} \not\models \Psi$ iff $\mathfrak{B} \models \Phi$.
- If Φ is $\bigvee \Xi$ then $\forall \xi \in \Xi$ we have $\mathfrak{A} \models \xi \Leftrightarrow \mathfrak{B} \models \xi$ which happens iff $\mathfrak{B} \models \Phi$. Likewise if Φ is $\bigwedge \Xi$.
- Suppose Φ is $\exists x \Psi(x)$ and that $\mathfrak{A} \models \Phi$. Then there exists $a \in A$ such that $\mathfrak{A}_a \models \Psi(a)$. By hypothesis $\exists$ has a winning strategy for $\mathrm{EF}_\omega(\mathfrak{A}, \mathfrak{B})$. Treating a as the 0th move of $\forall$ let $b \in B$ be the element that $\exists$ picks by way of her winning strategy. Then (a, b) is a winning position for $\exists$ in the game $\mathrm{EF}_\omega(\mathfrak{A}, \mathfrak{B})$. This is equivalent to $\exists$ having a winning strategy for the game $\mathrm{EF}_\omega(\mathfrak{A}_a, \mathfrak{B}_b)$ i.e. $\mathfrak{A}_a \sim_\omega \mathfrak{B}_b$. So by the induction hypothesis $\mathfrak{B}_b \models \Psi(b)$ i.e. $\mathfrak{B} \models \exists x \Psi(x)$. Thus $\mathfrak{B} \models \Phi$. Reversing the roles of $\mathfrak{A}$ and $\mathfrak{B}$ we see that $\mathfrak{B} \models \Phi$ also implies that $\mathfrak{A} \models \Phi$.

"$\Longleftarrow$": Conversely suppose $\mathfrak{A} \equiv_{\infty\omega} \mathfrak{B}$. We claim that $\mathfrak{A} \sim_\omega \mathfrak{B}$. If $\forall$ plays $a \in A$ then $\exists$ will respond $b \in \mathfrak{B}$ such that $\mathrm{tp}(a) = \mathrm{tp}(b)$[1] in $\mathcal{L}_{\infty\omega}(\tau)$. More precisely let[2]

$$\Phi := \{\varphi(x) \in \mathcal{L}_{\infty\omega}(\tau) : \mathfrak{A} \models \varphi(a/x) \text{ with (number of symbols in } \varphi) \leqslant 2^{|\mathfrak{A}|+|\mathfrak{B}|+\aleph_0}\}.$$

Then $\mathfrak{A} \models \bigwedge \Phi(a)$ i.e. $\mathfrak{A} \models \exists x \bigwedge \Phi(x)$. By assumption $\mathfrak{B} \models \exists x \bigwedge \Phi(x)$. Let $b \in B$ be a witness. Then $\exists$ responds to $\forall$ by choosing the element b. Then $\mathfrak{A}_a \equiv_{\infty\omega} \mathfrak{B}_b$. Continuing in this way we get a ω-sequence which is a win for $\exists$. Following this procedure is thus a winning strategy for $\exists$, so $\mathfrak{A} \sim_\omega \mathfrak{B}$. $\square$

So $\mathrm{EF}_\omega(\mathfrak{A}, \mathfrak{B})$ does characterize an equivalence between $\mathfrak{A}$ and $\mathfrak{B}$ but in the very strong infinitary logic of $\mathcal{L}_{\infty\omega}$.

We will now slightly modify the game with the aim of getting a new game that exactly characterizes (for finite signatures) elementary equivalence, i.e. equivalence in first-order logic.

What follows relies heavily on the notion of unnested formulae. For convenience we repeat the definition.

Definition. An **unnested atomic formula** is one of the form

- $x = c$, for $c \in \mathcal{C}_\tau$ and x a variable.
- $F\bar{x} = y$ where $F \in \mathcal{F}_\tau$ and $\bar{x}, y$ are variables.
- $R\bar{x}$, where $R \in \mathcal{R}_\tau$ and $\bar{x}$ are variables.
- $x = y$, where x and y are variables.

An **unnested formula** is built from the unnested atomic formulae by the usual rules.

Definition. The **unnested Ehrenfeucht-Fraïssé game** $\mathrm{EF}_\alpha[\mathfrak{A}, \mathfrak{B}]$ (note the square brackets) is the game where at stage $\beta < \alpha$, $\forall$ chooses en element from either $\mathfrak{A}$ or $\mathfrak{B}$ (i.e. $a_\beta \in A$ or $b_\beta \in B$) and $\exists$ responds with an element from the other structure. $\exists$ wins if for every *unnested* atomic formula $\varphi(x_{\beta|_{\beta<\alpha}})$

$$\mathfrak{A} \models \varphi(\bar{a}) \qquad \Longleftrightarrow \qquad \mathfrak{B} \models \varphi(\bar{b}).$$

We write $\mathfrak{A} \approx_\alpha \mathfrak{B}$ when $\exists$ has a winning strategy in $\mathrm{EF}_\alpha[\mathfrak{A}, \mathfrak{B}]$.

Remark. Most often the ordinal α in the above definition will either be finite or will be ω.

[1] $\mathrm{tp}(a)$ is the *type* of x it is the set of all $\mathcal{L}_{\infty\omega}(\tau)$ sentences in one variable which are true of a.

[2] the reason for the somewhat odd bound on the number of symbols in the definition of Φ is to ensure that Φ is actually a set.

Remark. If $\mathfrak{A} \sim_\alpha \mathfrak{B}$ then $\mathfrak{A} \approx_\alpha \mathfrak{B}$. This is clear since a play in $\mathrm{EF}_\alpha[\mathfrak{A}, \mathfrak{B}]$ is in fact a play of the old game $\mathrm{EF}_\alpha(\mathfrak{A}, \mathfrak{B})$. So a winning strategy in $\mathrm{EF}_\alpha(\mathfrak{A}, \mathfrak{B})$ is also a winning strategy in $\mathrm{EF}_\alpha[\mathfrak{A}, \mathfrak{B}]$.

The converse is not true. For example take $\tau = \{0, 1\}$ where 0 and 1 are constant symbols. Let $\mathfrak{A}$ be a τ-structure where A has one element and where $0^{\mathfrak{A}} = 1^{\mathfrak{A}}$ and let $\mathfrak{B}$ be τ-structure with B having two elements where $0^{\mathfrak{B}} \neq 1^{\mathfrak{B}}$. Then $\mathfrak{A} \approx_0 \mathfrak{B}$ but $\mathfrak{A} \not\sim_0 \mathfrak{B}$. To see that $\mathfrak{A} \approx_0 \mathfrak{B}$ we must see that $\mathfrak{A}$ and $\mathfrak{B}$ agree on all *unnested* atomic sentences. But there are none! So they vacuously agree. In the other game however, the (nested) atomic sentence $0 = 1$ is satisfied by $\mathfrak{A}$ but not by $\mathfrak{B}$. Note however, that the unnested Ehrenfeucht-Fraïssé game *can* tell $\mathfrak{A}$ and $\mathfrak{B}$ apart at level 1, i.e. $\mathfrak{A} \not\approx_1 \mathfrak{B}$. To see this suppose $\forall$ picks $a \in A$ then $\exists$ must pick $b \in B$. But then thinking of the formula $x{=}0$ we see that $\mathfrak{A} \models a{=}0$ and $\mathfrak{B} \not\models b{=}0$. So $\exists$ cannot win $\mathrm{EF}_1[\mathfrak{A}, \mathfrak{B}]$.

Question. What is the relation between $\approx_\omega$ and $\sim_\omega$? I.e. do there exist $\mathfrak{A}$ and $\mathfrak{B}$ such that $\mathfrak{A} \approx_\omega \mathfrak{B}$ and $\mathfrak{A} \not\sim_\omega \mathfrak{B}$?

Remark. There do exist $\mathfrak{A}$ and $\mathfrak{B}$ τ-structures (for finite signature τ) such that $\forall k < \omega\ \mathfrak{A} \approx_k \mathfrak{B}$ *but* $\mathfrak{A} \not\approx_\omega \mathfrak{B}$.

Here is an example. Let $\mathfrak{A} = (\mathbb{N}, <)$ and $\mathfrak{B} = (\mathbb{N} \oplus \mathbb{Z}, <)$ (where $\mathbb{N} \oplus \mathbb{Z}$ is the order gotten by adding a copy of $\mathbb{Z}$ after $\mathbb{N}$). Then $\mathfrak{A} \approx_k \mathfrak{B}$ for all $k < \omega$, *but* $\mathfrak{A} \not\approx_\omega \mathfrak{B}$. To see that $\mathfrak{A} \not\approx_\omega \mathfrak{B}$ imagine the case where $\forall$ picks all elements of $\mathbb{Z}$ (from $\mathfrak{B}$) doing down, then $\exists$ will run out of elements in $\mathbb{N}$ (from $\mathfrak{A}$) to pick.

Definition. For φ a formula, the **quantifier rank** $\mathrm{qr}(\varphi)$ is the number of nested quantifiers in φ. I.e.:

- If φ is atomic, then $\mathrm{qr}(\varphi) = 0$
- $\mathrm{qr}(\varphi \wedge \psi) = \mathrm{qr}(\varphi \vee \psi) = \max\{\mathrm{qr}(\varphi), \mathrm{qr}(\psi)\}$
- $\mathrm{qr}(\neg\varphi) = \mathrm{qr}(\varphi)$
- $\mathrm{qr}(\exists\varphi) = \mathrm{qr}(\varphi) + 1$.

Theorem 10.2. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be τ-structures. Then $\mathfrak{A} \equiv \mathfrak{B}$ if and only if for all finite $\tau' \subseteq \tau$ we have $\mathfrak{A}|_{\tau'} \approx_k \mathfrak{B}|_{\tau'}$ for all $k < \omega$.*

Remark. Clearly $\mathfrak{A} \equiv \mathfrak{B}$ iff for all finite $\tau' \subseteq \tau$ $\mathfrak{A}|_{\tau'} \equiv \mathfrak{B}|_{\tau'}$. Thus it suffices to prove the theorem in the case where τ is finite. The statement then becomes that $\mathfrak{A} \equiv \mathfrak{B}$ iff $\mathfrak{A} \approx_k \mathfrak{B}$ for all $k < \omega$.

Before giving a proof of this theorem we need an important lemma. Hodges calls it the *Fraïssé-Hintikka theorem* and notes that it is “the fundamental theorem about the equivalence relations $\approx_k$”. The theorem will follow as a corollary of the lemma.

Lemma. *For a finite signature τ and $k, n < \omega$, there is a finite set $\Theta_{n,k}$ of unnested formulae of quantifier rank $\leqslant k$ in n free variables $x_0, \dots, x_{n-1}$, such that*

0. *Distinct elements of $\Theta_{n,k}$ are inconsistent, i.e. for any $\eta, \theta \in \Theta_{n,k}$ then*

$$\models \forall \bar{x}(\eta \to \neg\theta)\,.$$[3]

1. *If $\varphi \in \mathcal{L}(\tau)$ has quantifier rank $\leqslant k$ and free variables $x_0, \dots, x_{n-1}$ then there is some subset $\Phi \subseteq \Theta_{n,k}$ such that $\models \forall \bar{x}(\varphi \leftrightarrow \bigvee \Phi)$.*
2. *Given $\mathfrak{A}, \mathfrak{B} \in \mathrm{Str}(\tau)$ then for any n-tuples $\bar{a} \in A^n$ and $\bar{b} \in B^n$, we have $\mathfrak{A}_{\bar{a}} \approx_k \mathfrak{B}_{\bar{b}}$ if and only if for each $\theta \in \Theta_{n,k}$,*

$$\mathfrak{A} \models \theta(\bar{a}) \qquad \Longleftrightarrow \qquad \mathfrak{B} \models \theta(\bar{b}).$$

Notation. For φ a formula define $\varphi^{[0]} := \varphi$ and $\varphi^{[1]} := \neg\varphi$.

Proof. We first construct $\Theta_{n,k}$ by recursion on k. [Note: this does not mean that we fix n. In the induction step we will use the $n+1$ level]

Let Φ be the set of unnested atomic formulae in $\mathcal{L}(\tau)$ in variables $x_0, \dots, x_{n-1}$. This set is finite. This is because τ is finite and to construct an unnested atomic formulae we are only allowed to introduce one symbol from τ.

To get $\Theta_{n,0}$ we will go through every way one might choose to make each instance of $\varphi \in \Phi$ either true or false, and then take conjunctions of these formulae. More precisely we let

$$\Theta_{n,0} := \left\{ \bigwedge_{\varphi \in \Phi} \varphi^{[s(\varphi)]} \;\middle|\; s : \Phi \longrightarrow \{0,1\} \right\}$$

and then

$$\Theta_{n,k+1} := \left\{ \bigwedge_{\varphi \in Y} \exists x_n \ \varphi(\bar{x}, x_n) \wedge \bigwedge_{\psi \in Z} \forall x_n \ \neg\psi(\bar{x}, x_n) \;\middle|\; \text{for } Y, Z \text{ a partition of } \Theta_{n+1,k} \right\}.$$

This finishes the construction of the sets $\Theta_{n,k}$ for arbitrary $n, k < \omega$.

Now we must check that conditions *0)*, *1)* , and *2)* are satisfied. First note that $\Theta_{n,k}$ is indeed finite and all elements are unnested and have quantifier rank $\leqslant k$.

0. Condition *0)* is reasonably clear. For $k = 0$ and $s, t : \Phi \longrightarrow \{0,1\}$ with $s \neq t$ there is some φ such that $s(\varphi) \neq t(\varphi)$ then

$$\left(\bigwedge_{\psi \in \Phi} \psi^{[s(\psi)]} \right) \longrightarrow \varphi^{[s(\varphi)]}$$

[3] The notation $\models \psi$ just means that for every τ-structure $\mathfrak{A}$ we have $\mathfrak{A} \models \psi$.

and

$$\left(\bigwedge_{\psi\in\Phi} \psi^{[t(\psi)]}\right) \longrightarrow \varphi^{[t(\varphi)]}.$$

Since $\varphi^{[s(\varphi)]}$ and $\varphi^{[t(\varphi)]}$ are explicitly inconsistent we see that $\bigwedge \psi^{[s(\psi)]}$ and $\bigwedge \psi^{[t(\psi)]}$ are inconsistent as well.

For the level $k+1$, suppose $Y \neq Y'$ and let $\eta \in Y \setminus Y'$. Now consider two formulae from $\Theta_{n,k+1}$. Then

$$\left(\bigwedge_{\varphi\in Y} \exists x_n\ \varphi \wedge \bigwedge_{\psi\in Y^c} \forall x_n\ \neg\psi\right) \longrightarrow \exists x_n\ \eta$$

whereas

$$\left(\bigwedge_{\varphi\in Y'} \exists x_n\ \varphi \wedge \bigwedge_{\psi\in (Y')^c} \forall x_n\ \neg\psi\right) \longrightarrow \forall x_n \neg\eta.$$

and since $\forall x_n \neg\eta \leftrightarrow \neg\exists x_n\ \eta$ we have an explicit inconsistency. By induction condition *0)* holds for all the sets $\Theta_{n,k}$.

1. To see that *1)* holds, note that if φ is an unnested formula of quantifier rank 0 in n free variables, then φ is a boolean combination of elements of Φ and so equivalent to some element of $\Theta_{n,0}$.
   [Case $qr(\varphi) \leqslant k+1$?????]
2. We show that condition *2)* holds by induction on k.
   - For $k = 0$. $(\mathfrak{A}, \bar{a}) \approx_0 (\mathfrak{B}, \bar{b})$ means that for ψ an unnested atomic τ-formula, $\mathfrak{A} \models \psi(\bar{a}) \iff \mathfrak{B} \models \psi(\bar{b})$. But the formulae in $\Theta_{n,0}$ are exactly the atoms in the boolean algebra generated by unnested atomic formulae. So if $(\mathfrak{A}, \bar{a})$ and $(\mathfrak{B}, \bar{b})$ agree on the unnested atomic formulae then they will agree on all elements of $\Theta_{n,0}$, and vice versa.
   - At stage $k+1$ we will take one implication at a time. First suppose $(\mathfrak{A}, \bar{a}) \approx_{k+1} (\mathfrak{B}, \bar{b})$. We show that for all $\varphi \in \Theta_{n,k+1}$, $\mathfrak{A} \models \varphi(\bar{a})$ implies $\mathfrak{B} \models \varphi(\bar{b})$. By symmetry we will also get that $\mathfrak{B} \models \varphi(\bar{b})$ implies $\mathfrak{A} \models \varphi(\bar{a})$.
     Let $\varphi \in \Theta_{n,k+1}$. By construction φ is
     $$\bigwedge_{\eta\in Y} \exists x_n\ \eta \wedge \bigwedge_{\xi\in Y^c} \forall x_n\ \neg\xi$$
     for some subset $Y \subseteq \Theta_{n+1,k}$. Suppose $\mathfrak{A} \models \varphi(\bar{a})$. For $\eta \in Y$ this implies that $\mathfrak{A} \models \exists x_n \eta(\bar{a}, x_n)$. Let $c \in A$ be a witness to this, i.e. $\mathfrak{A} \models \eta(\bar{a}, c)$. By hypothesis there exists some $d \in \mathfrak{B}$ such that $(\mathfrak{A}, \bar{a}, c) \approx_k (\mathfrak{B}, \bar{b}, d)$. Then by the induction hypothesis $\mathfrak{B} \models \eta(\bar{b}, d)$ so $\mathfrak{B} \models \exists x_n \eta(\bar{b}, x_n)$.

So for each $\eta \in Y$ we have $\mathfrak{B} \models \exists x_n \eta(\bar{b}, x_n)$. Likewise for $\xi \in Y^c$, if $\mathfrak{B} \not\models \forall x_n \neg\xi(\bar{b}, x_n)$ then $\mathfrak{B} \models \exists x_n \xi(\bar{b}, x_n)$ and by same argument we have that $\mathfrak{A} \models \exists x_n \xi(\bar{a}, x_n)$. Since this is not true by assumption we must have $\mathfrak{B} \models \forall x_n \neg\xi(\bar{b}, x_n)$. Thus $\mathfrak{B} \models \varphi(b)$. By symmetry of the roles of $\mathfrak{A}$ and $\mathfrak{B}$ we have that $\mathfrak{A} \models \varphi(\bar{a})$ iff $\mathfrak{B} \models \varphi(\bar{b})$ for all $\varphi \in \Theta_{n,k+1}$.
Now for the converse implication. Suppose $(\mathfrak{A}, \bar{a})$ and $(\mathfrak{B}, \bar{b})$ agree on all of the $\Theta_{n,k+1}$ formulae. We must show $\mathfrak{A} \approx_{k+1} \mathfrak{B}$, i.e. that $\exists$ has a winning strategy in $\mathrm{EF}_{k+1}[\mathfrak{A}, \mathfrak{B}]$. Suppose $\forall$ plays $c \in A$. As $\Theta_{n+1,k}$ partitions A^{n+1} (by property *1)* of this lemma), there is exactly one formula $\eta \in \Theta_{n+1,k}$ such that $\mathfrak{A} \models \eta(\bar{a}, c)$. Now as $\Theta_{n,k+1}$ partitions A^n there is exactly one formula $\varphi \in \Theta_{n,k+1}$ such that $\mathfrak{A} \models \varphi(\bar{a})$. Then

$$\varphi(\bar{x}) \longrightarrow \exists x_n \eta(\bar{x}, x_n)$$

since φ either implies $\exists x_n \eta(\bar{x}, x_n)$ or $\forall x_n \neg\eta(\bar{x}_n, n)$, but we know that $\mathfrak{A} \models \eta(\bar{a}, c)$. By hypothesis $(\mathfrak{A}, \bar{a})$ and $(\mathfrak{B}, \bar{b})$ agree on formulae from $\Theta_{n,k+1}$ so $\mathfrak{B} \models \varphi(\bar{b})$. This in turn implies that $\mathfrak{B} \models \exists x_n \eta(\bar{b}, x_n)$. Let $d \in B$ be a witness. Then $\exists$ will play d. By the induction hypothesis $(\mathfrak{A}, \bar{a}, c) \approx_k (\mathfrak{B}, \bar{b}, d)$. Likewise if $\forall$ picks some $d \in B$ then $\exists$ can find $c \in A$ such that $(\mathfrak{A}, \bar{a}, c) \approx_k (\mathfrak{B}, \bar{b}, d)$. Thus $(\mathfrak{A}, \bar{a}) \approx_{k+1} (\mathfrak{B}, \bar{b})$.
By induction we now have the desired equivalence.

□

We can now prove the theorem as a corollary. For convenience we state the result again.

Theorem 10.3. *For τ finite and $\mathfrak{A}, \mathfrak{B} \in \mathrm{Str}(\tau)$ the following are equivalent.*

- $\mathfrak{A} \equiv \mathfrak{B}$
- $\mathfrak{A} \approx_k \mathfrak{B}$ *for all* $k < \omega$.

Proof. Suppose first that $\mathfrak{A} \equiv \mathfrak{B}$. We show by induction on k that $\mathfrak{A} \approx_k \mathfrak{B}$ for all k. For $k = 0$ we have $\mathfrak{A} \equiv \mathfrak{B}$ implies $\mathfrak{A} \sim_0 \mathfrak{B}$, in particular $\mathfrak{A} \approx_0 \mathfrak{B}$.

Now for $k+1$. Suppose $\forall$ picks $b \in B$. Let $\varphi \in \Theta_{1,k}$ be the unique element of $\Theta_{1,k}$ such that $\mathfrak{B} \models \varphi(b)$. Then $\mathfrak{B} \models \exists x_0 \varphi(x_0)$. This is a sentence, and so by assumption $\mathfrak{A} \models \exists x_0 \varphi(x_0)$. Let $a \in A$ be a witness. Thus $\mathfrak{A} \models \varphi(a)$. So $(\mathfrak{A}, a) \models \psi$ if and only if $(\mathfrak{B}, b) \models \psi$ for all $\psi \in \Theta_{1,k+1}$ (since both $\mathfrak{A}$ and $\mathfrak{B}$ don't satisfy any other of the elements of $\Theta_{1,k+1}$ apart from φ). Now by property *2)* of the lemma we have that $(\mathfrak{A}, a) \approx_k (\mathfrak{B}, b)$. Now since b was arbitrary (and the roles for $\mathfrak{A}$ and $\mathfrak{B}$ were unimportant) we have $\mathfrak{A} \approx_{k+1} \mathfrak{B}$. By induction we are done.

Conversely, suppose $\mathfrak{A} \approx_k \mathfrak{B}$ for all $k < \omega$. We must show that $\mathfrak{A} \equiv \mathfrak{B}$. We show by induction on r that if $\varphi \in \mathcal{L}(\tau)$ is unnested and $\mathrm{qr}(\varphi) \leqslant r$ then $\mathfrak{A}$ and $\mathfrak{B}$ agree on φ. Since we have already seen that all formulae are equivalent to unnested formulae this will finish the proof.

For $r = 0$, φ is an unnested atomic formula. Then since $\mathfrak{A} \approx_0 \mathfrak{B}$, $\mathfrak{A}$ and $\mathfrak{B}$ must agree on φ. Similarly for φ a boolean combination of unnested atomic formulae.

For $r + 1$, suppose φ is $\exists x \theta(x)$ with $\mathrm{qr}(\theta) \leqslant r$. Suppose $\mathfrak{A} \models \varphi$ and let $a \in A$ be a witness, i.e. $\mathfrak{A} \models \theta(a)$. Let $\psi \in \Theta_{1,r}$ be such that $\mathfrak{A} \models \psi(a)$. ψ is unique by *1)* above. Since $\mathfrak{A} \approx_{r+1} \mathfrak{B}$ there exists $b \in B$ such $(\mathfrak{A}, a) \approx_k (\mathfrak{B}, b)$, i.e. $\mathfrak{B} \models \psi(b)$. But since $\mathrm{qr}(\theta) \leqslant r$ we have by property *1)* of the lemma, that $\theta \leftrightarrow \bigvee_{\eta \in Y} \eta$ for some $Y \subseteq \Theta_{1,r}$. Thus $\psi(b) \longrightarrow \theta(b)$. So $\mathfrak{B} \models \exists x \theta(x)$, i.e. $\mathfrak{B} \models \varphi$. This completes the proof. $\square$

Lecture 11

11.1 Games and Products

Last time we proved that for finite signatures elementary equivalence is equivalent to $\exists$ having a winning strategy in the $\mathrm{EF}_k[\mathfrak{A}, \mathfrak{B}]$ for all $k < \omega$.

Corollary. *If τ has no relation symbols and $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{B}_1$ and $\mathfrak{B}_2$ are τ-structures such that $\mathfrak{A}_1 \equiv \mathfrak{B}_1$ and $\mathfrak{A}_2 \equiv \mathfrak{B}_2$, then $\mathfrak{A}_1 \times \mathfrak{A}_2 \equiv \mathfrak{B}_1 \times \mathfrak{B}_2$.*

Proof. (Sketch) This basically follows from the fact that

$$\mathrm{EF}_k[\mathfrak{A}_1 \times \mathfrak{A}_2, \mathfrak{B}_1 \times \mathfrak{B}_2] = \mathrm{EF}_k[\mathfrak{A}_1, \mathfrak{B}_1] \times \mathrm{EF}_k[\mathfrak{A}_2, \mathfrak{B}_2].$$

This depends on the "correct" (i.e. categorical) definition of the products and of equality between games. $\square$

Hodges gives an application of this result to groups.

Corollary. *Let G_1, G_2 be elementarily equivalent groups, and let H be some group. Then $G_1 \times H$ is elementarily equivalent to $G_2 \times H$.*

The notes for this lecture are very short. This is because the material from the first part of lecture 11 was incorporated into the notes for lecture 10, and similarly the last part of lecture 11 was incorporated into the notes for lecture 12.

Lecture 12

12.1 Quantifier Elimination for $\mathrm{Th}(\mathbb{Z})$ as an Ordered Group

Let $\tau = \{+, -, 0, 1, <\}$ where $+$ and $-$ are binary function symbols, 0 and 1 are constants and $<$ is a binary relation. We will consider the theory of the integers as a discretely ordered group. We claim that this theory is axiomatized by the theory T of discretely ordered abelian groups G having $G/nG \cong \mathbb{Z}/n\mathbb{Z}$ for each $n \in \mathbb{Z}$. More precisely let T be the theory with the following axioms:

- ordered abelian group axioms as usual
- discretely ordered: $\forall x\, \neg(0 < x < 1)$
- $0 < 1$
- for each $n \in \mathbb{Z}$
$$\forall x \left(\bigvee_{j=0}^{n-1} \exists y (x = j + ny) \right)$$
where j is short for $1 + 1 + \cdots + 1$ (j times) and ny is short for $y + y + \cdots + y$ (n times).

Let $G \models T$, then since G is discretely ordered G/nG is generated by 1 and so the last axiom schema forces that $G/nG \cong \mathbb{Z}/n\mathbb{Z}$ for all n.

Definition. We define the **complexity**, $c(t)$ of a τ-term t to be essentially the number of additions in t. More precisely let

- $c(0) = c(1) = 1$.
- $c(x_i) = 1$.
- $c(-t) = c(t)$.
- $c(t + s) = c(t) + c(s)$.

Now we define the elimination set.

Definition. Let $\Xi_{n,k}$ be the set of formulae in n variables $x_0, \ldots, x_{n-1}$ of the form

- $t(\bar{x}) > 0$ where $c(t) \leqslant k$
- $s(\bar{x}) \equiv j \pmod{k!}$ where $c(s) \leqslant k$.

We let Ξ be the union $\bigcup \Xi_{n,k}$.

Theorem 12.1. *(Presburger)* T *axiomatizes* $\mathrm{Th}(\mathbb{Z}, +, -, 0, 1, <)$ *and* Ξ *is an elimination set.*

In particular T is a *complete* theory. The proof of the theorem will actually yield an effective procedure for converting a general formula to an equivalent formula in Ξ. So we will get *decidability* for the theory as well.

Remark. In fact the decidability result for T follows (by Gödel's Completeness Thoerem) from the first statement since $\mathrm{Th}(\mathbb{Z}, +, -, 0, 1, <)$ is complete.

The general approach of the proof will be the following. We show that equivalence relative to Ξ can be used to set up a back-and-forth system. This we know is enough to determine elementary equivalence, which gives the first statement. We also know that every formula is equivalent to a disjunction of formulae from the set $\Theta = \bigcup \Theta_{n,k}$ (constructed in Lecture 10). So if we can show that the equivalence relation given from Ξ is finer that the relation given by Θ then every element of Θ can be expressed as a disjunction of elements of Ξ. Since Θ was enough for an elimination set, it follows that Ξ is too.

Let us first define the "equivalence relation given by Ξ".

Definition. For $\mathfrak{A}$ and $\mathfrak{B}$ models of T and for $\bar{a} \in A^n$ and $\bar{b} \in B^n$, we say that

$$(\mathfrak{A}, \bar{a}) \sim_k^{\Xi} (\mathfrak{B}, \bar{b})$$

iff for all $\xi \in \Xi_{n,k}$ we have $\mathfrak{A} \models \xi(\bar{a}) \iff \mathfrak{B} \models \xi(\bar{b})$.

Now our goal is to show that there is a sequence of numbers $(k_i)_{i=0}^{\infty}$ such that $k_0 < k_1 < \cdots$ and such that we can carry out the back-and-forth construction if we know that we have the $\sim_k^{\Xi}$ for all k. More precisely we want

- if $(\mathfrak{A}, \bar{a}) \sim_{k_0}^{\Xi} (\mathfrak{B}, \bar{b})$ then $(\mathfrak{A}, \bar{a}) \approx_0 (\mathfrak{B}, \bar{b})$, and
- if $(\mathfrak{A}, \bar{a}) \sim_{k_{i+1}}^{\Xi} (\mathfrak{B}, \bar{b})$ and if $c \in A$ then there exists $d \in B$ such that $(\mathfrak{A}, \bar{a}, c) \sim_{k_i}^{\Xi} (\mathfrak{B}, \bar{b}, d)$. Vice versa: if $d \in B$ then there exists $c \in A$ such that $(\mathfrak{A}, \bar{a}, c) \sim_{k_i}^{\Xi} (\mathfrak{B}, \bar{b}, d)$.

The existence of such a sequence (k_i) will then imply that $\sim_k^{\Xi}$ is finer than $\approx_i$ which is what we want. We need two technical lemmas. The first shows that k_0 may be chosen to be 3.

Lemma. *If* $(\mathfrak{A}, \bar{a}) \sim_3^{\Xi} (\mathfrak{B}, \bar{b})$ *then* $(\mathfrak{A}, \bar{a}) \approx_0 (\mathfrak{B}, \bar{b})$

Proof. We must check for all unnested atomic formulae: $x_j < x_i$, $x_k = x_i + x_j$, $x_i = 0$, $x_j = 1$, and $x_i = x_j$. As an example we check $\mathfrak{A} \models a_k = a_i + a_j \iff \mathfrak{B} \models b_k = b_i + b_j$. Using the axioms of ordered abelian groups

$$a_k{=}a_i{+}a_j \quad \iff \quad a_k{-}(a_i{+}a_j){=}0 \quad \iff \quad \neg(a_k{-}(a_i{+}a_j){>}0) \wedge \neg((a_i{+}a_j){-}a_k{>}0).$$

Let $t(\bar{x}) := x_k - (x_i + x_j)$ and $u(\bar{x}) := -t(\bar{x})$ be terms. Both have complexity 3. By hypothesis $\neg(t(\bar{a}) > 0$ and so by definition of $\sim_3^{\Xi}$ we have $(\mathfrak{B}, \bar{b}) \models \neg(t(\bar{x}) > 0)$ and likewise $(\mathfrak{B}, \bar{b}) \models \neg(u(\bar{x}) > 0)$ so $b_k = b_i + b_k$. □

So we let $k_0 := 3$. Now to go up a step is a bit more complicated. We shall let $k_m := m^{2m}$. This suffices by the following lemma.

Lemma. *If* $(\mathfrak{A}, \bar{a}) \sim_{m^{2m}}^{\Xi} (\mathfrak{B}, \bar{b})$ *(*$m \geqslant 3$*) and* $c \in A$ *then there exists* $d \in B$ *such that* $(\mathfrak{A}, \bar{a}, c) \sim_m^{\Xi} (\mathfrak{B}, \bar{b}, d)$*. Similarly if* $d \in B$ *then there exists* $c \in A$ *such that* $(\mathfrak{A}, \bar{a}, c) \sim_m^{\Xi} (\mathfrak{B}, \bar{b}, d)$.

Proof. We will deal with congruence issues first and the with order issues second.

Let $c \in A$. We want to understand the congruence relations that c might have relative to terms when we plug in $\bar{a}$. We only consider terms of complexity $m - 1$. Consider the set

$$\begin{aligned} \Gamma := \quad & \{t(\bar{x}) + ix_n \equiv j \pmod{m!} \mid \\ & c(t) \leqslant m-1,\ i \leqslant m,\ 0 \leqslant j \leqslant m! \text{ and } \mathfrak{A} \models t(\bar{a}) + ic \equiv j \pmod{m!}\} \end{aligned}$$

As $(\mathfrak{A}, \bar{a}) \sim_{m^{2m}}^{\Xi} (\mathfrak{B}, \bar{b})$, for each t of complexity $\leqslant m-1$ we have $t(\bar{a}) \equiv t(\bar{b}) \pmod{m!}$. This statement makes sense since $\mathfrak{A}/m!\mathfrak{A} \cong \mathbb{Z}/m!\mathbb{Z}$ and $\mathbb{Z}/m!\mathbb{Z} \cong \mathfrak{B}/m!\mathfrak{B}$, so we can identify elements of $\mathfrak{A}$ and $\mathfrak{B}$ with their image under the isomorphism. Let $\alpha : \mathfrak{A}/m!\mathfrak{A} \longrightarrow \mathbb{Z}/m!\mathbb{Z}$ and $\beta : \mathfrak{B}/m!\mathfrak{B} \longrightarrow \mathbb{Z}/m!\mathbb{Z}$ be the isomorphisms. Now since $\alpha(c)$ satisfies all formulae of Γ we have that $e := \beta^{-1}(\alpha(c)) \in \mathfrak{B}$ also satisfies all formulae in Γ. So we have found $e \in \mathfrak{B}$ which looks like c up to congruence mod $m!$. Without loss of generality we can assume $0 \geqslant e < m!$.

Our final goal is to modify e, while preserving its congruence mod $m!$ so that it also looks like c in the ordering. We must find $f \in \mathfrak{B}$ such that $d = e + f(m!)$ works.

We deal with assertions of the form

$$t(\bar{a}) + ic > 0$$

where the complexity of t is $\leqslant m-1$ and $0 < i \leqslant m$. Multiplying by $\frac{m!}{i}$ we reduce to assertions of the form

$$\frac{m!}{i}t(\bar{a}) + m!c > 0.$$

Setting $u(\bar{a}) := \frac{m!}{i}t(\bar{a}) + m!c$ we have that the complexity of u is $\leqslant (m-1)m! < m^{2m}$. Consider the set

$$\{t(\bar{a}) \mid c(t) \leqslant (m-1)m!\}.$$

This is a finite set. Let $t(\bar{a})$ be chosen from this set so that $t(\bar{a}) < m!c$ is maximally so (i.e. there is no other term $t'(\bar{a})$ such that $t(\bar{a}) < t'(\bar{a}) < m!c$). Similarly let $u(\bar{a})$ be chosen so that $u(\bar{a}) \geqslant m!c$ minimally so. If one of t or u doesn't exist, then we just ignore the corresponding part of the following argument. Now we have

$$t(\bar{a}) \leqslant m!c \leqslant u(\bar{a}).$$

Since $(\mathfrak{A}, \bar{a}) \sim^{\Xi}_{m^{2m}} (\mathfrak{B}, \bar{b})$ we have that

$$t(\bar{a}) \equiv t(\bar{b}) \pmod{(m!)^2}$$

and

$$u(\bar{a}) \equiv u(\bar{b}) \pmod{(m!)^2}.$$

Also, by choice of e we have

$$m!c \equiv m!e \pmod{(m!)^2}.$$

Thus there exists $g \in \mathfrak{B}$ such that

$$g \equiv m!e \equiv m\ c \pmod{(m!)^2}.$$

Letting $d := \frac{g}{m!}$ gives the desired element of $\mathfrak{B}$, so that $(\mathfrak{A}, \bar{a}, c) \sim^{\Xi}_{m} (\mathfrak{B}, \bar{b}, d)$. This completes the proof. $\square$

The theorem now follows from the lemmas and the remarks above.

12.2 Automorphisms

We move on to discuss the relationship between reduces (and expansions) and automorphisms.

We will need a topology on our groups of automorphisms.

Definition. Given a set X let $\mathrm{Sym}(X) := \{\sigma : \sigma : X \to X \text{ is a bijection}\}$, be the group of permutations of X.

Remark. $\mathrm{Sym}(X)$ be be regarded as the automorphism group of the structure $\mathfrak{X}$ in the empty signature, with $\mathrm{dom}(\mathfrak{X}) = X$.

Notation. For $\sigma \in \mathrm{Sym}(X)$ and $\bar{a} \in X^n$ we write $\sigma\bar{a}$ for $(\sigma(a_0), \ldots, \sigma(a_{n-1}))$.

$\mathrm{Sym}(X)$ has a topology on it.

Definition. The **basic open sets** $U_{\bar{a},\bar{b}}$ in $\mathrm{Sym}(X)$ have the form

$$U_{\bar{a},\bar{b}} := \{\sigma \in \mathrm{Sym}(X) : \sigma\bar{a} = \bar{b}\}$$

for $\bar{a}, \bar{b} \in X^n$. The open sets of the topology are the unions of the basic open sets.

Remark. $U_{\bar{a},\bar{b}}$ are actually closed since

$$\mathrm{Sym}(X) \setminus U_{\bar{a},\bar{b}} = \bigcup_{\bar{c} \neq \bar{b}} U_{\bar{a},\bar{c}}.$$

So the sets $U_{\bar{a},\bar{b}}$ are *clopen*.

Remark. $U_{\bar{a},\bar{b}}$ is a coset of the stabilizer subgroup $\mathrm{Sym}(X)_{\bar{a}}$ (and also a coset of $\mathrm{Sym}(X)_{\bar{b}}$).

Remark. The point sets are closed. I.e. for any $\sigma \in \mathrm{Sym}(X)$

$$\{\sigma\} = \bigcap_{a \in X} U_{a,\sigma(a)}$$

is closed.

Remark. The topology we have given makes the action

$$\mu : \mathrm{Sym}(X) \times X \longrightarrow X$$

continuous when X is given the discrete topology. In fact it is the coarsest such topology. To see this let $V \subseteq X$ be a basic open set, i.e. $V = \{x\}$ for some $x \in X$. Then

$$\mu^{-1}(V) = \{(\sigma, y) | \sigma(y) = x\} = \bigcup_{y \in X} U_{y,x} \times \{x\}$$

which is open in the product topology.

If $\mathfrak{A}$ is a τ-structure then $\mathrm{Aut}(\mathfrak{A})$ is a subgroup of $\mathrm{Sym}(A)$. More generally if $\mathfrak{A}'$ is a τ'-structure and $\tau \subseteq \tau'$ then $\mathrm{Aut}(\mathfrak{A}')$ is a subgroup of $\mathrm{Aut}(\mathfrak{A}'|_\tau)$.

Theorem 12.2. $\mathrm{Aut}(\mathfrak{A})$ *is a closed subgroup of* $\mathrm{Sym}(A)$.

Proof. Let $\sigma \in \overline{\mathrm{Aut}(\mathfrak{A})}$. We want to show that $\sigma \in \mathrm{Aut}(\mathfrak{A})$. Let $\varphi(\bar{x})$ be any $\mathcal{L}(\tau)$-formula. We must show that for any $\bar{a}$ from $\mathfrak{A}$

$$\mathfrak{A} \models \varphi(\bar{a}) \qquad \Longleftrightarrow \qquad \mathfrak{A} \models \varphi(\sigma\bar{a}).$$

Suppose $\mathfrak{A} \models \varphi(\bar{a})$. Let $\bar{b} := \sigma\bar{a}$. Since $\sigma \in \overline{\mathrm{Aut}(\mathfrak{A})}$ we have that $U_{\bar{a},\bar{b}} \cap \mathrm{Aut}(\mathfrak{A})$ is non-empty. Let δ be in this intersection. Then $\delta(\bar{a}) = \bar{b} = \sigma(\bar{a})$. Thus,

$$\mathfrak{A} \models \varphi(\bar{a}) \qquad \Longleftrightarrow \qquad \mathfrak{A} \models \varphi(\delta(\bar{a})) \qquad \Longleftrightarrow \qquad \mathfrak{A} \models \varphi(\sigma\bar{a}).$$

So $\sigma \in \mathrm{Aut}(\mathfrak{A})$. $\square$

Lecture 13

13.1 Automorphism Groups of Structures

Last time we defined a topology on $\mathrm{Sym}(X)$. We showed that the automorphism group of a structure, $\mathrm{Aut}(\mathfrak{A})$, is a closed subgroup of the permutation group of the domain, $\mathrm{Sym}(A)$. It follows that if $\mathfrak{A}^+$ is an expansion of $\mathfrak{A}$ then $\mathrm{Aut}(\mathfrak{A}^+)$ is a closed subgroup of $\mathrm{Aut}(\mathfrak{A})$ with respect to the subspace topology on $\mathrm{Aut}(\mathfrak{A})$. This has a converse. Any closed subgroup of $\mathrm{Aut}(\mathfrak{A})$ can be realized as the automorphism group of an extension of $\mathfrak{A}$.

Proposition. *Let $\mathfrak{A}$ be a τ-structre and H a closed subgroup of $\mathrm{Aut}(\mathfrak{A})$. Then there is an extension of signatures $\tau^+ \supseteq \tau$ and an extension $\mathfrak{A}^+$ of $\mathfrak{A}$ to τ^+, such that $H = \mathrm{Aut}(\mathfrak{A}^+)$.*

Proof. The action of H on the set $A = \mathrm{dom}(\mathfrak{A})$ is used to determine some new relation symbols. For each $n \in \omega$ and for each H-orbit $X \subseteq A^n$ let $R_X \in \mathcal{R}_{\tau^+}$ be a new relation symbol of arity n. We let $\mathfrak{A}^+$ be the extension of $\mathfrak{A}$ where

$$R_X^{\mathfrak{A}^+} := X.$$

We claim $H = \mathrm{Aut}(\mathfrak{A}^+)$.

First let $h \in H$. Let $R_X \in \tau^+$ be one of the new relation symbols, and let $a \in A^n$ such that $\mathfrak{A}^+ \models R_X(a)$. Then $X = H.a \subseteq A^n$. So $h.a \in X$, i.e. $\mathfrak{A}^+ \models R_X(ha)$. Conversely if $\mathfrak{A} \models R_X(ha)$ then $ha \in X$ and so $a = h^{-1}(ha) \in X$ as well. Since $h \in \mathrm{Aut}(\mathfrak{A})$ and h fixes all new relation symbols we have that $h \in \mathrm{Aut}(\mathfrak{A}^+)$. So $H \leqslant \mathrm{Aut}(\mathfrak{A}^+)$.

Now let $\sigma \in \mathrm{Aut}(\mathfrak{A}^+)$. We will show that σ is in $\overline{H}$ and so by assumption in H (since H is closed). Let $U \ni \sigma$ be an open set. We may assume $U = U_{a,b}$ for some $a, b \in A^n$ so that $\sigma a = b$. Let $X = Ha$ be the H-orbit of a. Since $\sigma \in \mathrm{Aut}(\mathfrak{A}^+)$ we have

$$\mathfrak{A}^+ \models R_X(a) \Longrightarrow \mathfrak{A}^+ \models R_X(\sigma a)$$

so $\mathfrak{A}^+ \models R_X(b)$ i.e. $b \in X$. So there is some $h \in H$ such that $ha = b$. But this means that $h \in H \cap U_{a,b}$. In particular $H \cap U_{a,b} \neq \varnothing$, so every open set containing σ meets H, so $\sigma \in \overline{H} = H$. $\square$

Remark. As previously mentioned, for a τ-structure $\mathfrak{A}$ and $a \in A^n$ the type of a written $\mathrm{tp}(a)$ is defined to be

$$\mathrm{tp}(a) := \mathrm{Th}(\mathfrak{A}_a) = \{\varphi(\bar{x}) \in \mathcal{L}(\tau) : \mathfrak{A} \models \varphi(\bar{a})\}$$

i.e. all formulae which are true (in $\mathfrak{A}$) of the tuple a.

If there is some $\sigma \in \mathrm{Aut}(\mathfrak{A})$ and $a, b \in A^n$ with $\sigma a = b$ then $\mathrm{tp}(a) = \mathrm{tp}(b)$.

The converse is true in the structure $\mathfrak{A}^+$ constructed in the above proof. I.e. $\mathrm{tp}(a) = \mathrm{tp}(b)$ if and only if there is some $\sigma \in \mathrm{Aut}(\mathfrak{A}^+)$ such that $\sigma a = b$. For general τ-structures, $\mathfrak{B}$, this is not the case. There may be tuples a and b with $\mathrm{tp}(a) = \mathrm{tp}(b)$ and yet $\mathrm{Aut}(\mathfrak{B})a \neq \mathrm{Aut}(\mathfrak{B})b$.

Example. Let $\tau = \{E\}$ be the theory of a single equivalence relation. Let $\mathfrak{A}$ be a τ-structure such that there are exactly two equivalence classes, one of size $\aleph_0$ and the other of size $\aleph_1$. Let a and b be elements of A which are in distinct equivalence classes. Then $\mathrm{tp}(a) = \mathrm{tp}(b)$. Yet there can be no automorphism carrying a to b, since such an automorphism would have to carry one equivalence class to the other.

Remark. In the expansion $\mathfrak{A}^+$ constructed in the proof of the proposition the relation $a \sim b$ iff $\mathrm{tp}(a) = \mathrm{t}(b)$ is definable. This is a very unusual property for a structure.

Question. What does it mean about the theory T if in every model of T (for all $n \in \omega$), the equivalence relation $a \sim b$ iff $\mathrm{tp}(a) = \mathrm{tp}(b)$ is definable?

For example in the theory of equality this is true. Also in the theory of dense linear orders it is true. The condition fails for the theory of the reals as a field.

This question will be answered later in the course.

In the following we need a general lemma about topological groups.

Lemma. *Let be G a topological group and $H \leqslant G$ a subgroup and $U \leqslant G$ an open subset. If $U \subseteq H$ then H is open.*

Proof. H is the union of cosets of U, i.e.

$$H = \bigcup_{h \in H} hU$$

and since multiplication by $h \in H$ is a homeomorphism of H it follows that hU is open. Thus H is open. $\square$

In the case where the structures under consideration are countable there is a tighter connection between structure and automorphisms.

Notation. For $b \in A^n$ we shall denote the stabilizer $\mathrm{Aut}(\mathfrak{A})_{(b)}$ by $\mathrm{Aut}(\mathfrak{A}/b)$.

Theorem 13.1. *Let $\mathfrak{A}$ be a countable τ-structure, and $H \leqslant \mathrm{Aut}(\mathfrak{A})$ a closed subgroup. The following are equivalent.*

i) H *is open.*
ii) $|\mathrm{Aut}(\mathfrak{A})/H| \leqslant \aleph_0$.
iii) $|\mathrm{Aut}(\mathfrak{A})/H| < 2^{\aleph_0}$.

Remark. Note that there are at most $2^{\aleph_0}$ elements of $|\mathrm{Aut}(\mathfrak{A})/H|$.

Proof. i) $\Rightarrow$ ii). If H is open then it contains a basic open set, i.e. there exists $a, b \in A^n$ such that $U_{a,b} \subseteq H$. Now as observed last time $U_{a,b}$ is a coset of the stabilizer of b. Since H is a group it must contain the stabilizer itself. Thus

$$\mathrm{Aut}(\mathfrak{A}/b) \leqslant H.$$

So

$$\mathrm{Aut}(\mathfrak{A}).b \cong \mathrm{Aut}(\mathfrak{A})/\mathrm{Aut}(\mathfrak{A}/b)$$

as $\mathrm{Aut}(\mathfrak{A})$-sets. So

$$|\mathrm{Aut}(\mathfrak{A}).b| = |\mathrm{Aut}(\mathfrak{A})/\mathrm{Aut}(\mathfrak{A}/b)| \geqslant |\mathrm{Aut}(\mathfrak{A})/H|.$$

But $A^n \supseteq \mathrm{Aut}(\mathfrak{A}).b$ so $|\mathrm{Aut}(\mathfrak{A})/H| \leqslant |A^n| \leqslant \aleph_0$.

ii) $\Rightarrow$ iii). Clear.

iii) $\Rightarrow$ i). This step will require some work. We shall prove the contrapositive. We assume that H is not open and use this to show that the index of H in $\mathrm{Aut}(\mathfrak{A})$ has size $2^{\aleph_0}$. We build a tree inside of $\mathrm{Aut}(\mathfrak{A})$ which remains a tree when we mod out by H.

We will construct a sequence $(a_i)_{i\in\omega}$ of finite sequences from A, and a sequence $(\sigma_i)_{i\in\omega}$ from $\mathrm{Aut}(\mathfrak{A})$. For each $T \subseteq \{0, 1, \ldots, n-1\}$, say $T = \{i_1 < i_2 < \cdots < i_l\}$ we define

$$\sigma_T := \prod_{i\in T} \sigma_i = \sigma_{i_1}\sigma_{i_2}\cdots\sigma_{i_l}.$$

We shall arrange that the following hold for the sequences $(a_i)_{i\in\omega}$ and $(\sigma_i)_{i\in\omega}$:

- for $i > j$ we have $\sigma_i(a_j) = a_j$.
- for each $n \in \omega$ and $S, T \subseteq \{0, 1, \ldots, n-1\}$ if $S \neq T$ then $\sigma_S \not\equiv \sigma_T \pmod{H}$ on $\{a_0, \ldots, a_{n-1}\}$.

Remark. If the first condition is satisfied then it does make sense to define σ_T even for infinite $T \subseteq \omega$, as long as we restrict to the a_i's. I.e. the map $\sigma_T : \{a_i\} \to A^*$ defined by $\sigma_T(x) = \prod_{i\in T} \sigma_i(x)$ is well-defined. Since $\mathrm{Aut}(\mathfrak{A})$ is closed in $\mathrm{Sym}(A)$

then for each $T \subseteq \omega$ there is some $\widetilde{\sigma_T} \in \mathrm{Aut}(\mathfrak{A})$ such that $\widetilde{\sigma_T}|_{\{a_i\}} = \sigma_T$. So if we succeed in constructing the sequences (a_i) and (σ_i) to satisfy the two conditions then we can find $2^{\aleph_0}$ many automorphisms which are different mod H.

We define the sequences by induction. The 0 case and the $n+1$ cases are the same. So we just do the $n+1$ step.

Suppose that $(a_i)_{i<n}$ and $(\sigma_i)_{i<n}$ have been obtained (for $n = 0$ this just means the sequences are empty). We look for a_n and σ_n.

By hypothesis H is *not* open. By the lemma H does not contain any open subgroups. In particular H does not contain the stabilizer of the sequence constructed thus far. I.e.

$$\mathrm{Aut}(\mathfrak{A}/(a_0, \ldots, a_{n-1})) \not\subseteq H$$

so there is some $\sigma_n \in \mathrm{Aut}(\mathfrak{A}/(a_0, \ldots, a_{n-1})) \setminus H$.

We now claim that there is some $a_n \in A^m$ such that for all $h \in H$ we have $h(a_n) \neq \sigma_n(a_n)$. This is true since if not, then for all $a, b \in A^m$ (and all $m \in \omega$) if

$$\sigma_n \in U_{a,b}$$

then there is some $h \in U_{a,b} \cap H$ so that σ_n is in the closure of H. But H was assumed closed! So $\sigma_n \in H$, which is a contradiction. So we can find a_n such that σ_n and h disagree on a_n (for all $h \in H$).

By induction we have defined the sequences $(a_i)_{i<\omega}$ and $(\sigma_i)_{i<\omega}$. We must check that they satisfy the two conditions. The first property is clear since $\sigma_n \in \mathrm{Aut}(\mathfrak{A}/(a_0, \ldots, a_{n-1}))$ so σ_n acts trivially on a_i for $i < n$.

To check the second property suppose $S, T \subseteq \{0, \ldots, n\}$ with $S \neq T$. Let j be the first place they differ. With out loss of generality assume $j \in S$.

If $j < n$ then

$$\sigma_S(a_j) = \sigma_{S\setminus\{n\}}(a_j) \quad \text{and} \quad \sigma_T(a_j) = \sigma_{T\setminus\{n\}}(a_j)$$

since σ_n acts on a_j trivially for $j < n$. By induction we may assume $\sigma_{S\setminus\{n\}}$ and $\sigma_{T\setminus\{n\}}$ are inequivalent mod H on the set $\{a_0, \ldots, a_{n-1}\}$.

Now suppose $j = n$. Then $S = T \cup \{n\}$. Suppose there is some $h \in H$ such that $\sigma_S = \sigma_T h$ on $\{a_0, \ldots, a_n\}$. Then

$$\sigma_T h(a_n) = \sigma_S(a_n) = \sigma_T \sigma_n(a_n)$$

so multiplying by σ_T^{-1} we have $h(a_n) = \sigma_n(a_n)$ which is a contradiction with the construction og σ_n.

By the earlier remarks we have now shown that $|\mathrm{Aut}(\mathfrak{A})/H| \geqslant 2^{\aleph_0}$. Since $\mathfrak{A}$ is countable we must therefore have that $|\mathrm{Aut}(\mathfrak{A})/H| = 2^{\aleph_0}$. □

The above result has a more model-theoretic interpretation which we now develop.

We have seen that closed subgroups of automorphism groups come from expansions. Suppose $\tau \subseteq \tau^+$ is an extension of signatures and $\mathfrak{A}$ is a τ-structure and $\mathfrak{A}^+$ is an expansion. For each τ-automorphism σ of $\mathfrak{A}$ we can find an extension $\mathfrak{A}^{+\sigma}$ of $\mathfrak{A}$ such that $\sigma : \mathfrak{A}^+ \longrightarrow \mathfrak{A}^{+\sigma}$ is an isomorphism of τ^+ structures. To define $\mathfrak{A}^{+\sigma}$ let

- $R^{\mathfrak{A}^{+\sigma}} := \sigma(R^{\mathfrak{A}^+})$ for all $R \in \mathcal{R}_{\tau^+}$.
- $c^{\mathfrak{A}^{+\sigma}} := \sigma(c^{\mathfrak{A}^+})$ for all $c \in \mathcal{C}_{\tau^+}$.
- $f^{\mathfrak{A}^{+\sigma}}(b) := \sigma(f^{\mathfrak{A}^+}(\sigma^{-1}(b)))$ for all $f \in \mathcal{F}_{\tau^+}$.

Conversely if $\widetilde{\mathfrak{A}}$ is an expansion of $\mathfrak{A}$ to τ^+ such that $\mathfrak{A}^+ \cong \widetilde{\mathfrak{A}}$ then there is some $\sigma \in \mathrm{Aut}(\mathfrak{A})$ such that $\widetilde{\mathfrak{A}} = \mathfrak{A}^{+\sigma}$ (just pick an isomorphism $\mathfrak{A}^+ \cong \widetilde{\mathfrak{A}}$).

Moreover $\mathfrak{A}^+ = \mathfrak{A}^{+\sigma}$ if and only if $\sigma \in \mathrm{Aut}(\mathfrak{A}^+)$. So we can identify the set of expansions of $\mathfrak{A}$ which are isomorphic to $\mathfrak{A}^+$, with the set of cosets $\mathrm{Aut}(\mathfrak{A})/\mathrm{Aut}(\mathfrak{A}^+)$.

So now restating the theorem in these terms we get.

Theorem 13.2. *Let $\tau \subseteq \tau^+$ be an expansion of signatures and $\mathfrak{A}$ a countable τ-structure, and $\mathfrak{A}^+$ and expansion to τ^+. The following are equivalent.*

i) $\mathrm{Aut}(\mathfrak{A}^+)$ *is an open subgroup of* $\mathrm{Aut}(\mathfrak{A})$.
ii) There are at most $\aleph_0$ expansions of $\mathfrak{A}$ which are isomorphic to $\mathfrak{A}^+$.
iii) There are strictly less than $2^{\aleph_0}$ expansions which are isomorphic to $\mathfrak{A}^+$.

furthermore if these conditions are satisfied then there exists $m \in \omega$ and $a \in A^m$ such that

$$\mathrm{Aut}(\mathfrak{A}/a) \subseteq \mathrm{Aut}(\mathfrak{A}^+).$$

As a corollary we get.

Corollary. *For $\mathfrak{A}$ a countable τ-structure the following are equivalent.*

i) $|\mathrm{Aut}(\mathfrak{A})| < \aleph_0$.
ii) $|\mathrm{Aut}(\mathfrak{A})| < 2^{\aleph_0}$.
iii) There is some m and $b \in A^m$ such that $\mathfrak{A}_b$ is rigid, i.e. $\mathrm{Aut}(\mathfrak{A}_b) = \{\mathrm{id}\}$.

Proof. We have seen the equivalence between *i)* and *ii)*.

Let $\tau^+ := \tau_A$ and $\mathfrak{A}^+ := \mathfrak{A}_A$, then $\mathrm{Aut}(\mathfrak{A}^+) = \{\mathrm{id}\}$. If $|\mathrm{Aut}(\mathfrak{A})| < 2^{\aleph_0}$ then by the theorem $\{\mathrm{id}\} \leqslant \mathrm{Aut}(\mathfrak{A}^+) \leqslant \mathrm{Aut}(\mathfrak{A})$ is open! So, again by the theorem, there is some $b \in A^m$ such that $\mathrm{Aut}(\mathfrak{A}/b) \leqslant \mathrm{Aut}(\mathfrak{A}^+) = \{\mathrm{id}\}$, os $\mathfrak{A}_b$ is rigid. ☐

Lecture 14

14.1 Interpretations

Definition. An interpretation Γ of the ρ-structure $\mathfrak{B}$ in the τ-structure $\mathfrak{A}$ is given by

- a τ-formula $\partial_\Gamma(x_0, \ldots, x_{l-1})$
- for each unnested ρ-atomic formula $\varphi(y_0, \ldots, y_{m-1})$ a τ-formula

$$\varphi_\Gamma(x_{0,0}, \ldots, x_{l-1,0}; x_{0,1}, \ldots, x_{l-1,1}; \ldots\ldots x_{0,m-1}, \ldots, x_{l-1,m-1})$$

- and a surjective function
$$\pi : \partial_\Gamma(\mathfrak{A}) \longrightarrow \mathfrak{B}$$
such that for all $a, b \in \partial_\Gamma(\mathfrak{A})$ then $\pi(a) = \pi(b)$ if and only if $\varphi_\Gamma(a, b)$ where φ is the ρ-atomic formula $y_0 = y_1$.

The condition on the map π is just that it pulls back the equality relation on $\mathfrak{B}$ to the interpretation (via Γ) of the equality relation on $\mathfrak{A}$.

We give a couple of examples.

Example. A classic example of an interpretation is that of the complex numbers in the reals. Here we interpret a complex number $z \in \mathbb{C}$ as a pair of real numbers (a, b) (which we think of as $z = a + ib$) with addition and multiplication defined appropriately.

Formally we interpret $(\mathbb{C}, +, \cdot, 0, 1)$ in $(\mathbb{R}, +, \cdot, 0, 1)$ as follows; Let $\partial_\Gamma(x_0, x_1)$ be any true statement for example $x_0 = x_0$. Thus $\partial_\Gamma(\mathbb{R}) = \mathbb{R}^2$. Here are some of the crucial interpretations of the unnested formulae

- $(y = 0)_\Gamma$ will be $(x_0 = 0 \wedge x_1 = 0)$.
- $(y = 1)_\Gamma$ will be $(x_0 = 1 \wedge x_1 = 0)$.
- $(y_2 = y_0 + y_1)_\Gamma$ will be $(x_{0,0} + x_{0,1} = x_{0,2} \wedge x_{1,0} + x_{1,1} = x_{1,2})$
- $(y_2 = y_0 \cdot y_1)_\Gamma$ will be $(x_{0,2} = x_{0,0}x_{0,1} - x_{1,0}x_{1,1} \wedge x_{1,0}x_{0,1} + x_{0,0}x_{1,1})$

finally the map $\pi : \partial_\Gamma(\mathbb{R}) \longrightarrow \mathbb{C}$ is given by $\pi(a_0, a_1) = a_0 + a_1\sqrt{-1}$. Here the usual equality relation in $\mathbb{C}$ pulls back to the coordinate-wise equality relation on $\partial_\Gamma(\mathbb{R}) = \mathbb{R}^2$ as it should!

Example. Set theory is stronger than arithmetic. I.e. we can also interpret arithmetic inside of set theory.

Let $\tau = \{\in\}$ be the signature of set theory and let $\mathbb{V}$ be a model of ZFC. Let ρ be the language of arithmetic, $\rho = \{\leqslant, +, \cdot, 0, 1\}$ and let $\mathfrak{B} = (\mathbb{N}, +, \cdot, 0, 1, \leqslant)$. We let $\partial_\Gamma(x)$ be the τ-formula which says "x is a natural number", this can be formally expressed in the language $\mathcal{L}(\tau)$ but we will not do so now. Now addition and multiplication can be given their usual set-theoretical interpretations (which again we will not properly write out).

Example. The example of $\mathbb{C}$ interpreted in $\mathbb{R}$ generalizes to any *finite* field extension. I.e. if L/K is a finite field extension then $(L, +, \cdot, 0, 1)$ is interpretable in $(K, +, \cdot, 0, 1)$.

In the definition of interpretation we only required there to be interpretations of unnested atomic formulae, but in fact there is a natural way to associate any ρ-formula to a τ-formula.

Proposition. *Given an interpretation Γ of $\mathfrak{B}$ in $\mathfrak{A}$ there is a natural function*

$$(-)_\Gamma : \mathcal{L}(\rho) \longrightarrow \mathcal{L}(\tau)$$

such that $\mathfrak{A} \models (\varphi)_\Gamma(\bar{a})$ if and only if all a_i satisfy ∂_Γ and $\mathfrak{B} \models \varphi(\pi\bar{a})$. The association is given inductively by

- *for φ atomic unnested $(\varphi)_\Gamma$ is φ_Γ (as given in the definition of an interpretation)*
- $(\varphi \wedge \psi)_\Gamma$ *is* $(\varphi)_\Gamma \wedge (\Psi)_\Gamma$
- $(\neg\varphi)_\Gamma$ *will be* $\neg(\varphi)_\Gamma \wedge \bigwedge \partial_\Gamma(-)$
- $(\exists x\varphi)_\Gamma$ *will be* $(\exists y_0, \dots, y_{l-1})(\partial_\Gamma(\bar{y}) \wedge \varphi_\Gamma(\bar{y}))$

Proof. Immediate from the construction of Γ. □

Given a collection of formulae in $\mathcal{L}(\tau)$: ∂_Γ and φ_Γ for φ an unnested formula in $\mathcal{L}(\rho)$ then we want a theory T_Γ which says that these formulae give an interpretation. I.e. T_Γ asserts that for any $\mathfrak{A}$ which models T_Γ then the data ∂_Γ and φ_Γ define an interpretation. More precisely T_Γ must say

- If φ has n free variables and ∂_Γ has m free variables then φ_Γ has mn free variables.

- $(y_0 = y_1)_\Gamma$ is an equivalence relation $\sim$ on $\partial_\Gamma(-)$.
- for each $f \in \mathcal{F}_\rho$ if φ is $f(x) = y$ then T_Γ must say that

$$\forall \bar{u} \exists \bar{v} \varphi_\Gamma(\bar{u}, \bar{v}) \wedge \forall \bar{u}, \forall v, w(\varphi_\Gamma(\bar{u}, v) \wedge \varphi_\Gamma(\bar{u}, w) \longrightarrow v \sim w)$$

- For each constant $c \in \mathcal{C}_\rho$, if φ is $y = c$, then T_Γ must say that

$$\exists \bar{x} \varphi_\Gamma(\bar{x}) \wedge \forall \bar{x}, \bar{y} \varphi_\Gamma(\bar{x}) \wedge \varphi_\Gamma(\bar{y}) \longrightarrow x \sim y$$

  and

$$x \sim y \wedge \varphi_\Gamma(x) \longrightarrow \varphi_\Gamma(y)$$

- For $R \in \mathcal{R}_\rho$ then if $\varphi(x)$ is $R(x)$ we have that T_Γ must say that

$$\forall \bar{u}, \bar{v} \varphi_\Gamma(\bar{u}) \wedge \bar{u} \sim \bar{v} \longrightarrow \varphi_\Gamma(\bar{v}).$$

Proposition. *If $\mathfrak{A} \models T_\Gamma$ then Γ is an interpretation of $\Gamma(\mathfrak{A}) := \mathfrak{B}$ where* $\mathrm{dom}(\mathfrak{B}) := \partial_\Gamma(\mathfrak{A})/\sim$. *Here we have*

- *for constants c we have $c^{\mathfrak{B}} := [\bar{b}]_\sim$ for any $\bar{a} \in \partial_\Gamma(\mathfrak{A})$ such that $\mathfrak{A} \models (x{=}c)_\Gamma(\bar{a})$.*
- $([\bar{a}_0]_\sim, \ldots, [\bar{a}_{n-1}]_\sim) \in R^{\mathfrak{B}}$ *iff* $\mathfrak{A} \models (R(x))_\Gamma(\bar{a})$
- *and* $f^{\mathfrak{B}}([\bar{a}]_\sim) = [b]_\sim$ *iff* $\mathfrak{A} \models (f(\bar{x}) = y)_\Gamma(\bar{a}, b)$.

Proof. We defined T_Γ so as to say exactly what this proposition is saying. □

Example. If $\mathfrak{A}'$ is a definitional expansion on $\mathfrak{A}$ then the definitional expansion *is* an interpretation of $\mathfrak{A}'$ in $\mathfrak{A}$.

A useful observation (which we will now prove) is that an interpretation Γ preserves elementary substructures.

Proposition. *If $(\partial_\Gamma, \{\varphi_\Gamma : \varphi$ unnested ρ-formula$\})$ is given and $\mathfrak{A} \preccurlyeq \mathfrak{A}'$ where $\mathfrak{A}' \models T_\Gamma$ then $\Gamma(\mathfrak{A}) \preccurlyeq \Gamma(\mathfrak{A}')$.*

Proof. Since $\mathfrak{A} \preccurlyeq \mathfrak{A}'$ we have $\partial_\Gamma(\mathfrak{A}) \subseteq \partial_\Gamma(\mathfrak{A}')$. Also $\sim$ the equivalence relation (given by $(x = y)_\Gamma$) on ∂_Γ is an equivalence relation on both $\partial_\Gamma(\mathfrak{A})$ and on $\partial_\Gamma(\mathfrak{A}')$. Furthermore, again since $\mathfrak{A} \preccurlyeq A'$ the restriction of $\sim$ on $\partial_\Gamma(\mathfrak{A}')$ to $\partial_\Gamma(\mathfrak{A})$ is just the old $\sim$.

So the inclusion

$$\partial_\Gamma(\mathfrak{A}) \longrightarrow \partial_\Gamma(\mathfrak{A}')$$

induces an inclusion

$$\partial_\Gamma(\mathfrak{A})/\sim^{\mathfrak{A}} \longrightarrow \partial_\Gamma(\mathfrak{A}')/\sim^{\mathfrak{A}'}.$$

The rest of the proof now follows from the earlier proposition: For any unnested formula φ in $\mathcal{L}(\rho)$ and tuple $\bar{a}$ from $\partial_\Gamma(\mathfrak{A})$ we have

$$\mathfrak{A} \models (\varphi)_\Gamma(\bar{a}) \qquad \text{iff} \qquad \Gamma(\mathfrak{A}) \models \varphi([\bar{a}]_\sim)$$

by the proposition. But by elementary extension we have

$$\mathfrak{A} \models (\varphi)_\Gamma(\bar{a}) \qquad \text{iff} \qquad \mathfrak{A}' \models (\varphi)_\Gamma(\bar{a})$$

and so again by the proposition we have

$$\mathfrak{A}' \models (\varphi)_\Gamma(\bar{a}) \qquad \text{iff} \qquad \Gamma(\mathfrak{A}') \models \varphi([\bar{a}]_\sim)$$

so $\Gamma(\mathfrak{A}) \preccurlyeq \Gamma(\mathfrak{A}')$. □

If one can interpret a class of ρ structures in some other class of τ-structures, then one can pass elementary embedding from one class to the other.

Interpretations induce continuous homomorphisms between automorphism groups. To prove this we first need a general lemma about topological groups.

Lemma. *Let* G *and* H *be topological groups and* $\alpha : G \longrightarrow H$ *a homomorphism. Then* α *is continuous if and only if* α *is continuous at the identity.*

Proof. The forward direction is clear.

Suppose α is continuous at the identity $1_G \in G$. Let $g \in G$ and let $U \subseteq H$ be an open subset containing $\alpha(g)$. Then translating U by $\alpha(g)^{-1}$ we see that $1_H \in \alpha(g)^{-1}U$. Now $\alpha(g)^{-1}U$ is also open since translation is a homeomorphism $H \to H$. Now by assumption there is some V open in G such that $1_G \in V$ and $\alpha(V) \subseteq \alpha(g)^{-1}U$. Thus gV contains g (and is open) and $\alpha(gV) \subseteq U$. □

Proposition. *To an interpretation* Γ *of* $\mathfrak{B}$ *in* $\mathfrak{A}$ *there is an associated continuous homomorphism*

$$\Gamma : \mathrm{Aut}(\mathfrak{A}) \longrightarrow \mathrm{Aut}(\mathfrak{B})$$

Proof. We first define the homomorphism.

Let σ be an automorphism of $\mathfrak{A}$. First note that σ must preserve $\partial_\Gamma(\mathfrak{A})$. I.e. $\mathfrak{A} \models \partial_\Gamma(\bar{a})$ if and only if $\mathfrak{A} \models \partial_\Gamma(\sigma\bar{a})$.

Now the equivalence relation $\sim$ is also defined by some formula, so σ also preserves this. I.e. $a \sim b$ iff $\sigma a \sim \sigma b$.

Thus σ induces a function, $\hat{\sigma}$ of equivalence classes $\partial(\mathfrak{A})/\sim$. Now the by the isomorphism $(\partial(\mathfrak{A})/\sim) \cong \mathfrak{B}$ we get a (bijective) function $\Gamma(\sigma) : \mathfrak{B} \to \mathfrak{B}$.

We must check that it is also an automorphism. It suffices to check that $\Gamma(\sigma)$ preserves unnested ρ-formulae. Let φ be an unnested ρ-formula and $\mathfrak{B} \models \varphi(\bar{b})$. This is equivalent to $\mathfrak{A} \models (\varphi)_\Gamma(\bar{a})$ (where $\bar{a} = \pi(\bar{b})$) which is equivalent to $\mathfrak{A} \models (\varphi)\Gamma(\sigma\bar{a})$ and finally this is equivalent to $\mathfrak{B} \models \varphi(\Gamma(\sigma)(\bar{b})$.

Finally we must also check the continuity of $\Gamma : \mathrm{Aut}(\mathfrak{A}) \longrightarrow \mathrm{Aut}(\mathfrak{B})$. For this we use the lemma: It suffices to check continuity at the identity. Let U be open subset of $\mathrm{Aut}(\mathfrak{B})$ containing $\Gamma(id_\mathfrak{A})$. Without loss of generality we may assume that U is a basic open set around $\mathrm{id}_{\mathfrak{B}}$, i.e. take U to be the stabilizer of $\bar{b}$ for some $\bar{b}$ from $\mathfrak{B}$. Let $\bar{a}$ be a finite tuple of $\mathfrak{A}$ such that $\bar{b} = \pi\bar{a}$ (which is possible since π is surjective). Then $\Gamma(\sigma)(U_{\bar{a},\bar{a}}) \subseteq U_{\bar{b},\bar{b}}$. So Γ is continuous. $\square$

Question. Suppose that Γ is an interpretation of $\mathfrak{B}$ in $\mathfrak{A}$ and Δ is an interpretation of $\mathfrak{A}$ in $\mathfrak{B}$. Must $\Delta \circ \Gamma : \mathrm{Aut}(\mathfrak{A}) \longrightarrow \mathrm{Aut}(\mathfrak{A})$ be an automorphism?

These and many other related questions have been heavily studied, see for example [3] and [4].

Lecture 15

15.1 Elimination of Imaginaries

Definition. An **imaginary element** of a τ-structure $\mathfrak{A}$ is a class $[a]_E$ where $a \in A^n$ and E is a definable equivalence relation on A^n.

So an imaginary element can be thought of as an element of a quotient of a definable set by a definable equivalence relation. Thus thinking in terms of the category of definable sets[1], passing to the imaginaries means allowing this category to have quotients.

Example (Trivial equivalence relation). If $a \in A^n$ then we may identify a with the class $[a]_{=_{A^n}}$ under the (definable) equivalence relation given simply by equality.

Definition. A τ-structure $\mathfrak{A}$ **eliminates imaginaries** if, for every definable equivalence relation E on A^n there exists definable function $f : A^n \longrightarrow A^m$ such that for $x, y \in A^n$ we have

$$xEy \iff f(x) = f(y)$$

Remark. The definition given above is what Hodges calls *uniform elimination of imaginaries*.

Remark. If $\mathfrak{A}$ eliminates imaginaries, then for any definable set X and definable equivalence relation E on X, there is a definable set Y and a definable bijection $f : X/E \longrightarrow Y$. Of course this is not literally true, we should rather say that there is a definable map $f' : X \longrightarrow Y$ such that f' is invariant on the equivalence classes defined by E.

So elimination of imaginaries is saying that quotients exists in the category of definable sets.

Remark. If $\mathfrak{A}$ eliminates imaginaries then for any imaginary element $[a]_E = \tilde{a}$ there is some tuple $\hat{a} \in A^m$ such that $\tilde{a}$ and $\hat{a}$ are **interdefinable** i.e. there is a formula $\varphi(x, y)$ such that

[1] This is the category whose objects are definable sets and morphism are definable functions.

- $\mathfrak{A} \models \varphi(a, \hat{a})$,
- If $a'Ea$ then $\mathfrak{A} \models \varphi(a', \hat{a})$,
- If $\varphi(b, \hat{a})$ then bEa,
- If $\varphi(a, c)$ then $c = \hat{a}$.

To get the formula φ we use the function f given by the definition of elimination of imaginaries, and let $\varphi(x, y)$ be $f(x) = y$, (note: then $\hat{a} = f(a)$).

Almost conversely, if for every $\mathfrak{A}' \equiv \mathfrak{A}$ every imaginary in $\mathfrak{A}'$ is interdefinable with a **real** (i.e. non-imaginary) tuple then $\mathfrak{A}$ eliminates imaginaries. We will prove this after proving the compactness theorem.

Example. For any structure $\mathfrak{A}$, every imaginary in $\mathfrak{A}_A$ is interdefinable with a sequence of real elements.

Example. Let $\mathfrak{A} = (\mathbb{N}, <, \equiv \pmod 2)$. Then $\mathfrak{A}$ eliminates imaginaries. For example, to eliminate the "odd/even" equivalence relation, E, we can define $f : \mathbb{N} \longrightarrow \mathbb{N}$ by mapping x to the least z such that xEz. I.e. f is define by the formula

$$f(x) = y \iff xEy \wedge \forall z[xEz \longrightarrow y < z \vee y = z].$$

In the above example we claim furthermore $\mathfrak{A}$ eliminates all other equivalence relations. This is because $\mathfrak{A}$ has *definable choice functions*.

Definition. $\mathfrak{A}$ has **definable choice functions** if for any formula $\theta(\bar{x}, \bar{y})$ there is a definable function $f(\bar{y})$ such that

$$\forall \bar{y} \exists \bar{x} [\theta(\bar{x}, \bar{y}) \longleftrightarrow \theta(f(\bar{y}), \bar{y})]$$

(i.e. f is a *skolem function* for θ) and such that

$$\forall \bar{y} \forall \bar{z} [\forall \bar{x} (\theta(\bar{x}, \bar{y}) \longleftrightarrow \theta(\bar{x}, \bar{z})) \longrightarrow f(\bar{y}) = f(\bar{z})]$$

Proposition. *If $\mathfrak{A}$ has definable choice functions then $\mathfrak{A}$ eliminates imaginaries.*

Proof. Given a definable equivalence relation E on A^n let f be a definable choice function for $E(x, y)$. Since E is an equivalence relation we have $\forall y E(f(y), y)$ and

$$\forall y, z \; ([y]_E = [z]_E \longrightarrow f(y) = f(z))$$

thus $f(y) = f(z) \iff yEz$. $\square$

Example (continued). We now see that $\mathfrak{A} = (\mathbb{N}, <, \equiv \pmod 2)$ eliminates imaginaries. Basically since $\mathfrak{A}$ is well ordered, we can find a least element to witness membership of definable sets, hence we have definable choice functions.

Question. Suppose $\mathfrak{A}$ has Skolem functions. Must $\mathfrak{A}$ eliminate imaginaries?

Example. $\mathfrak{A} = (\mathbb{N}, \equiv \pmod 2)$ does *not* eliminate imaginaries.

First note that the only definable subsets of $\mathbb{N}$ are $\varnothing, \mathbb{N}, 2\mathbb{N}$ and $(2n+1)\mathbb{N}$. This is because $\mathfrak{A}$ has an automorphisms which switches $(2n+1)\mathbb{N}$ and $2\mathbb{N}$.

Now suppose $f : \mathbb{N} \longrightarrow \mathbb{N}^M$ eliminates the equivalence relation $\equiv \pmod 2$, i.e.

$$f(x) = f(y) \iff y \equiv x \pmod 2.$$

Then range(f) is definable and has cardinality 2. Since there are no definable subsets of $\mathbb{N}$ of cardinality 2, we must have $M > 1$. Now let $\pi : \mathbb{N}^M \longrightarrow \mathbb{N}$ be a projection. Then $\pi(\text{range}(f))$ is a finite nonempty definable subset of $\mathbb{N}$. But no such set exists.

Note that if we allow parameters in defining subsets, then $\mathfrak{A}$ does eliminate imaginaries.

Example. Consider a vector space V over a field K. We will put these together into a two-sorted structure $(V, K, +_V, 0_V, +_K, \cdot_K, \cdot_{K,V}, 0_K, 1_K)$ here the functions and constant are suitably defined. Now define, for $v, w \in V$,

$$v \sim w \iff \exists \lambda \in K \setminus 0 \; \lambda v = w$$

Then $V/\sim$ is the projective space $\mathbb{P}(V)$.

Question. Can we eliminate imaginaries in this case?

Proposition. *If the τ-structure $\mathfrak{A}$ eliminates imaginaries, then $\mathfrak{A}_A$ eliminates imaginaries.*

Proof. The idea is that an equivalence relation with parameters can be obtained as a fiber of an equivalence relation in more variables. More precisely, let $E \subseteq A^n$ be an equivalence relation definable in $\mathfrak{A}_A$. Let $\varphi(x, y; z) \in \mathcal{L}(\tau)$ and $a \in A^l$ be such that

$$xEy \iff \mathfrak{A} \models \varphi(x, y; a).$$

Now define

$$\psi(x, u, y, v) = \begin{cases} u = v \wedge \text{“}\varphi \text{ defines an equivalence relation“} & \text{or} \\ u \neq v & \text{or} \\ \text{“}\varphi(x, y, v) \text{ does not define an equivalence relation“} & \end{cases}$$

Where “φ defines an equivalence relation“ is clearly first-order expressible. Now ψ defines an equivalence relation on A^{n+l}. Letting $f : A^{n+l} \longrightarrow A^M$ eliminate ψ, then $f(-, a)$ eliminates E. ∎

15.2 Multi-Sorted Structures, $\mathfrak{A}^{eq}$

We saw that atomisation was a way to force elimination of quantifiers. Similarly one can force elimination of imaginaries, provided one is willing to work in a multi-sorted logic.

Given a τ-structure $\mathfrak{A}$ we will construct $\mathfrak{A}^{eq}$ as follows. For each definable equivalence relation E on A^n we have a sort S_E and a function symbol π_E interpreted as

$$\pi_E^{\mathfrak{A}^{eq}} : A^n \longrightarrow S_E^{\mathfrak{A}^{eq}} := A^n/E$$

mapping a to $[a]_E$. This shows $\mathfrak{A}^{eq}$ is interpreted[2] in $\mathfrak{A}$.

Conversely, $\mathfrak{A}$ can be interpreted in $\mathfrak{A}^{eq}$. Let $\partial_{()=}(x)$ be $x \in S_{=_A}$. Given an unnested τ-formula $\varphi(x_0, \dots, x_{l-1})$ consider E_φ defined by

$$\bar{x} E_\varphi \bar{y} \quad \Longleftrightarrow \quad (\varphi(\bar{x}) \longleftrightarrow \varphi(\bar{y}))$$

then we have

$$\pi_{E_\varphi} : A^l \longrightarrow S_{E_\varphi}$$

This almost works.

Question. How can we define, in $\mathfrak{A}^{eq}$, the class $[\bar{a}]_{E_\varphi}$ where $\mathfrak{A} \models \varphi(\bar{a})$?

[2]Here we mean interpreted in the sense of multi sorted structures.

Lecture 16

16.1 Compactness

We shall now prove the compactness theorem. The proof we give is by the so-called Henkin construction. Later in the course we will give a second proof making use of ultrafilters.

Definition. A theory T is said to be **finitely satisfiable** if whenever $T_0 \subseteq T$ is a *finite* subset of T there exists a model $\mathfrak{A}$ of T_0.

Theorem 16.1 (The Compactness Theorem)**.** *Let τ be a signature and T a τ-theory. If T is finitely satisfiable then T is satisfiable, i.e. then there exists a model $\mathfrak{A} \models T$.*

Remark. Compactness, as proved here, is a property of *first-order logic*, i.e. the sentences of T are assumed to come from the first-order language $\mathcal{L}_{\omega\omega}(\tau)$.

We will use the technique of using the language itself to build a structure satisfying T.

We first prove some lemma's allowing us to reduce the problem of finding models of T to that of finding models of a certain nice extension of T.

Definition. A τ-theory T **has Henkin constants** if for each formula $\varphi(x) \in \mathcal{L}(\tau)$ with one free variable x, there is a constant symbol $c \in \mathcal{C}_\tau$ such that

$$T \vdash \exists x\ \varphi(x) \longleftrightarrow \varphi(c).$$

Henkin constants act ass witnesses to all existential sentences, iff they are implied by T.

Proposition. *For any signature τ there is a signature τ^{Hen} expressible as $\bigcup_{i=0}^{\infty} \tau_{(i)}^{\mathrm{Hen}}$ and a theory T^{Hen} in $\mathcal{L}(\tau^{\mathrm{Hen}})$ such that*

1.) τ^{Hen} *is an expansion by constants of* τ.
2.) T^{Hen} *has Henkin constants.*

3.) *For each* $\mathfrak{A} \in \mathrm{Str}(\tau)$ *there is some non-zero* $\mathfrak{A}' \in \mathrm{Str}(\tau^{\mathrm{Hen}})$ *such that* $\mathfrak{A}' \models T^{\mathrm{Hen}}$ *and* $\mathfrak{A}$ *is the* τ*-reduct of* $\mathfrak{A}'$.

Proof. We define $\tau^{\mathrm{Hen}}_{(n)}$ and T^{Hen}_n recursively in n. Let $\tau^{\mathrm{Hen}}_{(0)} := \tau$ and $T^{\mathrm{Hen}}_0 := \varnothing$. At stage n we expand $\tau^{\mathrm{Hen}}_{(n)}$ by adding constants only, indeed let

$$\mathcal{C}_{\tau^{\mathrm{Hen}}_{(n+1)}} := \mathcal{C}_{\tau^{\mathrm{Hen}}_{(n)}} \cup \{c_\varphi : \varphi \text{ is in } \mathcal{L}(\tau^{\mathrm{Hen}}_{(n)}) \text{ with exactly 1 free variable}\}.$$

We also expand T^{Hen}_n to state that the new constants c_φ act as witnesses, i.e. let

$$T^{\mathrm{Hen}}_{n+1} := T^{\mathrm{Hen}}_n \cup \{\exists x \varphi(x) \longleftrightarrow \varphi(c_\varphi) : \varphi \in \mathcal{L}(\tau^{\mathrm{Hen}}_{(n)}) \text{ with exactly 1 free variable}\}.$$

Then we define

$$\tau^{\mathrm{Hen}} := \bigcup_n \tau^{\mathrm{Hen}}_{(n)} \qquad \text{and} \qquad T^{\mathrm{Hen}} := \bigcup_n T^{\mathrm{Hen}}_n.$$

Clearly T^{Hen} has Henkin constants and τ^{Hen} is an expansion of τ by constants. This takes care of *1.)* and *2.)* in the proposition.

Now we show that τ^{Hen} and T^{Hen} satisfy property *3.)*. Let $\mathfrak{A}$ be non-empty τ-structure [1]. We will find $\mathfrak{A}_{(n)} \in \mathrm{Str}(\tau^{\mathrm{Hen}}_{(n)})$ such that $\mathfrak{A}_{(0)} = \mathfrak{A}$ and $\mathfrak{A}_{(n)} = \mathfrak{A}_{(n+1)}|_{\tau^{\mathrm{Hen}}_{(n)}}$ and such that $\mathfrak{A}_{(n)} \models T^{\mathrm{Hen}}_n$.

For $n = 0$ let $\mathfrak{A}_{(0)} := \mathfrak{A}$, then have $\mathfrak{A}_{(0)} \models T^{\mathrm{Hen}}_0 = \varnothing$.

Given $\mathfrak{A}_{(n)}$, for each $\varphi \in \mathcal{L}(\tau^{\mathrm{Hen}}_{(n)})$ with free variable x if $\mathfrak{A}_{(n)} \models \exists x \varphi$ let $a_\varphi \in \varphi(\mathfrak{A}_{(n)})$, if $\mathfrak{A}_{(n)} \models \neg\exists x \varphi(x)$ then let a_φ be arbitrary. *Here we have used the axiom of choice to pick the witnesses* a_φ. We interpret

$$c_\varphi^{\mathfrak{A}_{(n+1)}} := a_\varphi.$$

This ensures that $\mathfrak{A}_{(n+1)} \models T^{\mathrm{Hen}}_{n+1}$.

Finally let $\mathfrak{A}'$ be the unique τ^{Hen} structure with $\mathfrak{A}'|_{\tau^{\mathrm{Hen}}_{(n)}} = \mathfrak{A}_{(n)}$. $\mathfrak{A}'$ is the desired structure. □

Corollary. *If* T *is a finitely satisfiable* τ*-theory then* $T \cup T^{\mathrm{Hen}}$ *is a finitely satisfiable* τ^{Hen}*-theory.*

Proof. Let $S' \subseteq T \cup T^{\mathrm{Hen}}$ be finite. Then $S := S' \cap T$ is finite. By hypothesis there is some $\mathfrak{A}$ such that $\mathfrak{A} \models S$. By the proposition there is some expansion $\mathfrak{A}' \in \mathrm{Str}(\tau^{\mathrm{Hen}})$ such that $\mathfrak{A}' \models T^{\mathrm{Hen}}$ and such that $\mathfrak{A}'|_\tau = \mathfrak{A}$. This then implies that $\mathfrak{A}' \models S$. But since $S' \subseteq S \cup T^{\mathrm{Hen}}$ we see that $\mathfrak{A}' \models S'$. □

[1] If $\mathfrak{A}$ is empty then condition *3.)* is vacuously satisfied since it is never the case that $\exists\varphi(x)$ is true.

Thus given any finitely satisfiable theory T we can canonically expand the language and the theory to get a finitely satisfiable theory T^{Hen} *which has Henkin constants*. So in proving the compactness theorem it will suffice to consider only the case where T has Henkin constants.

We shall make one more reduction of the problem before proving the compactness theorem. This time we show how to extend to complete theories. Recall that we say a theory T is **complete** if for every sentence φ either $\varphi \in T$ or $\neg\varphi \in T$.

Proposition. *If* T *is a finitely satiable* τ*-theory then there is a* complete *extension* $\widetilde{T} \supseteq T$ *which is still finitely satisfiable.*

Proof. We use Zorn's lemma to find a maximal finitely satisfiable extension of T and then argue by maximality that this extension must be complete.

Indeed let $\mathcal{P}$ be the partially ordered (by inclusion) set of extensions $T' \supseteq T$ which are finitely satisfiable. $\mathcal{P}$ is non-empty since $T \in \mathcal{P}$. Taking a non-empty chain in $\mathcal{P}$ then the union of the chain is also an element of $\mathcal{P}$ since any finite subset of the union is contain in one of the elements of the union and therefore satisfiable. Therefore by Zorn's lemma there is a maximal finitely satisfiable extension $\widetilde{T} \supseteq T$.

Now we claim that $\widetilde{T}$ is complete. Suppose by way of contradiction that φ is a $\mathcal{L}(\tau)$-sentence and such that both φ and $\neg\varphi$ are *not* in $\widetilde{T}$. Then both

$$\widetilde{T} \cup \{\varphi\} \supsetneq \widetilde{T} \supseteq T$$

and

$$\widetilde{T} \cup \{\neg\varphi\} \supsetneq \widetilde{T} \supseteq T$$

so neither $\widetilde{T} \cup \{\varphi\}$ nor $\widetilde{T} \cup \{\neg\varphi\}$ are elements of $\mathcal{P}$. Since they both contain T, the only way they can avoid being in $\mathcal{P}$ is if they are not finitely satisfiable. So there is some $U, V \subseteq \widetilde{T}$ such that $U \cup \{\varphi\}$ and $V \cup \{\neg\varphi\}$ are not satisfiable. But now $U \cup V \subseteq \widetilde{T}$ is finite hence satisfiable. Let $\mathfrak{A}$ be a model of $U \cup V$. Now either φ or $\neg\varphi$ holds in $\mathfrak{A}$, either way we have a contradiction since one of $U \cup \{\varphi\}$ and $V \cup \{\neg\varphi\}$ will be satisfied by $\mathfrak{A}$. ☐

So starting with any finitely satisfiable theory T in any signature τ we can expand the signature and the theory to get a $\tau' \supseteq \tau$ and $T' \supseteq T$ which has Henkin constants. We can now extend further to another τ'-theory $T'' \supseteq T'$ which is complete. Then T'' still has Henkin constants and is also complete. If T'' is satisfiable then we can take a reduction back to τ to see that T is also satisfiable.

Proposition. *If* T *is a finitely satisfiable theory with Henkin constants, then there exists a model* $\mathfrak{A}$ *of* T*. In fact we may take* $\mathrm{dom}(\mathfrak{A})$ *to be* $\{c^{\mathfrak{A}} : c \in \mathcal{C}_\tau\}$.

Proof. Let $\mathcal{C} := \mathcal{C}_\tau$ be the set of constant symbols. We define a relation $\sim$ on $\mathcal{C}$ by $c \sim d$ iff the sentence $c{=}d$ is in T. We will show that $\sim$ is an equivalence relation and then let $\mathrm{dom}(\mathfrak{A})$ be the set of equivalence classes.

First, to see that $\sim$ is an equivalence relation we check the three axioms. They all follow the same pattern so let us just show the reflexivity: Let $c \in \mathcal{C}$. Since T is complete either $c{=}c$ or $c{\neq}c$ is in T. But T is also finitely satisfiable and since there is no model satisfying $c{\neq}c$ we must have $c{=}c \in T$, and so $c \sim c$.

We now define a τ-structure $\mathfrak{A}$ with domain $\mathcal{C}/\sim$. For $c \in \mathcal{C}_\tau$ let $c^{\mathfrak{A}} := [c]_\sim$. For $f \in \mathcal{F}_\tau$ of arity n, then given $c_0, \dots, c_{n-1} \in \mathcal{C}_\tau$ then

$$f^{\mathfrak{A}}([c_0]_\sim, \dots, [c_{n-1}]_\sim) = [d]_\sim$$

if $T \vdash f(c_0, \dots, c_{n-1}) = d$ for $d \in \mathcal{C}_\tau$. For $R \in \mathcal{R}_\tau$ of arity n and $c_0, \dots, c_{n-1} \in \mathcal{C}_\tau$ then

$$([c_0]_\sim, \dots, [c_{n-1}]_\sim) \in R^{\mathfrak{A}} \qquad \text{iff} \qquad R(c_0, \dots, c_{n-1}) \in T.$$

Of course we must check that f is actually a function, and that it is well-defined. Likewise we must also show that $R^{\mathfrak{A}}$ is well-defined.

To see that f is a function consider the formula $\varphi(x)$ given by $f(c_0, \dots, c_{n-1}) = x$. Since T has Henkin constants we have

$$T \vdash \exists x \varphi(x) \longleftrightarrow \varphi(d)$$

for some $d \in \mathcal{C}_\tau$. Since T is complete and finitely satisfiable it must be the case that $T \vdash \exists x f(c_0, \dots, c_{n-1}) = x$. Thus $T \vdash \varphi(d)$ and so $f^{\mathfrak{A}}$ is defined. To show that $f^{\mathfrak{A}}$ and $R^{\mathfrak{A}}$ are *well*-defined uses the same style of arguments.

We now have a τ-structure $\mathfrak{A}$. Finally we show that $\mathfrak{A}$ is a model of T. We work by induction on the complexity of the sentence φ to show that $\mathfrak{A} \models \varphi$ if and only if $\varphi \in T$. Without loss of generality we may assume φ is unnested.

If φ is an unnested atomic sentence then by construction of $\mathfrak{A}$ we see that $\mathfrak{A} \models \varphi$ if and only if $\varphi \in T$.

If φ is $\theta \wedge \psi$, then $\mathfrak{A} \models \varphi$ iff $\mathfrak{A} \models \theta$ and $\mathfrak{A} \models \psi$ which by the inductive hypothesis happens iff $\theta \in T$ and $\psi \in T$. But $\theta \in T$ and $\psi \in T$ iff $\theta \wedge \psi \in T$ since otherwise we violate the assumption that T is finitely satisfiable and complete. A similar argument works for the other boolean combinations.

Now suppose φ is $\exists x \theta$. Then

$$\begin{aligned} \mathfrak{A} \models \varphi \qquad & \text{iff} \qquad \exists a \in A,\ \mathfrak{A} \models \theta(a) && (16.1) \\ & \text{iff} \qquad \exists c \in \mathcal{C}_\tau,\ \mathfrak{A} \models \theta(c) && (16.2) \\ & \text{iff} \qquad \exists c \in \mathcal{C},\ \theta(c) \in T && (16.3) \end{aligned}$$

So if $\mathfrak{A} \models \varphi$ then since T is complete and finitely satisfiable we must have $\varphi \in T$. Furthermore if $\varphi \in T$ then by the above biimplications and using that T has Henkin constants we see that $\mathfrak{A} \models \varphi$.

Thus, $\mathfrak{A}$ is a model of T. □

Exercise. Let $\tau = \{E\}$ where E is a binary relation symbol. Use the compactness theorem to show that $\mathfrak{A} \preccurlyeq \mathfrak{B}$ where $\mathfrak{A}$ and $\mathfrak{B}$ are τ-structures such that $E^{\mathfrak{A}}$ and $E^{\mathfrak{B}}$ are both equivalence relations. In $\mathfrak{A}$ there is exactly one equivalence class of size n for each $n \in \omega$, and $\mathfrak{B}$ extends $\mathfrak{A}$ by having one new infinite equivalence class.

16.2 Upward Löwenheim-Skolem

Corollary (Upward Löwenheim-Skolem)**.** *If $\mathfrak{A}$ is infinite and λ any infinite cardinal and $|\mathfrak{A}| \leqslant \lambda$, then there exists $\mathfrak{B}$ such that $\mathfrak{A} \preccurlyeq \mathfrak{B}$ and $|\mathfrak{B}| = \lambda$.*

Proof. Let $\tau' \supseteq \tau$ be an expansion by constants

$$\mathcal{C}_{\tau'} := \mathcal{C}_{\tau} \cup \{c_\alpha : \alpha < \lambda\}.$$

Consider the theory

$$T := \mathrm{Th}(\mathfrak{A}_A) \cup \{c_\alpha \neq c_\beta : \alpha < \beta\}.$$

We claim that T is finitely satisfiable: Let $S \subseteq T$ be finite. Then S mentions only finitely many c_α's say, $\{c_{\alpha_1}, \dots, c_{\alpha_m}\}$. Let $\mathfrak{A}'$ be the τ'_A-structure with $\mathfrak{A}'|_{\tau_A} = \mathfrak{A}_A$ and $c_{\alpha_i}^{\mathfrak{A}'} := a_i$ where we pick distinct elements $a_1, \dots, a_m$ from A (which is possible since A is infinite). Then $\mathfrak{A}' \models S$.

By compactness there exists $\mathfrak{B}'$ a model of T. By construction $\mathfrak{B}'|_{\tau_A} \equiv \mathfrak{A}_A$ so $\mathfrak{A} \preccurlyeq \mathfrak{B}'|_\tau$ (by the elementary diagram lemma). Now $\{c_\alpha^{\mathfrak{B}} : \alpha < \lambda\} \subseteq B$ and so $|\mathfrak{B}| \geqslant \lambda$. If $|\mathfrak{B}|$ is too big we can use the Downward Löwenheim-Skolem theorem to get the right size. □

Lecture 17

17.1 Compactness Applied

The compactness theorem proved last time has many consequences and will be used constantly from now on.

For instance if some formula $\varphi(x, y)$ always defines a *finite* subset (in all structures) then in fact there is some bound on the size of the definable subsets. More precisely we have the following proposition.

Proposition. *Let* T *be a* τ*-theory and* $\varphi(x, y) \in \mathcal{L}(\tau)$ *(where* x *and* y *are tuples). If for all* $\mathfrak{A} \models T$ *and* $a \in A$ *we have that* $\varphi(\mathfrak{A}, a)$ *is finite, then there is some* $k \in \omega$ *such that* $T \vdash \forall y \exists^{\leqslant k} x\ \varphi(x, y)$.

Proof. We expand the signature by adding constants $\{c_i : i \in \omega\} \cup \{d\}$. Call this new signature τ'. Consider the τ'-theory

$$S := T \cup \{\varphi(c_i, d) : i \in \omega\} \cup \{c_i \neq c_j : i < j\}.$$

Suppose, towards contradiction, that the proposition is false, i.e. there is no bound $k \in \omega$. Then we claim that S is finitely satisfiable. By compactness we then get a model $\mathfrak{B}'$ of S. Taking the reduct $\mathfrak{B}'|_\tau$ we will see that $\varphi(\mathfrak{B}, a)$ is infinite for some a.

So we must show that S is finitely satisfiable. Let $S_0 \subseteq S$ be finite. Then there is some $N \in \omega$ such that

$$S_0 \subseteq T \cup \{\varphi(c_i, d) : i \leqslant N\} \cup \{c_i \neq c_j : i < j < N\}.$$

By assumption there is no $k \in \omega$ that bounds $\varphi(\mathfrak{A}, a)$, so there is some $\mathfrak{A} \models T$ and a from A such that $|\varphi(\mathfrak{A}, a)| > N$. Let $b_0, \ldots, b_{N-1}$ be N distinct elements from $\varphi(\mathfrak{A}, a)$. Expand $\mathfrak{A}$ to $\mathfrak{A}'$ (a τ'-structure) by

$$d^{\mathfrak{A}'} := a$$

and

$$c_i^{\mathfrak{A}'} := \begin{cases} b_i & \text{if } i < N \\ a & \text{otherwise} \end{cases}$$

(the choice a for $c_i^{\mathfrak{A}'}$ where $i \geqslant N$ is arbitrary, we could choose which ever element we want). Now $\mathfrak{A}' \models S_0$ since $\mathfrak{A} = \mathfrak{A}'|_\tau \models T$ and $\mathfrak{A} \models \varphi(b_i, a)$ for $i < N$ and $\mathfrak{A} \models c_i \neq c_j$ for $i < j < N$. Thus S_0 is satisfiable.

By compactness there is a model $\mathfrak{B}'$ of S. Let $b_i := c_i^{\mathfrak{B}'}$ and $a := d^{\mathfrak{B}'}$. Let $\mathfrak{B} := \mathfrak{B}'|_\tau$. Then $\mathfrak{B} \models T$ and the infinitely many distinct b_i's are all in the set $\varphi(\mathfrak{B}, a)$ which is a contradiction. This completes the proof. □

This proposition demonstrates a weakness of first-order logic. First-order logic cannot tell the difference between "arbitrarily large but finite" and "infinite". If we want to say "finite" then we must say "finite and bounded by k" for some k. The contrapositive is also interesting, namely that if we have a first-order theory with arbitrarily large finite models, then there is an infinite model. Of course this makes essential use of the first-order setting.

Example. Can one deduce the existence of an infinite well-ordered set from the existence of arbitrarily large finite well-ordered sets?

The most immediate approach using the compactness theorem never gets off the ground since being a well-ordered set is not a first-order property! It is however a second-order property, but second-order logic is not compact.

Instead what we can do is look at all sentences satisfied by all of the well-orders $\mathfrak{A}_n := (\{0, 1, \ldots, n-1\}, <)$. I.e. let T be the theory

$$\mathrm{Th}(\{\mathfrak{A}_n : n \in \omega\}).$$

Question. What is T? In fact T is the theory of discrete linear order with first and last elements. Can you prove this?

Applying compactness to

$$T \cup \{c_i \neq c_j : i \neq j, i, j \in \omega\}$$

to obtain a model $\mathfrak{A}$ of T which is infinite.

The infinite model $\mathfrak{A}$ will not be a well-order, however it will contain an infinite well-order. Indeed $\mathfrak{A}$ will have a first element, say b (bottom) and last element, say t (top). t will have predecessors $P^n(t)$ and b will have successors $S^n(b)$ for all $n \in \omega$. Since $\mathfrak{A}$ is a linear order $S^n(b) \neq P^m(t)$ for any n, m, thus we get an infinite descending chain $t > P(t) > P^2(t) > \ldots$.

However the subset $b < S(b) < S^2(b) < \ldots$ *is* an infinite well-order.

We mentioned that being a well-order is not first-order expressible. This has not actually been proven yet.

Proposition. *If* $(X, <)$ *is an infinite linear ordered set then there exists* $(Y, <)$ *such that* $(X, <) \equiv (Y, <)$ *and such that* $(Y, <)$ *is* not *well-ordered.*

Proof. First extend the signature by constants: $\tau' = \{<\} \cup \{c_i : i \in \omega\}$. Let T be the τ'-theory

$$\mathrm{Th}((X,<)) \cup \{c_i > c_j : i < j\}$$

(note the "reversal" of the ordering of the c_i's). Now T is finitely satisfiable [proof: by finding an appropriate finite subset of X (which we assumed was infinite) which serves to give is a finite decreasing chain]. By compactness there is some infinite model Y' of T. Taking the reduct back to τ we have $(Y,<) \equiv (X,<)$ and Y has an infinite descending chain. $\square$

The compactness theorem is very strong. As an example of its usefulness consider Ax's Theorem (Problem 13. Sec. 5.1 of Hodges). A variant of this theorem is the following.

Theorem 17.1 (Ax). *If $f : \mathbb{C}^n \longrightarrow \mathbb{C}^n$ is given by polynomials and f has prime order, then f has a fixed point.*

17.2 Types

Definition. Given a τ-structure $\mathfrak{A}$, $n \in \omega$, $\bar{a} \in A^n$ and $B \subseteq A$ then the **type of** $\bar{a}$ **over** B, $\mathrm{tp}^{\mathfrak{A}}(\bar{a}/B)$ is

$$\mathrm{Th}(\mathfrak{A}_B, \bar{a})$$

thought of in the language $\mathcal{L}(\tau_{B,x_1,\dots,x_n})$ where the x_i's are constant symbols which must only be substituted with the a_i's.

Notation. If the structure $\mathfrak{A}$ is clear from context then we write $\mathrm{tp}(\bar{a}/B)$ instead of $\mathrm{tp}^{\mathfrak{A}}(\bar{a}/B)$. Also we sometimes omit the bar above a even if a is a tuple.

Informally the type of a over B is the set of all formulae (with parameters from B) which are true of a inside $\mathfrak{A}$. Concretely we have

$$\mathrm{tp}^{\mathfrak{A}}(\bar{a}/B) = \{\varphi(x_1, \dots, x_n; \bar{b}) : \mathfrak{A} \models \varphi(\bar{a}, \bar{b}) \text{ with } \bar{b} \text{ from } B \text{ and } \varphi \text{ from } \mathcal{L}(\tau)\}.$$

More generally,

Definition. An n**-type** over B (relative to $\mathfrak{A}$) is a complete finitely satisfiable theory in $\mathcal{L}(\tau_{B,x_1,\dots,x_n})$ extending $\mathrm{Th}(\mathfrak{A}_B)$.

Definition. Given an n-type p we say that $a \in A^n$ **realizes** p if $p = \mathrm{tp}(a/B)$. If there is such an element in A^n then we say that p is **realized** in $\mathfrak{A}$. If there is no such element then we say that $\mathfrak{A}$ **omits** p.

We can always find an elementary superstructure wherein a given type is realized:

Proposition. *If $p(x_1, \dots, x_n)$ is an n-type, then there is some $\mathfrak{A}'$ which is an elementary extension of $\mathfrak{A}$ and $a \in (A')^n$ such that $p = \mathrm{tp}^{\mathfrak{A}'}(a/B)$.*

Proof. We use compactness. Let

$$T := p \cup \mathrm{Th}(\mathfrak{A}_A)$$

(recall that $\mathrm{Th}(\mathfrak{A}_A) = \mathrm{eldiag}(\mathfrak{A})$ by definition). We claim that T is finitely satisfiable. Let $T_0 \subseteq T$ be finite. Let

$$T_0 \cap p = \{\varphi_1(\bar{x}), \dots, \varphi_l(\bar{x})\}$$

We will show that $(T_0 \cap p) \cup \mathrm{Th}(\mathfrak{A}_A)$ has a model by showing that there exists $\bar{a} \in A^n$ such that $(\mathfrak{A}_A, \bar{a}) \models (T_0 \cap p) \cup \mathrm{Th}(\mathfrak{A}_A)$. So we want $a_1, \dots, a_n \in A$ such that

$$\mathfrak{A}_A \models \bigwedge_{i=1}^{l} \varphi_i(a_1, \dots, a_n).$$

We can find this if and only if

$$\mathfrak{A}_B \models \exists \bar{x} \bigwedge_{i=1}^{l} \varphi_i(\bar{x})$$

(remember that the φ_i's only involve parameters from B) which is true if and only if

$$\exists \bar{x} \bigwedge \varphi_i(\bar{x}) \in \mathrm{Th}(\mathfrak{A}_B) \subseteq p$$

which is true since p is complete.

So by compactness there is some $\mathfrak{A}' \models T$. So $\mathfrak{A}' \models \mathrm{Th}(\mathfrak{A}_A)$ and so $\mathfrak{A} \preccurlyeq \mathfrak{A}'$. Furthermore

$$p = \mathrm{tp}(x_1^{\mathfrak{A}'}, \dots, x_n^{\mathfrak{A}'}/B)$$

since p is complete and $p \subseteq \mathrm{tp}(x_1^{\mathfrak{A}'}, \dots, x_n^{\mathfrak{A}'}/B)$, so p is realized in $\mathfrak{A}'$. □

Example. We give an example where a type is omitted. Let $\mathfrak{A} = (\mathbb{Q}, <)$ and let $B = \mathbb{Q}$. Let $C := \{q \in \mathbb{Q} : q < \sqrt{2}\}$ and $p(x)$ be the 1-type given by the complete extension of

$$\mathrm{Th}(\mathfrak{A}_\mathbb{Q}) \cup \{q < x | q \in C\} \cup \{x < q : q \in \mathbb{Q} \setminus C\}$$

Now p is finitely satisfiable since given any finite $p_0 \subseteq p$ we only mention finitely many

$$q_1 < \cdots < q_n$$

and of these q_i's some are in C and some are not. Letting q_m be the maximal q_i contained in C then q_{m+1} is not in C. By density of $\mathbb{Q}$ there is some element

r between q_m and q_{m-1} such that $\varphi(r)$ holds for all $\varphi \in p_0$. However p is not realized in $\mathfrak{A}$ since this would require $\sqrt{2} \in \mathbb{Q}$.

One way of realizing p in this case would be to let $\mathfrak{A}' = (\mathbb{Q} \cup \{\sqrt{2}\}, <)$ then $\mathfrak{A} \preccurlyeq \mathfrak{A}'$ and $\mathfrak{A}'$ realizes p.

It is worthwhile studying all types together as a topological space.

Definition. Given a τ-structure $\mathfrak{A}$ and $B \subseteq A$ and $n \in \omega$ the **Stone space** $S_n(B)$ (also denoted $S_X(B)$) is the set

$$\{p : p \text{ an } n\text{-type over } B \text{ relative to } \mathfrak{A}\}.$$

We topologies $S_n(B)$ by letting the basic open sets be

$$(\varphi) := \{p \in S_n(B) : \varphi \in p\}$$

for $\varphi \in \mathcal{L}(\tau_{B,x_1,\ldots,x_n})$.

Lecture 18

18.1 Type Spaces

Let us recall the definition of $S_n(B)$ from last time.

Definition. Given a τ-structure $\mathfrak{B}$, $A \subseteq \mathrm{dom}(\mathfrak{B})$ and $n \in \omega$ then the **Stone space of $\mathfrak{B}$ over** A is

$$S_n(A) := \{p : p \text{ is an } n\text{-type over } A\}$$

The basic open sets in $S_n(A)$ are of the form

$$(\varphi) := \{p \in S_n(A) : \varphi \in p\} \qquad \text{for } \varphi \in \mathcal{L}(\tau_{A,x_1,\dots,x_n})$$

Remark. The spaces $S_n(A)$ are also called **type spaces**.

A topological space X is **totally disconnected** if for all distinct elements a and b of X there exists an open partition U, V such that $a \in U$ and $b \in V$.

Proposition. *$S_n(A)$ is a totally disconnected compact space.*

Proof. For *totally disconnectedness*: Take $p \neq q$ from $S_n(A)$. Then there is some φ in the symmetric difference of p and q. Suppose $\varphi \in p \setminus q$. Then $p \in (\varphi)$ and $q \in (\neg\varphi)$ since types are complete. Now $S_n(A) = (\varphi) \cup (\neg\varphi)$ and $(\varphi) \cap (\neg\varphi) = \varnothing$. So $S_n(A)$ is totally disconnected.

Now for *compactness*: Suppose $\mathcal{U}$ is an open cover of $S_n(A)$. We may assume that $\mathcal{U}$ consists of basic open sets, i.e. there is some set Φ of formulae in $\mathcal{L}(\tau_{A,x_1,\dots,x_n})$ such that $\mathcal{U} = \{(\varphi) : \varphi \in \Phi\}$.

Suppose towards contradiction that there is no finite subcover of $\mathcal{U}$ exists. Consider the theory

$$T := \mathrm{Th}(\mathfrak{B}_A) \cup \{\neg\varphi(\bar{x}) : \varphi \in \Phi\}.$$

We claim that T is satisfiable. If not then by compactness there is some finite $\Phi_0 \subseteq \Phi$ such that

$$\mathrm{Th}(\mathfrak{B}_A) \cup \{\neg\varphi : \varphi \in \Phi_0\}$$

would be inconsistent. I.e.

$$\mathrm{Th}(\mathfrak{B}_A) \vdash \neg \bigwedge_{\varphi \in \Phi_0} \neg \varphi(\bar{x})$$

which implies

$$\mathrm{Th}(\mathfrak{B}_A) \vdash \bigvee_{\varphi \in \Phi_0} \varphi(\bar{x})$$

which is true if and only if

$$\mathfrak{B}_A \models \forall \bar{x} \bigvee_{\varphi \in \Phi_0} \varphi(\bar{x})$$

Then for any type p in $S_n(A)$ since $p \supseteq \mathrm{Th}(\mathfrak{B}_A)$ we must have

$$p \vdash \bigvee_{\varphi \in \Phi_0} \varphi(\bar{x}).$$

Now since p is complete it must satisfy one of the $\varphi(x)$, i.e.

$$p \vdash \varphi(\bar{x})$$

for some $\varphi(\bar{x}) \in \Phi_0$. Another way of saying this is that $p \in (\varphi)$. But then p models $\mathrm{Th}(\mathfrak{B}_A) \cup \{\neg\varphi : \varphi \in \Phi_0\}$ contrary to the assumption that this theory is inconsistent.

Now applying compactness to the theory T we get a model $(\mathfrak{C}, \bar{b}) \models T$. Letting $q = \mathrm{tp}(\bar{b}/A)$ then $q \supseteq T$, so $q \notin \bigcup_{\varphi \in \Phi}(\varphi)$, but this is a contradiction since

$$q \in S_n(A) \setminus \bigcup_{\varphi \in \Phi} (\varphi).$$

which was assumed empty. □

Remark. In particular $S_n(A)$ is Hausdorff. Also note that the basic open sets (φ) are clopen.

Remark. Another way of showing that $S_n(A)$ is compact would be to use Tychonoff's theorem (the product of compact spaces is compact). Then one would consider the map

$$S_n(A) \longrightarrow \prod_{\varphi \in \mathcal{L}(\tau_{A,\bar{x}})} \{0, 1\}$$

which sends $p \in S_n(A)$ to it's characteristic function,

$$\chi_p(\varphi) = \begin{cases} 0 & \text{if } \varphi \notin p \\ 1 & \text{if } \varphi \in p \end{cases}$$

Giving $\{0,1\}$ the discrete topology, and $\prod\{0,1\}$ the product topology we get, by Tychonoff's theorem that the product is compact. The map above is continuous, and injective. Furthermore the image is closed (hence compact) and the map is actually a homeomorphism onto it's image.

Notation (temporary). If the signature is ambiguous then we denote by $S_n^\tau(A)$ the space of n-types of $\mathfrak{B}$ over $A \subseteq \mathfrak{B}$ where $\mathfrak{B}$ is a τ-structure.

Consider $\tau \subseteq \tau'$, an extension of signatures $\mathfrak{B}'$ a τ'-structure and $A \subseteq B' := \mathrm{dom}(\mathfrak{B}')$. Then there is a restriction map

$$|_\tau : S_n^{\tau'}(A) \longrightarrow S_n^\tau(A)$$

sending $p \in S_n^{\tau'}(A)$ to $p|_\tau := p \cap \mathcal{L}(\tau_{A,\bar{x}})$.

Proposition. *The above map is continuous and surjective.*

Proof. surjective: Let $q \in S_n^\tau(A)$. We claim that $q \cup \mathrm{eldiag}(\mathfrak{B}')$ is consistent. If not then there is some finite set $\Xi \subseteq \mathrm{eldiag}(\mathfrak{B}')$ and finite $Q \subseteq q$ such that $\Xi \cup Q$ is inconsistent. We have

$$\Xi = \{\xi_1(b_1), \dots, \xi_m(b_m)\}, \qquad \xi_i \in \mathcal{L}(\tau'_A) \text{ and } b_i \text{ a tuple from } B$$

and

$$Q = \{\varphi_1(\bar{x}), \dots, \varphi_l(\bar{x})\}, \qquad \varphi_i \in \mathcal{L}(\tau_A).$$

By padding we can assume all the tuples b_i are the same. Further, by conjunction we may assume that $m = l = 1$ so

$$\Xi = \{\xi(b)\} \qquad \text{and} \qquad Q = \{\varphi(\bar{x})\}.$$

Since $\Xi \cup Q$ is inconsistent we have

$$\xi(b) \vdash \neg\varphi(\bar{x})$$

(where $\bar{x} = x_1, \dots, x_n$ are new constants, not appearing in $\mathcal{C}_{\tau'_B}$) so have

$$\xi(b) \vdash \forall\bar{x}\neg\varphi(\bar{x}).$$

But $\forall\bar{x}\neg\varphi(\bar{x})$ is a *sentence* in $\mathcal{L}(\tau_A)$. Now $\mathrm{Th}(\mathfrak{B}'_A|_\tau) \subseteq \mathrm{Th}(\mathfrak{B}'_B)$ and so $\mathfrak{B}'_B \models \xi(b)$ so $\mathfrak{B}'_B \models \forall\bar{x}\neg\varphi(\bar{x})$, which implies that $\forall\bar{x}\neg\varphi(\bar{x})$ is in $\mathrm{Th}(\mathfrak{B}'_A|_\tau)$. Since $q \in S_n^\tau(A)$ we have $q \supseteq \mathrm{Th}(\mathfrak{B}'_A|_\tau)$. So $\forall\bar{x}\neg\varphi(\bar{x}) \in q$ but $q \vdash \varphi(\bar{x})$ which is a contradiction. Thus, by compactness, $q \cup \mathrm{eldiag}(\mathfrak{B}')$ is consistent.

Let $(\mathfrak{B}'', \bar{b}) \models \mathrm{eldiag}(\mathfrak{B}') \cup q$. Set $p = \mathrm{tp}^{\tau'}(b/A)$. Then $p|_\tau = q$. So the map is surjective.

Continuity: Let $U \subseteq S_n^\tau(A)$ be a basic open set, say $U = (\varphi)^\tau$ for some $\varphi \in \mathcal{L}(\tau_{A,\bar{x}})$. Then $(-)|_\tau^{-1}(U) = (\varphi)^{\tau'}$, which is also basic open. Thus the map is continuous. $\square$

Corollary. *Given* $A \subseteq B \subseteq C$ *there is a continuous onto map* $S_n(B) \longrightarrow S_n(A)$. *S*

So restriction of parameters defines a surjective continuous map.

Example (Finite type spaces). In the language of equality (i.e. $\tau = \varnothing$) we have $|S_1(\varnothing)| = 1$, $|S_2(\varnothing)| = 2$ and $|S_3(\varnothing)| = 5$.

Example (Countable type spaces). Let $\tau = \{E\}$ where E is a binary relation symbol. Let the τ-structure $\mathfrak{A} = (A, E)$ be such that E is an equivalence relation with exactly one equivalence class of size n for each $n \in \omega$ and no other equivalence classes. Then we claim that $|S_1(\varnothing)| = \aleph_0$. We can find $\aleph_0$ distinct elements of $S_1(\varnothing)$ by considering formulae (with one free variable x) expressing the number of elements related to x. For instance we could let $\varphi_n(x)$ be $\exists^{=n} y\ (y \neq x) \wedge E(y, x)$. This shows that $|S_1(\varnothing)| \geqslant \aleph_0$. We will not prove the other inequality. One way to approach this would be to prove some form of quantifier-simplification and then check that there are $\leqslant \aleph_0$ types.

Example (Maximal type space). Let τ be the signature of ordered fields and let $\mathfrak{A} = (\mathbb{Q}, +, \cdot, <, 0, 1)$. Then we claim $|S_n(\varnothing)| = 2^{\aleph_0}$. This is because we can define all rational numbers, and using these we can define all cuts of $\mathbb{Q}$, these are all consistent, taking the completion of these we see that there are at least $2^{\aleph_0}$ distinct types. However there cannot be any more since the language is countable and each type is a subset of the language.

Remark. Could there be a countable language $\mathcal{L}$ where the number of types over the empty set lies strictly between $\aleph_0$ and $2^{\aleph_0}$? The negative answer is known as *Vaught's Conjecture*.

Given a signaure τ and a set $\Delta(\bar{x}; \bar{y})$ of τ-formulae where $\bar{x}$ is a finite tuple of new variables and $\bar{y}$ is arbitrary. Given a τ-structure $\mathfrak{A}$, $b \in A^n$, $C \subseteq A$ we define a Δ-type, $\mathrm{tp}^{\Delta}(b/C)$ as

$$\begin{aligned} \mathrm{tp}^{\Delta}(b/C) := \quad & \{\delta(x, d) : \delta(x, y) \in \Delta(x, y), d \text{ from } C, \text{ and } \mathfrak{A} \models \delta(b, d)\} \\ & \cup \{\neg\delta(x, d) : \delta(x, y) \in \delta(x, y)\Delta(x, y), d \text{ from } C, \text{ and } \mathfrak{A} \models \neg\delta(b, d)\} \end{aligned}$$

Then we let the set of Δ-types, $S_n^{\Delta}(C)$ be the set of all maximal consistent sets of formulae of the form $\delta(x, d) \vee \neg\delta(x, d)$ as δ ranges through Δ and d ranges through C.

As before there is a natural restriction map

$$|_{\Delta} : S_n(C) \longrightarrow S_n^{\Delta}(C)$$

which is continuous and surjective.

Remark (Concerning stability theory). We have defined Δ-types in the way that Shelah defines them. This is a more syntactic way. There are however some semantic properties that one would expect them to have which they do not have. There is a subtle fix that can be found in Pillay's book.

Proposition. *Let $\mathfrak{A}$ be a τ-structure. Let $\Delta_n(x_1, \dots, x_n, \bar{y})$ be a set of formulae. Suppose that for all n the restriction map*

$$|_\Delta : S_n(\varnothing) \longrightarrow S_n^\Delta(\varnothing)$$

is a bijection. Then $\Delta = \bigcup_n \Delta_n$ is an elimination set for $\mathfrak{A}$.

Proof. Let $\varphi \in \mathcal{L}(\tau_{x_1,\dots,x_n})$ be a τ-formula. We must show that φ is equivalent to a boolean combination of elements of Δ. Consider the theory

$$T := \mathrm{Th}_\tau(\mathfrak{A}) \ \cup \ \{\varphi(a) \wedge \neg\varphi(b)\} \ \cup \ \{\delta(a) \longleftrightarrow \delta(b) : \delta \in \Delta_n(\bar{x})\}.$$

where a and b are new n-tuples of constant symbols.

T must be inconsistent. For suppose $(\mathfrak{A}', a, b) \models T$. Then $\mathrm{tp}(a) \neq \mathrm{tp}(b)$ while $\mathrm{tp}^\Delta(a) = \mathrm{tp}^\Delta(b)$. By hypothesis this cannot happen since $S_n(\varnothing)$ is in bijection with $S_n^\Delta(\varnothing)$.

So (by compactness) there is some finite part, of T that is inconsistent. I.e. there are some $\delta_1, \dots, \delta_l \in \Delta$ such that

$$T^* := \mathrm{Th}(\mathfrak{A}) \ \cup \ \{\varphi(a) \wedge \neg\varphi(b)\} \ \cup \ \{\delta_i(a) \longleftrightarrow \delta_i(b) : i \leqslant l\}$$

is inconsistent.

Notation. Recall some notation previously used: $\theta^1 := \theta$ and $\theta^{-1} := \neg\theta$.

For $s \in \{-1, 1\}^l$ let,

$$\Phi_s := \bigwedge_{i=1}^{l} \delta_i^{s(i)}$$

and

$$\Psi := \bigvee_{\{s \ : \ \mathfrak{A} \models \ \exists x \varphi(x) \wedge \Phi_s(x)\}} \Phi_s$$

Note that Ψ is a boolean combination of elements of Δ. Now we claim that Ψ is equivalent to φ. Suppose a from $\mathfrak{A}$ satisfies Ψ, i.e. $\mathfrak{A} \models \Psi(a)$. So for some $s \in \{-1, 1\}^l$ we have

$$\mathfrak{A} \models \Phi_s(a)$$

and by definition of Ψ we have, for the same s that

$$\mathfrak{A} \models \exists x \varphi(x) \wedge \Phi_s$$

Let b be from $\mathfrak{A}$ such that

$$\mathfrak{A} \models \varphi(b) \wedge \Phi_s(b).$$

So we have $\mathfrak{A} \models \Phi_s(a) \wedge \Phi_s(b)$, thus for all $i \leqslant l$ we get

$$\mathfrak{A} \models \delta_i(a) \longleftrightarrow \delta_i(b)$$

Now, since T^* is inconsistent, it follows that

$$\mathfrak{A} \models \varphi(a) \longleftrightarrow \varphi(b)$$

and since $\mathfrak{A} \models \varphi(b)$ we finally have $\mathfrak{A} \models \varphi(a)$. $\square$

This is a powerful technique for proving quantifier elimination. If you can show – by automorphism arguments or some sort of semantic analysis – that some set Δ of formulae is enough to distinguish all types, then it is also enough to distinguish all formulae. This is the way that one proves quantifier elimination for more complicated structures.

Lecture 19

19.1 Amalgamation

Amalgamations are useful for realizing many types all at once inside one structure.

We will prove the Elementary Amalgamation Theorem as a consequence of the Compactness Theorem. Let τ be some signature and $\mathfrak{A}, \mathfrak{B} \in \mathrm{Str}(\tau)$ and $C \subseteq A$ and $C' \subseteq B$ together with a bijection $f : C \to C'$.

Notation. By $\mathfrak{A}_C \equiv \mathfrak{B}_{C'}$ we mean the usual expect that whenever some constant symbol $c \in C$ is used in a formula on the $\mathfrak{A}$-side then the corresponding constant symbol $f(c) \in C'$ is used on the $\mathfrak{B}$-side.

Theorem 19.1 (Elementary Amalgamation)**.** *With τ, $\mathfrak{A}$ and $\mathfrak{B}$, C and f as above, if $\mathfrak{A}_C \equiv \mathfrak{B}_{C'}$ then there exists a τ-structure $\mathfrak{D}$ such that*

- *There are elementary embeddings $\iota : \mathfrak{A} \to \mathfrak{D}$ and $j : \mathfrak{B} \to \mathfrak{D}$, such that $\iota|_C = j \circ f$*

I.e. that the following diagram commutes:

$$\begin{tikzcd}
& \mathfrak{D} & \\
\mathfrak{A} \arrow[ur, hook, dashed, "\iota"] & & \mathfrak{B} \arrow[ul, hook', dashed, "j"'] \\
& C \arrow[ul, hook] \arrow[ur, hook'] &
\end{tikzcd}$$

Proof. First, without loss of generality we may assume that $A \cap B = C$, as sets, by identifying C and C' via f and then by replacing A and B with new disjoint copies whose only overlap is C.

We aim to use the elementary diagram lemma. Consider the $\mathcal{L}(\tau_{A \cup B})$-theory

$$T := \mathrm{eldiag}(\mathfrak{A}) \cup \mathrm{eldiag}(\mathfrak{B}).$$

A model of T would then (upon reduction to τ) give us what we want.

Suppose T does not have a model. By compactness there is some finite inconsistent subset of T. I.e. we would have

$$\psi_1(a^{(1)}), \ldots, \psi_n(a^{(n)}) \in \mathrm{eldiag}(\mathfrak{A}) \qquad \text{and} \qquad \varphi_1(b^{(1)}), \ldots, \varphi_m(b^{(m)}) \in \mathrm{eldiag}(\mathfrak{B})$$

where $\psi_i, \varphi_i \in \mathcal{L}(\tau_C)$ and where $a^{(i)}, b^{(i)}$ are tuples from A and B respectively, such that

$$\{\psi_1(a^{(1)}), \ldots, \psi_n(a^{(n)}), \varphi_1(b^{(1)}), \ldots, \varphi_m(b^{(m)})\}$$

is inconsistent.

We now make some adjustments to make things more maneageble.

- We may assume no constant symbols from C appear as coordinates of $a^{(i)}$ or $b^{(i)}$ since if they did then we could absorb them into the formulae ψ_i or φ_i from $\mathcal{L}(\tau_C)$.
- We may assume that $a^{(i)} = a^{(j)} =: a$ and $b^{(i)} = b^{(j)} =: b$ for all i, j, by padding the ψ_i's and φ_i's with dummy variables.
- We may assume that $n = m = 1$, i.e. that there is only one φ_i and ψ_i. This is because the elementary diagram is closed under conjunctions, so letting

  $$\varphi := \bigwedge \varphi_i \qquad \text{and} \qquad \psi := \bigwedge \psi_i$$

  amounts to the same thing.

So after these reductions we now have that the $\mathcal{L}(\tau_C)$-theory

$$\{\psi(a), \varphi(b)\}$$

is inconsistent (note that $\psi(a) \in \mathrm{eldiag}(\mathfrak{A})$ and $\varphi(b) \in \mathrm{eldiag}(\mathfrak{B})$). Thus we have

$$\models \psi(a) \longrightarrow \neg\varphi(b).$$

Now $\psi(a) \in \mathrm{eldiag}(\mathfrak{A})$ so $\mathfrak{A}_A \models \psi(a)$. Then for any choice of expansion of $\mathfrak{A}_A$ to $\tau_{A \cup \{b'\}}$ we must have

$$\mathfrak{A}_{A,b'} \models \neg\varphi(b').$$

(Here we have used that A and B are disjoint apart from C.) Since this holds for *any* way we interpret b' in $\mathfrak{A}$, it follows that

$$\mathfrak{A}_A \models \forall y \neg\varphi(y).$$

Now $\forall y \neg\varphi(y)$ is a τ_C-sentence. So $\forall y \neg\varphi(y) \in \mathrm{Th}(\mathfrak{A}_C)$. But by assumption $\mathfrak{A}_C \equiv \mathfrak{B}_C$ and so $\mathrm{Th}(\mathfrak{A}_C) = \mathrm{Th}(B_C)$, so that

$$\mathfrak{B} \models \forall y \neg\varphi(y).$$

But $\mathfrak{B}_B \models \varphi(b)$ and so $\mathfrak{B} \models \exists y\, \varphi(y)$, which is a contradiction.

Therefore T is consistent. Letting $\mathfrak{D}^+$ be a model of T we get the desired $\mathfrak{D}$ as $\mathfrak{D}^+|_\tau$. $\square$

Thus, given a common subset (or even substructure) and two extensions which from their first-order theories look the same relative to the common subset, then they can be amalgamated into a common elementary extension.

Example. Let $\tau = \tau_{abeliangroup}$, $\mathfrak{A} = (\mathbb{R}, +, -, 0)$ and $\mathfrak{B} = (\mathbb{Q}, +, -, 0)$ and $C = \{1\}$. Then we are in the case of the theorem, i.e. $\mathfrak{A}_C \equiv \mathfrak{B}_C$. So we can amalgamate $\mathfrak{A}$ and $\mathfrak{B}$. Now the result of this would be $\mathfrak{B}$. But then $A \cap B = \mathbb{Q}$ would strictly contain C. In fact there is no way of avoiding this.

This example shows that there can be some obstruction which prevents us from amalgamating freely over C, i.e. such that images of A and B inside D have too big an overlap. Another example may illuminate the problem.

Example. Let $\tau = \tau_{fields}$, $\mathfrak{A} = (\mathbb{Q}(t)^{alg}, +, \cdot, 0, 1)$ and $\mathfrak{B} = (\mathbb{C}, +, \cdot, 0, 1)$ and $C = \mathbb{Q}$. Once again in the amalgamation $\mathfrak{D}$, of $\mathfrak{A}$ and $\mathfrak{B}$ over C, the sets A and B (inside D) will strictly contain C. For instance $\sqrt[3]{5}$ will have to be in this intersection.

Definition. Let $\mathfrak{A}$ be a τ-structure and $X \subseteq \operatorname{dom}(\mathfrak{A})$. An element $a \in \operatorname{dom}(\mathfrak{A})$ is **algebraic over** X if there is a first-order formula $\varphi(x, \bar{y}) \in \mathcal{L}(\tau)$ and a tuple $\bar{b}$ from X such that $\mathfrak{A} \models \varphi(a, \bar{b}) \wedge \exists^{\leqslant n} x \varphi(x, \bar{b})$ for some $n \in \omega$. We write $\operatorname{acl}_{\mathfrak{A}}(X)$ for the set of elements of $\mathfrak{A}$ that are algebraic over X. We say that X is **algebraically closed** if $\operatorname{acl}_{\mathfrak{A}}(X) = X$. Another way of stating this is that whenever φ is a formula with parameters from X and $\varphi(\mathfrak{A})$ is finite, then $\varphi(\mathfrak{A}) \subseteq X$.

Theorem 19.2. *Let $\mathfrak{A}$, $\mathfrak{B}$ and C be as in the previous theorem (where we have identified C and C'). Assume further that C is algebraically closed as a subset of $\mathfrak{A}$. Then there is a τ-structure $\mathfrak{D}$ such that*

- *There are elementary embeddings $\iota : \mathfrak{A} \to \mathfrak{D}$ and $j : \mathfrak{B} \to \mathfrak{D}$, such that $\iota|_C = j \circ f$*
- $\iota(A) \cap j(B) = \iota(C)$.

I.e. we have the following commutative diagram

$$\begin{tikzcd} & \mathfrak{D} & \\ \mathfrak{A} \arrow[ur, hook, dashed, "\iota"] & & \mathfrak{B} \arrow[ul, hook', dashed, "j"'] \\ & C \arrow[ul, hook] \arrow[ur, hook'] & \end{tikzcd} \qquad \textit{where} \qquad \iota(A) \cap j(B) = \iota(C).$$

Proof. We follow much the same proof as the Elementary Amalgamation Theorem above. Therefore this will only be a sketch, to show the main differences between the two proofs.

As before we assume $A \cap B = C$.

Now let

$$T := \mathrm{eldiag}(\mathfrak{A}) \cup \mathrm{eldiag}(\mathfrak{B}) \cup \{a \neq b : a \in A \setminus C \text{ and } b \in B \setminus C\}.$$

We aim to show that T is consistent, this will clearly suffice for the theorem. As before, if T is not consistent then we get $\psi(\bar{a})$ and $\varphi(\bar{b})$ such that $\psi, \varphi \in \mathcal{L}(\tau_C)$ and $\psi(\bar{a}) \in \mathrm{eldiag}(\mathfrak{A})$ and $\varphi(\bar{b}) \in \mathrm{eldiag}(\mathfrak{B})$, *where all coordinates of* $\bar{a}$ *and* $\bar{b}$ *are not from* C. Furthermore we now also know that

$$\{\psi(\bar{a})\} \cup \{\varphi(\bar{b})\} \cup \{\bigwedge_{i,j \leqslant n} a_i \neq b_j\}$$

(for some $n \in \omega$) is inconsistent. I.e. we have

$$\models \psi(\bar{a}) \longrightarrow \left(\neg\varphi(\bar{b}) \ \vee \ \bigvee_{i,j \leqslant n} a_i = b_j\right).$$

Since the elements of the tuple $\bar{b}$ are not from C any expansion of $\mathfrak{A}_A$ to $\tau_{A \cup \{\bar{b}'\}}$ must have

$$\mathfrak{A}_{A,\bar{b}'} \models \neg\varphi(\bar{b}') \ \vee \ \bigvee_{i,j \leqslant n} a_i = b_j'.$$

Thus by definition of the universal quantifier,

$$\mathfrak{A}_A \models \forall y_1, \ldots, y_n \ \left(\neg\varphi(\bar{y}) \ \vee \ \bigvee_{i,j \leqslant n} a_i = y_j\right).$$

This implies that

$$\mathfrak{A} \models \exists x_1, \ldots, x_n \forall y_1, \ldots, y_n \left(\neg\varphi(\bar{y}) \ \vee \ \bigvee_{i,j \leqslant n} x_i = y_j\right)$$

But now $\exists \bar{x} \forall \bar{y} \left(\neg\varphi(\bar{y}) \ \vee \ \bigvee_{i,j \leqslant n} x_i = y_i\right)$ is a sentence in $\mathcal{L}(\tau_C)$. Since $\mathfrak{A}$ satisfies this sentence, it is an element in $\mathrm{Th}(\mathfrak{A}_C)$ which by assumption is equal to $\mathrm{Th}(\mathfrak{B}_C)$.

But, $\mathfrak{B}_B \models \varphi(\bar{b})$. Now we claim that for each $j \leqslant n$ the set

$$\{b' \mid \exists y_1, \ldots, y_{j-1}, y_{j+1}, \ldots, y_n \ \varphi(y_1, \ldots, y_{j-1}, b', y_{j+1}, \ldots, y_n)\}$$

is infinite: Since if not then, letting $\theta_j(\bar{x})$ be $\varphi(x_1, \ldots, x_{j-1}, b', x_{j+1}, \ldots, x_n)$, we would have that $\theta_j(\mathfrak{B})$ is finite. Now $b_j \in \theta_j(\mathfrak{B})$ and $\theta_j \in \mathcal{L}(\tau_C)$ so, *since* C *is* *algebraically closed*, $\theta_j(\mathfrak{B}) \subseteq C$ so $b_j \in C$. But this contradicts the assumption that all coordinates of $\bar{b}$ were not from C.

So with these infinitely many b''s we see that

$$\mathfrak{B}_C \models \forall \bar{x} \exists \bar{y}\ \varphi(\bar{y})\ \wedge\ \bigwedge y_i \neq x_j.$$

But $\mathfrak{A}_C \equiv \mathfrak{B}_C$ and we already saw that $\mathfrak{A}_C$ does not satisfy the above sentence. Thus we have a contradiction, and so T must be consistent. □

We can also amalgamate many models at the same time.

Corollary. *If* $\{\mathfrak{B}_i\}_{i \in I}$ *is a nonempty set of* τ*-structures with* $C \subseteq \mathfrak{B}_i$ *a common subset, and* $(\mathfrak{B}_i)_C \equiv (\mathfrak{B}_j)_C$ *for all* $i, j \in I$*, then there exists a* τ*-structure* $\mathfrak{D}$ *which is an elementary extension of all the* $\mathfrak{B}_i$*'s. If furthermore* C *is algebraically closed, then as before we can arrange that* $B_i \cap B_j = C$ *(inside of* D*) for all* $i \neq j$ *from* I*.*

Proof. We do the proof without assuming C is algebraically closed. The modifications in the case where C is algebraically closed are much the same as before.

We use compactness together with induction. Let

$$T := \bigcup_{i \in I} \mathrm{eldiag}(\mathfrak{B}_i).$$

We must check that T is satisfiable. By compactness T is consistent if and only if, for each finite $J \subseteq I$

$$\bigcup_{i \in J} \mathrm{eldiag}(\mathfrak{B}_i)$$

is consistent. This we can check by induction on $|J|$. If $|J| = 1$ then this is clear. For $|J| = n+1$ let $J = J' \cup \{j\}$. By induction there is some $\mathfrak{D}_{J'}$ such that for all $j' \in J'$ we have $\mathfrak{B}_{j'} \preccurlyeq \mathfrak{D}_{J'}$. Now using the Elementary Amalgamation Theorem we can amalgamate $\mathfrak{D}_{J'}$ together with $\mathfrak{B}_j$ over C. □

We can use these results to realize as many types as we want.

Corollary. *Given any* τ*-structure* $\mathfrak{C}$ *there exists some elementary extension* $\mathfrak{D}$ *of* $\mathfrak{C}$ *such that for all* $p \in S_1(C)$*,* p *is realized in* $\mathfrak{D}$*.*

Proof. For $p \in S_1(C)$ we have seen that we can realize it in some extension, say $\mathfrak{B}_p$, where $\mathfrak{C} \preccurlyeq \mathfrak{B}_p$. Using the above corollary with the family $\{\mathfrak{B}_p : p \in S_1(C)\}$ we get the existence of $\mathfrak{D}$ such that for all $p \in S_1(C)$ we have $\mathfrak{C} \preccurlyeq \mathfrak{B}_p \preccurlyeq \mathfrak{D}$. Thus every $p \in S_1(C)$ is realized in $\mathfrak{D}$. □

Remark. In the above corollary we could have taken some subset of $S_1(C)$ and realized all types from this subset. The proof is the same.

This will later be used to build *saturated* models, where every type over every "reasonably small" subset of the model, is realized in the model.

Lecture 20

20.1 Heirs and Coheirs

Notation. Recall that $S_n(A)$ denotes the space of n-types over A. The union $\bigcup_{n=1}^{\infty} S_n(A)$ is written simply as $S(A)$.

Definition. Let $\mathcal{M}$ be a τ-structure and $A \subseteq B \subseteq \mathrm{dom}(\mathcal{M})$. Given $p \in S(A)$ and $q \in S(B)$ with $p \subseteq q$ (as sets), then q is an **heir** of p if, for each formula $\varphi(x, y) \in \mathcal{L}(\tau_A)$, if there exists b from B such that $\varphi(x, b) \in q$, then there exists some b' from A such that $\varphi(x, b') \in p$.

Roughly, q is an heir of p if every formula represented in q is already represented in p.

A related but different notion is that of coheir.

Definition. Given p and q as above, we say that q is a **coheir** of p if, for each formula $\theta(x)$ from q there exists $a \in A$ such that $\mathcal{M} \models \theta(a)$.

Remark (equivalent definition of coheir). q is a coheir of p if, for each formula $\varphi(x, y) \in \mathcal{L}(\tau_A)$ if $b \in B$ and $\varphi(x, b) \in q$ then there is some $a \in A$ such that $\mathcal{M} \models \varphi(a, b)$.

Notation. Given a set A and a tuple b we denote by Ab the set $A \cup \{b_i : b_i \text{ from } b\}$.

Remark (In what sense is the notion of *coheir* "co" to the notion of *heir*?). Given $A \subseteq \mathrm{dom}(\mathcal{M})$ and a, b tuples, then

$$\mathrm{tp}(a/Ab) \text{ is an heir of } \mathrm{tp}(a/A)$$

if and only if

$$\mathrm{tp}(b/Aa) \text{ is a coheir of } \mathrm{tp}(b/A).$$

Proof.

- To say that $\mathrm{tp}(a/Ab)$ is an heir of $\mathrm{tp}(a/A)$ is to say that for each $\varphi(x, y) \in \mathcal{L}(\tau_A)$ with $\varphi(x, b) \in \mathrm{tp}(a/Ab)$ there exists b' from A such that $\varphi(x, b') \in \mathrm{tp}(a/A)$,

- which is to say that for each $\varphi(x, y) \in \mathcal{L}(\tau_A)$ we have, $\mathcal{M} \models \varphi(a, b)$ if and only if there is some b' from A such that $\mathcal{M} \models \varphi(a, b')$,
- which is to say that, for each $\varphi(x, y) \in \mathcal{L}(\tau_A)$ then $\varphi(a, y) \in \mathrm{tp}(b/Ab)$ if and only if there exists b' from A such that $\mathcal{M} \models \varphi(a, b')$,
- which is to say that $\mathrm{tp}(b/Ab)$ is a coheir of $\mathrm{tp}(b/A)$.

□

We will now show that heirs and coheirs always exist (if A is the domain of a model).

Proposition. *If* $\mathcal{M} \preccurlyeq \mathcal{N}$ *and* $B \supseteq M$ *and* $p \in S(M)$ *then there exist* $q, r \in S(B)$ *such that* q *is an heir of* p *and* r *is a coheir of* p.

Proof. We shall write down "what we don't want". To get q we want an extension of p such that no formulae are represented which are not already represented by p. Thus consider the theory

$$Q := p \cup \{\neg\theta(x, b) : \theta(x, y) \in \mathcal{L}(\tau_M) \text{ not represented in } p, \text{ and with } b \text{ from } B\}$$

We claim: Q *is consistent.*

Proof. (of claim) If not then there exists a finite list $\theta_1(x, b_1), \ldots, \theta_n(x, b_n)$ of formulae where $\theta_i(x, y)$ is not represented in p, and some formula $\varphi(x) \in p$ such that

$$\vdash \varphi(x) \longrightarrow \bigvee_i \theta_i(x, b_i).$$

In particular, by generalization,

$$\mathcal{N}_B \models \exists y_1, \ldots, y_n \forall x \left(\varphi(x) \longrightarrow \bigvee_i \theta_i(x, y_i) \right).$$

But now $\exists y_1, \ldots, y_n \forall x \, (\varphi(x) \longrightarrow \bigvee_i \theta_i(x, y_i))$ is a τ_A sentence, so since $\mathcal{M} \preccurlyeq \mathcal{N}$ we have that $\mathcal{M}_B$ also satisfies this sentence. Thus there exists $b'_1, \ldots, b'_n$ from M such that

$$\mathcal{M}_B \models \forall x \left(\varphi(x) \longrightarrow \bigvee_i \theta_i(x, y_i) \right).$$

Since each $\theta_i(x, b'_i)$ is a τ_M-formula, either $\theta_i(x, b'_i) \in p$ or $\neg\theta_i(x, b'_i) \in p$ (since p is complete). Now since $\varphi \in p$ there must indeed be some i such that $\theta_i(x, b'_i) \in p$. But this means that $\theta_i(x, y)$ is represented in p, which is a contradiction. □

Now by compactness Q is consistent. Letting $(\mathcal{M}', a) \models Q$ (where a is meant to be substituted into the tuple x in the definition of Q) and $q := \text{tp}(a/B)$ we have $q \supseteq p$ and q is an heir of p.

Now to get the coheir, consider

$$R := p \cup \{\neg\theta(x, b) : \theta(x, y) \in \mathcal{L}(\tau_M) \text{ and there does not exist } a' \text{ from } M \text{ s.t. } \mathcal{M} \models \theta(a', b)\}.$$

Once again we claim: R *is consistent.*

Proof. (of claim) If not, then there exists some $\varphi(x) \in p$ and $\theta_1(x, b_1), \ldots, \theta_n(x, b_n)$ such that

$$\vdash \varphi(x) \longrightarrow \bigvee_i \theta_i(x, b_i)$$

and such that each $\theta_i(x, b_i)$ is *not* realized in $\mathcal{M}$.

Since φ is in p and p extends $\text{Th}(\mathcal{M}_M)$ and p is consistent, so

$$\mathcal{M}_M \models \exists x \varphi(x)$$

so let a be a witness, i.e. $\mathcal{M}_M \models \varphi(a)$. Since $\mathcal{M} \preccurlyeq \mathcal{N}$ and since

$$\vdash \varphi(x) \longrightarrow \bigvee_i \theta_i(x, b_i)$$

it follows that there is some i such that $\mathcal{N} \models \theta_i(a, b_i)$, i.e. $\theta_i(x, b_i)$ is realized in $\mathcal{M}$, which is a contradiction, thus the claim holds. □

Let $(\mathcal{M}', a) \models R$. Then $r = \text{tp}(a/B)$ is the desired coheir for p. □

So given p as in the theorem, we can extend it to get an heir (q) and a coheir (r).

Example. Let $\mathcal{M} = ((0, 1), <)$ thought of as an elementary substructure of $\mathfrak{B} = (\mathbb{R}, <)$ and let p be the type generated by the formula $\{x > a : a \in (0, 1)\}$. What do the heirs and coheirs of p look like?

Suppose $q \in S(\mathbb{R})$ is an heir of p, and suppose $r \in S(\mathbb{R})$ is a coheir of p. Is $x > 2$ in r or q?

It cannot be in r since $x > 2$ then there would have to be some $a \in (0, 1)$ such that $a > 2$ was satisfied by $\mathfrak{B}$. Of course there isn't. In general r is (generated by) $\{a < x : a < 1\}$.

Now $x > 2$ *is* in q since the formula $\psi(x, y) = \neg(x > y)$ is *not* represented in p. In general q is (generated by) $\{x > a : a \in \mathbb{R}\}$.

Note that the heirs and coheirs of p are different.

Example. In many cases it is in fact the case that there are at most two distinct coheirs. We can modify the above example slightly to get an example where p has two coheirs.

Let $\tilde{p} \in S(\mathbb{R})$ be $\{x > a : a \leqslant \frac{1}{2}\} \cup \{x < a : a > \frac{1}{2}\}$. Which says of x that is is "infinitesimally greater than $\frac{1}{2}$". Then considering $\mathbb{R}$ as an elementary substructure of a model, $\mathcal{M}$, which has infinitesimals, then there are two coheirs

$$q^+ = \{x > a : a \leqslant \frac{1}{2}, a \in M\} \cup \{x < a : a \geqslant \frac{1}{2}\}$$

and

$$q^- = \{x > a : a \in M \text{ such that } \forall r \in \mathbb{R} \text{ if } r > \frac{1}{2}, \text{ then } a < r\}.$$

20.2 Preservation Theorems

Earlier in the course we observed that theories which have certain simple syntactic characterizations are also preserved under certain semantic operations. For instance if T admits a universal (i.e. $\forall_1$) axiomatization, then the class of models of T is closed under substructures. We now turn to proving converses of these statements.

Theorem 20.1. *(Łos-Tarski) If T is a theory in some language $\mathcal{L}(\tau)$ then the following are equivalent.*

1) *If $\mathfrak{A} \subseteq \mathfrak{B}$ with $\mathfrak{B} \models T$, then $\mathfrak{A} \models T$.*
2) *There is a set U of universal sentences such that T and U have exactly the same models.*

We have already seen that the second condition implies the first. We shall prove the converse implication shortly.

Notation. (As in Hodges) Given τ-structures $\mathfrak{A}$ and $\mathfrak{B}$, we write $\mathfrak{A} \Rrightarrow_\Delta \mathfrak{B}$, for a set of τ-sentences, if for all $\delta \in \Delta$ we have

$$\mathfrak{A} \models \delta \qquad \Longrightarrow \qquad \mathfrak{B} \models \delta$$

The case where Δ is the set of existential (i.e. $\exists_1$) sentences is written $\mathfrak{A} \Rrightarrow_\exists \mathfrak{B}$.

We now prove that if $\mathfrak{A} \Rrightarrow_\exists \mathfrak{B}$ then $\mathfrak{A}$ is a substructure of (an elementary extension of) $\mathfrak{B}$.

Notation. In the following we shall use the notation $\mathrm{Diag}(\mathfrak{A})$, for the atomic diagram. Unlike the old notion ($\mathrm{diag}(\mathfrak{A})$) this will contain all quantifier-free formulae, i.e. it is closed under conjunction and disjunction. This slightly expands the notion used previously, but no extra information is used.

Proposition. *The following are equivalent.*

1) $\mathfrak{A} \Rrightarrow_\exists \mathfrak{B}$.
2) *There exists* $\mathfrak{C}$ *such that* $\mathfrak{B} \preccurlyeq \mathfrak{C}$ *and* $\mathfrak{A}$ *embeds into* $\mathfrak{C}$.

Proof. 2) $\Rightarrow$ 1). We have $\mathfrak{C} \equiv \mathfrak{B}$, so in particular, $\mathrm{Th}_\forall(\mathfrak{C}) = \mathrm{Th}_\forall(\mathfrak{B})$. We know that if $\mathfrak{A}$ is (isomorphic to) a substructure of $\mathfrak{C}$ then $\mathrm{Th}_\forall(\mathfrak{C}) \subseteq \mathrm{Th}_\forall(\mathfrak{A})$. In other words

$$\mathrm{Th}_\exists(\mathfrak{A}) \subseteq \mathrm{Th}_\exists(\mathfrak{C}) = \mathrm{Th}_\exists(\mathfrak{B})$$

which is to say $\mathfrak{A} \Rrightarrow_\exists \mathfrak{B}$.

1) $\Rightarrow$ 2). Consider the theory

$$T := \mathrm{eldiag}(\mathfrak{B}) \cup \mathrm{diag}(\mathfrak{A})$$

(making sure that the new constant symbols for A and B don't overlap). We claim that T is consistent. If not then, by compactness, there exists some $\varphi(\mathfrak{b}) \in \mathrm{eldiag}(\mathfrak{B})$ and $\psi(\mathfrak{a}) \in \mathrm{diag}(\mathfrak{A})$ such that $\mathfrak{b}$ and $\mathfrak{a}$ are new constant symbols and such that

$$\vdash \varphi(\mathfrak{b}) \longrightarrow \neg\psi(\mathfrak{a}).$$

In particular

$$\mathfrak{B} \models \forall x[\varphi(x) \longrightarrow \neg\psi(x)].$$

Since $\varphi(\mathfrak{b}) \in \mathrm{eldiag}(\mathfrak{B})$ the above implication shows that $\mathfrak{B} \models \forall x \neg\psi(x)$. But $\mathfrak{A} \Rrightarrow_\exists \mathfrak{B}$ and $\mathfrak{A} \models \exists x \psi(x)$, which is a contradiction. Thus T is consistent. Letting $\mathfrak{C} \models T$, we have that $\mathfrak{C}|_\tau$ is the desired structure. $\square$

Example. In the proof above we used the fact that $\mathrm{Th}_\forall(\mathfrak{C}) \subseteq \mathrm{Th}_\forall(\mathfrak{A})$ whenever $\mathfrak{A}$ is a substructure of $\mathfrak{C}$. We give an example where the containment is strict. Consider $\mathbb{Z} \subseteq \mathbb{R}$ where they are considered τ-structures for $\tau = \{<, 0, 1\}$. Then $\mathbb{Z}$ satisfies $\forall x[x = 1 \vee x = 0 \vee x < 0 \vee x > 1]$ which $\mathbb{R}$ does not.

We can now prove Theorem 20.1.

Proof. (of Łos-Tarski, Theorem 20.1) We have already seen 1) $\Rightarrow$ 2).

2) $\Rightarrow$ 1). Let $U := T_\forall$ the set of all universal consequences of T. We must show that $\mathrm{Mod}(T) = \mathrm{Mod}(U)$. If $\mathfrak{A} \models T$ then clearly $\mathfrak{A} \models U$. So $\mathrm{Mod}(T) \subseteq \mathrm{Mod}(U)$.

Now suppose $\mathfrak{A} \models U$. Let $S := T \cup \mathrm{Diag}(\mathfrak{A})$. We claim that S is consistent. If not then there exists $\varphi(\mathfrak{a}) \in \mathrm{Diag}(\mathfrak{A})$ such that $T \vdash \neg\varphi(\mathfrak{a})$. Thus $T \vdash \forall x \neg\varphi(x)$ (since $\mathfrak{a}$ was a new variable). Since φ is quantifier-free we now see that $\forall x \neg\varphi(x) \in T_\forall$. But sine $\mathfrak{A} \models T_\forall$ and since $\mathfrak{A} \models \exists x \varphi(x)$ we have a contradiction. So S is indeed consistent. Let $\mathfrak{C}$ be a model of S. Then $\mathfrak{A}$ embeds into $\mathfrak{C}$ and we have $\mathfrak{C} \models T$. By condition 1) we have that $\mathfrak{A} \models T$. This completes the proof. $\square$

There are many similar kinds of preservation theorems for different types of syntactic classes. Many examples can be found in Chang and Keisler's book [2].

Lecture 21

21.1 Interpolation

Theorem 21.1. *Given signatures $\tau_1, \tau_2 \supseteq \sigma$ such that $\tau_1 \cap \tau_2 = \sigma$, and given $\mathfrak{A}_1 \in \mathrm{Str}(\tau_1)$ and $\mathfrak{A}_2 \in \mathrm{Str}(\tau_2)$ such that $\mathfrak{A}_1|_\sigma \equiv \mathfrak{A}_2|_\sigma$, then there exists $\mathfrak{B}$ a $\tau_1 \cup \tau_2$-structure such that $\mathfrak{A}_1 \preccurlyeq \mathfrak{B}|_{\tau_1}$ and $\mathfrak{A}_2 \preccurlyeq \mathfrak{B}|_{\tau_2}$.*

Proof. Consider the theory

$$\mathrm{eldiag}(\mathfrak{A}_1) \cup \mathrm{eldiag}(\mathfrak{A}_2).$$

A model of this theory would suffice. If no such model exists then by compactness there are $\varphi(a) \in \mathrm{eldiag}(\mathfrak{A}_1)$ and $\psi(b) \in \mathrm{eldiag}(\mathfrak{A}_2)$ where a and b are new constants, $\psi, \varphi \in \mathcal{L}(\sigma)$ and

$$\vdash \varphi(a) \longrightarrow \neg\psi(b).$$

Thus any expansion of $\mathfrak{A}_1$ to an $\mathcal{L}(\tau_{1,b})$-structure must satisfy $\neg\psi(b)$, so $\mathfrak{A}_1 \models \forall x \neg\psi(x)$. Now $\forall x \neg\psi(x) \in \mathcal{L}(\sigma)$ and so $\mathfrak{A}_1|_\sigma \models \forall x \neg\psi(x)$. But since $\mathfrak{A}_1|_\sigma \equiv \mathfrak{A}_2|_\sigma$ we must have $\mathfrak{A}_2|_\sigma \models \forall x \neg\psi(x)$, contradicting the fact that $\psi(b) \in \mathrm{eldiag}(\mathfrak{A}_2)$. $\square$

From this theorem we get two syntactic consequences.

Notation. For T a τ-theory and $\sigma \subseteq \tau$ we denote by T_σ, the set of all σ-consequences of T, i.e. $\{\psi \in \mathcal{L}(\sigma) : T \vdash \psi\}$.

Corollary. *If $\sigma \subseteq \tau$ is an extension of signatures and T is a τ-theory, then a σ-structure $\mathfrak{A}$ satisfies T_σ if and only if there is a model $\mathfrak{B}$ of T such that $\mathfrak{A} \preccurlyeq \mathfrak{B}|_\sigma$.*

Proof. Let $\mathfrak{A} \models T_\sigma$. Consider the theory

$$T \;\cup\; \mathrm{eldiag}(\mathfrak{A})$$

(remember that $\mathrm{eldiag}(\mathfrak{A})$ is a σ-theory). If this were a consistent theory the we would be done. If not then, by compactness, there is some $\psi(a) \in \mathrm{eldiag}(\mathfrak{A})$ such that $T \cup \{\psi(a)\}$ is inconsistent. Here $\psi \in \mathcal{L}(\sigma)$ and a is a tuple of new constants. So we have

$$T \vdash \forall x \neg\psi(x)$$

i.e. $\forall x \neg\psi(x) \in T_\sigma$, contradicting the fact that $\mathfrak{A} \models T_\sigma$. Thus we get the desired model.

The converse implication is clear. $\square$

Remark. Note that the Corollary does *not* claim that $\mathfrak{A}$ is a reduct of a model of T. To see that this is false in general, consider $\sigma = \{<\}$ and $\tau = \{<, +, 0\}$ and T the theory of divisible ordered abelian groups. Then T_σ is the theory of dense linear orders without endpoints. Then $\mathbb{Q} \oplus \mathbb{R} \models T_\sigma$, but there is no way to order $\mathbb{Q} \oplus \mathbb{R}$ to make it satisfy T (since it is not homogeneous).

Corollary. *(Interpolation Theorem) Given $\tau_1, \tau_2 \supseteq \sigma$ with $\tau_1 \cap \tau_2 = \sigma$ and T_1, T_2 theories in $\mathcal{L}(\tau_1)$, $\mathcal{L}(\tau_2)$ respectively. If $T_1 \cup T_2$ is inconsistent, then there is a sentence $\psi \in \mathcal{L}(\sigma)$ such that $T_1 \vdash \psi$ and $T_2 \vdash \neg\psi$.*

Proof. Consider the theory $(T_1)_\sigma \cup (T_2)_\sigma$. If this is inconsistent then we're done. If it is consistent then let $\mathfrak{A}$ be a model. Note that $\mathfrak{A}$ is a σ-structure. By the Corollary there exists a model $\mathfrak{B}_1 \models T_1$ such that $\mathfrak{A} \preccurlyeq \mathfrak{B}_1|_\sigma$, and a model $\mathfrak{B}_2 \models T_2$ such that $\mathfrak{A} \preccurlyeq \mathfrak{B}_2|_\sigma$. But then $\mathfrak{B}_1|_\sigma \equiv \mathfrak{B}_2|_\sigma$ and so by the Theorem, there exists some $\tau_1 \cup \tau_2$-structure $\mathfrak{C}$ such that $\mathfrak{B}_1 \preccurlyeq \mathfrak{C}|_{\tau_1}$ and $\mathfrak{B}_2 \preccurlyeq \mathfrak{C}|_{\tau_2}$. But then $\mathfrak{C} \models T_1 \cup T_2$ contrary to our assumption. $\square$

As a Corollary we get the Craig Interpolation Theorem.

Theorem 21.2. *(Craig Interpolation) Given τ_1, τ_2 and σ as above, and $\varphi \in \mathcal{L}(\tau_1)$, $\psi \in \mathcal{L}(\tau_1)$. If $\varphi \vdash \psi$ then there exists $\theta \in \mathcal{L}(\sigma)$ such that $\varphi \vdash \theta$ and $\theta \vdash \psi$.*

One other consequence of the theorem is called Beth's Definability Theorem. It states roughly that explicit and implicit definability are equivalent for first-order logic.

Theorem 21.3. *(Beth's Definability Theorem) Let $\sigma \subseteq \tau$ be an extension of signatures, T a τ-theory and $\varphi(x) \in \mathcal{L}(\tau)$. Then the following are equivalent.*

1) *(Implicit) For all models $\mathfrak{A}, \mathfrak{B} \models T$, if $\mathfrak{A}|_\sigma = \mathfrak{B}|_\sigma$ then $\varphi(\mathfrak{A}) = \varphi(\mathfrak{B})$.*
2) *(Explicit) There is some $\psi \in \mathcal{L}(\sigma)$ such that $T \vdash \varphi \leftrightarrow \psi$.*

Proof. 2) $\Longrightarrow$ 1). Clear.

1) $\Longrightarrow$ 2). Consider the theory (in the extended language $\mathcal{L}(\tau_{a,b})$)

$$S := T \cup \{\varphi(a)\} \cup \{\neg\varphi(b)\} \cup \{\psi(a) \leftrightarrow \psi(b) : \psi \in \mathcal{L}(\sigma)\}.$$

If S is inconsistent then we're done, since by compactness there would be a finite set Ψ of $\mathcal{L}(\sigma)$-formulae such that

$$T \vdash \forall x, y \left[\bigwedge_{\psi \in \Psi} (\psi(x) \leftrightarrow \psi(y)) \longrightarrow (\varphi(x) \leftrightarrow \varphi(y)) \right].$$

Now set

$$\theta := \bigvee_{\substack{\mu \subseteq \Psi \text{ s.t.} \\ T \cup \{\varphi(x)\} \cup \bigwedge \mu \cup \bigwedge_{\psi \notin \mu} \neg\psi \\ \text{is consistent}}} \left(\bigwedge_{\psi \in \mu} \psi \wedge \bigwedge_{\psi \notin \mu} \neg\psi \right)$$

Then $T \vdash \varphi \leftrightarrow \theta$.

So suppose S is consistent. Let $(\mathfrak{C}, a, b)$ be a model of S. We will now apply Theorem 21.1. We define two new signatures. Let τ_1 and τ_2 be disjoint copies (over σ) of τ together with a new constant symbol c. More precisely we can decorate each symbol of $\tau \setminus \sigma$ with a subscript either 1 or 2. Thus τ_1 consists of symbols from σ together with symbols $x^{(1)}$ for all $x \in \tau \setminus \sigma$ and also contains the new symbol c. Similarly for τ_2. So $\tau_1, \tau_2 \supseteq \sigma_c$ and $\tau_1 \cap \tau_2 = \sigma_c$.

Let $\mathcal{M}$ be $(\mathfrak{C}, a)$ considered as a τ_1-structure, i.e. $c^{\mathcal{M}} = a$ and $x^{(1)} = x^{\mathfrak{C}}$. Similarly let $\mathcal{N}$ be $(\mathfrak{C}, b)$ considered as a τ_2-structure.

Now since we arranged that a and b have the same σ-type (since $(\mathfrak{C}, a, b) \models S$) we have that

$$\mathcal{M}|_{\sigma_c} \equiv \mathcal{N}|_{\sigma_c} \qquad \text{i.e. } (\mathfrak{C}|_\sigma, a) \equiv (\mathfrak{C}|_\sigma, b).$$

Now by Theorem 21.1 there exists $\mathfrak{D}$ a $(\tau_1 \cup \tau_2)$-structure such that $\iota_1 : \mathcal{M} \preccurlyeq \mathfrak{D}|_{\tau_1}$ and $\iota_2 : \mathcal{N} \preccurlyeq \mathfrak{D}|_{\tau_2}$ (elementary embeddings). Note that $\iota_1(a) = \iota_2(b)$ since $c^{\mathfrak{D}} = a$ and $c^{\mathfrak{D}} = b$.

Let $\mathfrak{A}$ be $\mathfrak{D}|_{\tau_1}$ regarded as a τ-structures (i.e. forgetting the constant c). Similarly let $\mathfrak{B}$ be $\mathfrak{D}|_{\tau_2}$ regarded as a τ-structure. Then we have a literal equality

$$\mathfrak{A}|_\sigma = \mathfrak{B}|_\sigma$$

since then are both equal to $\mathfrak{D}|_\sigma$. Now $\mathfrak{A}$ and $\mathfrak{B}$ both model T since $\mathfrak{C}$ was a model of T. But they disagree on φ, i.e. $\varphi(\mathfrak{A}) \neq \varphi(\mathfrak{B})$ since $\mathfrak{A} \models \varphi(c^{\mathfrak{D}})$ and $\mathfrak{B} \models \neg\varphi(c^{\mathfrak{D}})$. This contradicts the assumption (1), thus S must be inconsistent. This completes the proof. $\square$

21.2 Indiscernibles

Indiscernibles are a tool for analyzing structures by making them much more homogeneous. By making them more homogeneous we can take local information and expand it to get global information about the structures.

Definition. A sequence $(a_i)_{i\in\omega}$ in some τ-structure $\mathfrak{A}$ is an **indiscernible sequence** if for any formula $\varphi(x_0,\dots,x_{n-1})$ and $i_0 < \dots < i_{n-1}$ and $j_0 < \dots < j_{n-1}$ increasing sequences from ω then

$$\mathfrak{A} \models \varphi(a_{i_0},\dots,a_{i_{n-1}}) \longleftrightarrow \varphi(a_{j_0},\dots,a_{j_{n-1}}).$$

Remark. If $(a_i)_{i\in\omega}$ is an indiscernible sequences then the type of an increasing n-sequence $a_{i_0},\dots,a_{i_{n-1}}$ is constant, i.e. is the same for all such increasing n-sequences from $(a_i)_{i\in\omega}$.

In particular any two elements a_i and a_j from the sequence have the same type.

Definition. If the order does not matter then the sequence $(a_i)_{i\in\omega}$ is called an **indiscernible set**. More precisely the requirement is that for any set $J \subseteq \omega$ of size n, say $J = \{j_0,\dots,j_{n-1}\}$ then

$$\mathfrak{A} \models \varphi(a_0,\dots,a_{n-1}) \longleftrightarrow \varphi(a_{j_0},\dots,a_{j_{n-1}}).$$

Of course an indiscernible set is in particular an indiscernible sequence.

Example. If $a_0 < a_1 < \dots \in \mathbb{Q}$ then $(a_i)_{i\in\omega}$ is an indiscernible sequence in $\mathbb{Q}$ considered as an ordered structure. It is not an indiscernible *set*.

Example. If X is any infinite set in the language of equality (i.e. $\tau = \varnothing$) and $(a_i)_{i\in\omega}$ is any sequence without repititions from X then it is an indiscernible set. Alternatively if $(a_i)_{i\in\omega}$ is the constant sequence then it is also an indiscernible set.

Example. Let V be a vector space over a field k in the language of vector spaces $\tau = \{+, (\lambda)_{\lambda\in k}\}$. Then any linearly independent set $X \subseteq V$ is an indiscernible set. To see this note that we can extend X to a basis for V, and that a change of bases extends to an automorphism of V.

Our goal is to show the following:

Proposition. *If* T *is any* τ*-theory and* $\Sigma(x)$ *a set of* $\mathcal{L}(\tau_x)$*-formulae such that it is consistent that there exists a model* $\mathfrak{A}$ *of* T *such that* $\Sigma(\mathfrak{A})$ *is infinite, then there exists a model* $\mathfrak{B}$ *of* T *and a sequence* $(a_i)_{i\in\omega}$ *which is non-constant and is an indiscernible sequence such that* $\mathfrak{B} \models \Sigma(a_i)$ *for all* $i \in \omega$.

We shall begin the proof, but we will need Ramsey's theorem at some point. The proof of Ramsey's theorem will be given afterwards.

Proof. (Assuming Ramsey's Theorem) We write down what we want: Let S be the theory

$$T \cup \bigcup_{i=0}^{\infty} \Sigma(x_i) \cup \{x_i \neq x_j \mid i \neq j\}$$
$$\cup \{\psi(x_{i_0},\dots,x_{i_{n-1}}) \leftrightarrow \psi(x_{j_0},\dots,x_{j_{n-1}}) \mid \psi \in \mathcal{L}(\tau), i_0 < \dots < i_{n-1} \text{ and } j_0 < \dots < j_{n-1}\}$$

If S is consistent then we are done since the interpretations of the x_i's would be a non-constant indiscernible sequence.

Suppose therefore that S is not consistent. Then by compactness there is some finite fragment which is inconsistent. Then there is some $N \in \mathbb{N}$ such that the theory

$$\begin{aligned}&T \cup \{\theta(x_i) \mid i \leqslant N\} \cup \{x_i \neq x_j \mid i \neq j \leqslant N\}\\ &\cup \{\psi_k(x_{i_0}, \dots, x_{i_{n_k-1}}) \leftrightarrow \psi_k(x_{j_0}, \dots, x_{j_{n_k-1}}) \mid \psi_k \in \mathcal{L}(\tau), k \leqslant K,\\ &i_0 < \cdots < i_{n-1} \leqslant N \text{ and } j_0 < \cdots < j_{n-1} \leqslant N\}\end{aligned}$$

is inconsistent. We may assume (by way of padding) that there is some n such that $n_k = n$ for all k.

Now we know that there is some model $\mathfrak{A}$ of T such that θ has infinitely many realizations, i.e. $|\theta(\mathfrak{A})| \geqslant \aleph_0$. Let $b_0, b_1, \dots$ be a sequence of distinct elements from $\theta(\mathfrak{A})$.

Notation. The set $[\omega]^n$ consists of all strictly increasing n-tuples. I.e. $[\omega]^n := \{(l_1, \dots, l_n) \in \omega^n : l_1 < l_2 < \cdots < l_n\}$.

Define a function
$$f : [\omega]^n \longrightarrow \mathcal{P}(\{1, \dots, k\})$$
by
$$f(i_0, \dots, i_{n-1}) := \{k \mid \mathfrak{A} \models \psi_k(b_{i_0}, \dots, b_{i_{n-1}})\}$$

Now f is a function from $[\omega]^n$ to a finite set. By Ramsey's theorem (see below for statement and proof) there exists $H \subseteq \omega$ infinite and homogeneous, i.e. $f|_{[H]^n}$ is constant. Let $H = \{a_0 < a_1 < \dots\}$. Interpret x_i in $\mathfrak{A}$ as a_i. This will satisfy out purportedly inconsistent sub theory. This yields a contradiction and completes the proof (modulo Ramsey's theorem). □

We need fill the gap in the above proof.

Notation. In the course of the proof we introduced the notation $[\omega]^n$ for all increasing n-sequences from ω.

Theorem 21.4. *(Ramsey's Theorem) Given a function f from $[\omega]^n$ (for some $n \in \omega$) to a finite set, then there exists an infinite subset H of ω such that f is constant $[H]^n$*

Proof. We may assume that the codomain of f is in fact $\{0, \dots, N-1\}$ (where N is to the cardinality of the codomain).

Consider the structure $\mathfrak{A} = (\omega, <, \{k\}_{k \in \omega}, f)$, where $<$ is interpreted as the standard order on ω and $f^{\mathfrak{A}}$ is interpreted to be the same as the given function f expect that $f^{\mathfrak{A}}(b_0, \dots, b_{n-1}) = 0$ if $(b_0, \dots, b_{n-1}) \notin [\omega]^n$.

We will prove the theorem by induction on n.

For $n = 1$ the theorem follows from the pigeon hole principle.

For $n + 1$, suppose the theorem holds for all integers $\leqslant n$.

Take a *proper* elementary extension $\mathfrak{A}^*$ of $\mathfrak{A}$, which is possible by upward Löwenheim-Skolem. In particular $\mathfrak{A} \equiv \mathfrak{A}^*$. So $(\mathfrak{A}^*, <)$ is a linear order. Let $a \in \operatorname{dom}(\mathfrak{A}^*) \setminus \omega$ be a new element from $\mathfrak{A}^*$. Note that $a > n$ for every $n \in \omega$ since for all $n \in \omega$ the structure $\mathfrak{A}$ satisfies that n has exactly n predecessors, hence $\mathfrak{A}^*$ must satisfy this as well. But then we cannot have $a \leqslant n$ for any $n \in \omega$ and so by the linearity of the order we must have $a > n$. So a is an "infinite" number in $\mathfrak{A}^*$.

We construct an increasing sequence $m_0 < m_1 < \dots$ from ω. The first n elements are not important we just pick them such that $m_0 < m_1 < \cdots < m_{n-1}$. Now with $m_0 < \cdots < m_{j-1}$ constructed we search for an element $x > m_{j-1}$ such that

- for each $i_0 < \cdots < i_{n-1} \leqslant j - 1$ we have

$$f^{\mathfrak{A}^*}(m_{i_0}, \dots, m_{i_{n-1}}, a) = f^{\mathfrak{A}^*}(m_{i_0}, \dots, m_{i_{n-1}}, x).$$

I.e. x must behave like a with respect to the sequence $m_{i_0}, \dots, m_{i_{n-1}}$. Such an x will then be the j'th element of the sequence $m_0 < m_1 < \dots$. This puts finitely many constraints on x and so we can write it out as a first-order formula.

Consider the formula $\theta(x)$ given by

$$x > m_{j-1} \ \wedge \bigwedge_{i_0 < \cdots < i_{n-1} \leqslant j-1} f(m_{i_0}, \dots, m_{i_{n-1}}, x) = f^{\mathfrak{A}^*}(m_{i_0}, \dots, m_{i_{n-1}}, a)$$

[Note: the first instance of the symbol f in θ is just a symbol, the second instance "$f^{\mathfrak{A}^*}(...)$" is the actual value of $f^{\mathfrak{A}^*}$ on the tuple $(m_{i_0}, \dots, m_{i_{n-1}}, a)$, i.e. a number in $\{0, \dots, N-1\}$.]

We we have $\mathfrak{A}^* \models \theta(a)$ and so $\mathfrak{A}^* \models \exists x \theta(x)$. Now since $\mathfrak{A} \preccurlyeq \mathfrak{A}^*$ we have $\mathfrak{A} \models \exists x \theta(x)$. So let m_j be a witness, then we have the next element of the sequence: $m_0 < \cdots < m_{j-1} < m_j$.

Now we use the induction hypothesis: Define $g : [\omega]^n \longrightarrow \{0, \dots, N\}$ by

$$g(l_1, \dots, l_n) = f^{\mathfrak{A}^*}(m_{l_1}, \dots, m_{l_n}, a).$$

By the induction hypothesis there exists a homogenous set H such that g is constant on $[H]^n$. Now we claim that f is constant on the subset of H given by the l sequence m_l, i.e.

$$f|_{[\{m_l : l \in H\}]^{n+1}} \text{ is constant}$$

To see this, suppose $i_0 < \cdots < i_n$ and $j_0 < \cdots < j_n$ then

$$\begin{aligned} f(m_{l_{i_0}}, \dots, m_{l_{i_n}}) &= f^{\mathfrak{A}^*}(m_{l_{i_0}}, \dots, m_{l_{i_{n-1}}}, a) \\ &= g(l_{i_0}, \dots, l_{n-1}) \\ &= g(l_{j_0}, \dots, l_{j_{n-1}}) \\ &= f^{\mathfrak{A}^*}(m_{l_{j_0}}, \dots, m_{l_{j_{n-1}}}, a) \\ &= f(m_{l_{j_0}}, \dots, m_{l_{j_n}}). \end{aligned}$$

The theorem follows by induction. $\square$

Lecture 22

22.1 Finite Ramsey Theory

Notation. For $\kappa, \lambda, \mu, \nu$ cardinals, we write $\kappa \to (\lambda)^{\mu}_{\nu}$ if, for all functions $f : [\kappa]^{\mu} \longrightarrow \nu$ there exists $H \subseteq \kappa$ such that $|H| = \lambda$ and such that $f|_{[H]^{\mu}}$ is constant.

In Lecture 21 we proved Ramsey's Theorem which can be briefly stated as:

$$\aleph_0 \to (\aleph_0)^m_n \quad \text{holds for all } n, m \in \omega.$$

As a Corollary we get the Finite Ramsey's Theorem.

Corollary. *For all* $m, n, l \in \omega$ *there is some* $k \in \omega$ *such that* $k \to (l)^m_n$.

Proof. We use the compactness theorem. Fix $m, n, l \in \omega$. Suppose for contradiction that no k exists.

Consider the signature $\tau = \{<, f, \underline{0}, \underline{1}, \ldots, \underline{n-1}\}$, where f an m-ary function symbol. Then the following set of formulae is consistent

$$\begin{aligned} &\{x_i < x_j : i < j < \omega\} \cup \text{“} < \text{is a total order”} \cup \text{“}\underline{0} < \cdots < \underline{n-1}\text{”} \\ &\cup \left\{ \forall y_1, \ldots, y_m \bigvee_{i=0}^{n-1} f(\bar{y}) = \underline{i} \right\} \\ &\cup \text{“there is no homogeneous set of size } l \text{ for this function”} \end{aligned}$$

The last statement is in fact a collection of first order statements, we can write it as follows:

$$\left\{ \bigvee_{\substack{|S|=|T|=n \\ S,T \subseteq \{i_1,\ldots,i_l\}}} f(x_s) \neq f(x_t) \;:\; i_1 < \cdots < i_l \right\}$$

i.e. for each choice of l elements $i_1 < \ldots, < i_l$ from ω and increasing n-sequences from this set, f is not constant.

By the assumption that no k exists, the above theory is consistent. So let $\mathfrak{A}$ be a model and define $g : [\omega]^m \longrightarrow n$ by

$$(i_1 < \cdots < i_m) \longmapsto f^{\mathfrak{A}}(x_{i_1}^{\mathfrak{A}}, \ldots, x_{i_m}^{\mathfrak{A}}).$$

By (the Infinite) Ramsey's Theorem there exists $H \subseteq \omega$ such that $|H| = \aleph_0$ and such that $g|_{[H]^m}$ is constant. But this is a contradiction since then any size l-subset $\{x_h : h \in H\}$ is homogeneous. $\square$

So if one wants to find some homogeneous set of size l, then it can be done, provided one chooses a big enough domain.

We will return to Indiscernibles when we come to Morley's Theorem.

22.2 Fraïssé Constructions and Ages

Definition. For $\mathcal{K}$ a class of τ-structures the **age of** $\mathcal{K}$, written $\mathrm{age}(\mathcal{K})$, is the class of all finitely generated τ-structures $\mathfrak{A}$ such that there exists $\mathfrak{B} \in \mathcal{K}$ and an embedding of $\mathfrak{A}$ into $\mathfrak{B}$.

Notation. If $\mathcal{K} = \{\mathcal{M}\}$ consists of a single τ-structure then we write $\mathrm{age}(\mathcal{M})$ for the age of $\mathcal{K}$.

The way the age of $\mathcal{M}$ is defined above it will not be a set. We often will consider $\mathcal{K}$ *up to isomorphism* in which case it will be an actual set.

Remark. Often one requires τ to be a finite relational signature. In this case a finitely generated substructure is the same as a finite substructure.

Example. The age of $(\mathbb{Q}, <)$ is the class of finite linear orders.

Question. What is the age of the reals $(\mathbb{R}, +, -, \cdot, 0, 1)$ as a field? [Hint: finitely generated archimedeanly orderable rings].

Example. Let $\mathcal{K} = \{(H, 1, \cdot, ^{-1}) : H \text{ is a finite group}\}$ be the class of finite groups. Let $G = \coprod_{H \in \mathcal{K}} H$ be the direct limit of all finite groups. Then $\mathrm{age}(G) = \mathcal{K}$.

Question. For which $\mathcal{K}$ is there some single τ-structure $\mathcal{M}$ such that $\mathcal{K} = \mathrm{age}(\mathcal{M})$.

In the following we will provide a complete answer for this question.

One necessary condition is the following.

Definition. We say that $\mathcal{K}$ has the **Hereditary Property** (HP) if, when $\mathfrak{B} \in \mathcal{K}$ and $\mathfrak{A}$ a finitely generated substructure of $\mathfrak{B}$, then $\mathfrak{A} \in \mathcal{K}$.

This is however not sufficient. There is another necessary condition, which will turn out to be sufficient (together with HP).

Definition. We say that $\mathcal{K}$ has the **Joint Embedding Property** (JEP) if for all $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$ there is some $\mathfrak{C} \in \mathcal{K}$ such that $\mathfrak{A}$ and $\mathfrak{B}$ can both be embedded into $\mathfrak{C}$.

$$\begin{tikzcd} & \mathfrak{C} & \\ \mathfrak{A} \arrow[ur, hook, dashed] & & \mathfrak{B} \arrow[ul, hook', dashed] \end{tikzcd}$$

Note that JEP is necessary since for $\mathfrak{A}, \mathfrak{B} \in \text{age}(\mathcal{M})$ then the substructure generated by the generators of $\mathfrak{A}$ and $\mathfrak{B}$ is again a finitely generated substructure, i.e. an element of $\mathcal{M}$.

Proposition. *Let τ be a countable signature. If $\mathcal{K}$ is a countable set of finitely generated τ-strucutres that has HP and JEP, then there exists a τ-structure $\mathcal{M}$ such that $\mathcal{K} = \text{age}(\mathcal{M})$.*

Proof. $\mathcal{K}$ is countable, so let us list its elements $\mathcal{K} = \{\mathfrak{A}_i : i \in \omega\}$. We will build a sequence $(\mathfrak{B}_n)_{n\in\omega}$ by recursion. We arrange that

1) $(\mathfrak{B}_n)_{n\in\omega}$ is a chain, i.e. $\mathfrak{B}_0 \subseteq \mathfrak{B}_1 \subseteq \cdots$.
2) Each $\mathfrak{B}_i \in \mathcal{K}$.
3) For all $j < i$, $\mathfrak{A}_j$ is a substructure of $\mathfrak{B}_i$.

To start, we let $\mathfrak{B}_0 := \mathfrak{A}_0$.

At stage n in the construction we have constructed $\mathfrak{B}_{n-1}$ satisfying the conditions. By the Joint Embedding Property there is some $\mathfrak{B}_n \in \mathcal{K}$ such that $\mathfrak{A}_n$ and $\mathfrak{B}_{n-1}$ may be embedded into $\mathfrak{B}_n$.

This produces the sequence $(\mathfrak{B}_n)_{n\in\omega}$. Now let $\mathcal{M} := \bigcup_{n\in\omega} \mathfrak{B}_n$. Now we consider the age of $\mathcal{M}$.

For each $n \in \omega$, $\mathfrak{A}_n \subseteq \mathfrak{B}_n \subseteq \mathcal{M}$ so $\mathcal{K} \subseteq \text{age}(\mathcal{M})$. Now suppose $\mathfrak{C} \in \text{age}(\mathcal{M})$, say $\mathfrak{C}$ is generated by the finite set F. Since F is finite, there is some $n \in \omega$ such that $F \subseteq B_n$. So $\mathfrak{C}$ is a finitely generated substructure of $\mathfrak{B}_n$ and by the Hereditary Property of $\mathcal{K}$ we must have that $\mathfrak{C} \in \mathcal{K}$. Thus $\text{age}(\mathcal{M}) = \mathcal{K}$. $\square$

Remark. Note that we arranged that $\mathcal{M}$ has cardinality at most $\aleph_0$.

The next question we will consider is how unique $\mathcal{M}$ is. The answer will rely on the notion of ultrahomogeneity.

Lecture 23

23.1 Fraïssé's Theorem

Fraïssé's Theorem states that given a class $\mathcal{K}$ of finitely generated structures with HP, JEP and one further property (define below), we can construct a *unique* structure $\mathcal{M}$ such that the age of $\mathcal{M}$ will be $\mathcal{K}$ and such that $\mathcal{M}$ is countable and ultrahomogeneous (defined below).

Definition. We say that a class $\mathcal{K}$ of τ-structures, has the **Amalgamation Property** (AP) if, whenever we have a diagram

$$\begin{tikzcd}
\mathfrak{B} & & \mathfrak{C} \\
& \mathfrak{A} \arrow[ul, hook] \arrow[ur, hook] &
\end{tikzcd}$$

in $\mathcal{K}$, then there exists $\mathfrak{D} \in \mathcal{K}$ and embeddings ι and j such that

$$\begin{tikzcd}
& \mathfrak{D} & \\
\mathfrak{B} \arrow[ur, hook, dashed, "\iota"] & & \mathfrak{C} \arrow[ul, hook', dashed, "j"'] \\
& \mathfrak{A} \arrow[ul, hook] \arrow[ur, hook] &
\end{tikzcd}$$

is a commutative diagram.

Remark. In general the Amalgamation Property (AP) does not imply the Joint Embedding Property (JEP). As an example one can consider the class of fields. It does not have JEP since fields of different characteristic cannot be jointly embedded. However it does have the AP. [Details:...]

However, if τ is a relational signature, then HP and AP does imply JEP.

Remark. In general the Joint Embedding Property does not imply the Amalgamation Property. [Example:...]

Definition. Let τ be a finite relational signature. A τ-structure $\mathcal{M}$ is said to be **ultrahomogeneous** if whenever $\mathfrak{A}, \mathfrak{B} \subseteq \mathcal{M}$ are finite substructures and $f : \mathfrak{A} \to \mathfrak{B}$ is an isomorphism, then there exists $\sigma \in \mathrm{Aut}(\mathcal{M})$ such that $\sigma|_A = f$. I.e. any isomorphism between finite substructures can be extended to an automorphisms of $\mathcal{M}$.

Notation. Given $a \in M^n$ and $X \subseteq M$, the **quantifier-free type** of a over X if denoted by $\mathrm{qf.tp}(a/X)$. As usual if X is empty we write simply $\mathrm{qf.tp}(a)$.

Proposition. *Let τ be a finite relational signature. If $\mathcal{M}$ is ultrahomogeneous then* $\mathrm{Th}(\mathcal{M})$ *eliminates quantifiers.*

Proof. We first note that the types are determined by the quantifier-free types. I.e. we have that for $a, b \in M^n$ (say $a = (a_1, \ldots, a_n)$ and $b = (b_1, \ldots, b_n)$) then $\mathrm{qf.tp}(a) = \mathrm{qf.tp}(b)$ implies $\mathrm{tp}(a) = \mathrm{tp}(b)$. This is because, if $\mathrm{qf.tp}(a) = \mathrm{qf.tp}(b)$, then

$$f : \{a_1, \ldots, a_n\} \longrightarrow \{b_1, \ldots, b_n\} \qquad a_i \longmapsto b_i$$

is an isomorphism. By ultrahomogeneity we can extend f to an automorphism σ of $\mathcal{M}$. Thus $\sigma(a_i) = b_i$. Since automorphisms preserve all formulae it follows that $\mathrm{tp}(a) = \mathrm{tp}(b)$.

Now let $\theta(x_1, \ldots, x_n)$ be any formula. There are only finitely many atomic formulae $\psi_1, \ldots, \psi_l$ in n-variables. Let

$$\Theta := \left\{ S \subseteq \{1, \ldots, l\} : \mathcal{M} \models \exists \bar{x} \left(\theta(\bar{x}) \wedge \bigwedge_{i \in S} \psi_i(\bar{x}) \wedge \bigwedge_{i \notin S} \neg\psi_i(\bar{x}) \right) \right\}$$

Note that Θ is finite. Define θ^* to be the following formula

$$\theta^*(\bar{x}) := \bigvee_{S \in \Theta} \left(\bigwedge_{i \in S} \psi_i(\bar{x}) \wedge \bigwedge_{i \notin S} \neg\psi_i(\bar{x}) \right)$$

Note that θ^* is quantifier-free. We claim that

$$\mathrm{Th}(\mathcal{M}) \vdash \theta \longleftrightarrow \theta^*$$

Suppose $\mathcal{M} \models \theta(a)$. Then for some $S \subseteq \{1, \ldots, l\}$ we have

$$\mathcal{M} \models \bigwedge_{i \in S} \psi_i(a) \wedge \bigwedge_{i \notin S} \neg\psi_i(a).$$

So for this S we have $S \in \Theta$. Thus $\mathcal{M} \models \theta^*(a)$, so $\mathrm{Th}(\mathcal{M}) \vdash \theta \to \theta^*$.

Now suppose $\mathcal{M} \models \theta^*(b)$ so there is some $S \in \Theta$ such that

$$\mathcal{M} \models \left(\bigwedge_{i \in S} \psi_i(b) \wedge \bigwedge_{i \notin S} \neg\psi_i(b) \right)$$

and we also have that $\mathcal{M} \models \theta(b)$. Then $\mathrm{qf.tp}(a) = \mathrm{qf.tp}(b)$. So $\mathrm{tp}(a) = \mathrm{tp}(b)$ by the observation made in the beginning of the proof. Thus $\mathcal{M} \models \theta(a)$. This shows that $\mathrm{Th}(\mathcal{M}) \vdash \theta \leftrightarrow \theta^*$, so $\mathrm{Th}(\mathcal{M})$ eliminates quantifiers. $\square$

Proposition. *If $\mathcal{M}$ is ultrahomogeneous, then there are only finitely many n-types over $\varnothing$. I.e. $|S_n(\varnothing)| < \aleph_0$.*

Proof. Note that since τ is finite and relational it follows that the space $S_n^{qf}(\varnothing)$ of quantifier-free n-types, is finite. Ultrahomogeneity implies that the map

$$S_n(\varnothing) \longrightarrow S_n^{qf}(\varnothing)$$

(which we get from the quantifier-elimination) is a bijection. $\square$

We now Prove Fraïssé's Theorem. It consists of a uniqueness and existence claim. We begin with uniqueness.

Proposition. *(Fraïssé, Uniqueness) Let τ be a finite relational signature. Suppose $\mathcal{M}$ and $\mathcal{N}$ are countable ultrahomogeneous structures such that* $\mathrm{age}(\mathcal{M}) = \mathrm{age}(\mathcal{N})$ *and such that this common class has the Amalgamation Property. Then $\mathcal{M} \cong \mathcal{N}$.*

Remark. Note that $\mathcal{K} := \mathrm{age}(\mathcal{M}) = \mathrm{age}(\mathcal{N})$ automatically has HP and JEP. We will need these properties in the proof.

Proof. We follow the approach used to prove Cantor's Theorem saying $\mathrm{Th}(\mathbb{Q}, <)$ is $\aleph_0$-categorical, namely to list all elements, then build an isomorphism in finite stages.

List $M = \{m_i : i \in \omega\}$ and $N = \{n_i : i \in \omega\}$. We build, by recursion, na increasing sequence

$$f_0 \subseteq f_1 \subseteq \cdots$$

of isomorphisms between finite substructures of $\mathcal{M}$ and $\mathcal{N}$. We arrange that at stage l we have

- $\mathrm{dom}(f_l) \supseteq \{m_i : i < l\}$
- $\mathrm{range}(f_l) \supseteq \{n_i : i < l\}$.

At stage 0 we let $f_0 := \varnothing$. Since τ is relational this is indeed an isomorphism between finite substructures.

Now at stage l we have defined f_{l-1}. By the Amalgamation Property we get the following commutative diagram inside $\mathcal{K} := \operatorname{age}(\mathcal{M}) = \operatorname{age}(\mathcal{N})$:

$$\begin{tikzcd}
& \mathfrak{D} & \\
\mathfrak{B} = \operatorname{dom}(f_{l-1}) \cup \{m_{l-1}\} \arrow[ur, hook, dashed] & & \mathfrak{C} = \operatorname{range}(f_{l-1}) \arrow[ul, hook', dashed] \\
& \mathfrak{A} = \operatorname{dom}(f_{l-1}) \arrow[ul, hook] \arrow[ur, hook'] &
\end{tikzcd}$$

Now by definition of $\mathcal{K}$, $\mathfrak{D}$ is isomorphic to some finite substructure $\mathfrak{D}'$ of $\mathcal{M}$ and to some finite substructure $\mathfrak{D}''$ of $\mathcal{N}$. Let us choose isomorphisms $\alpha : \mathfrak{D} \longrightarrow \mathfrak{D}'$ and $\beta : \mathfrak{D} \longrightarrow \mathfrak{D}''$. There is a copy of $\mathfrak{B}$ inside of $\mathfrak{D}'$ since there is a copy of $\mathfrak{B}$ inside of $\mathfrak{D}$. Now $\alpha|_B : \mathfrak{B} \longrightarrow \mathfrak{B}' := \alpha(\mathfrak{B})$ is an isomorphism of finite substructures of $\mathcal{M}$. By ultrahomogeneity this map extends to an automorphisms $\sigma \in \operatorname{Aut}(\mathcal{M})$. Now $\sigma^{-1}(\mathfrak{D}') \supseteq \mathfrak{B}$ and $\sigma^{-1}(\mathfrak{D}') \cong \mathfrak{D}$. Thus, renaming if necessary, we may assume that $\mathfrak{D} \supseteq \mathfrak{B}$. Arguing in a similar way, but inside $\mathcal{N}$ we may assume $\mathfrak{D}'' \supseteq \mathfrak{C}$.

Now we define $f_{l-\frac{1}{2}} : \mathfrak{D} \longrightarrow \mathfrak{D}''$ to be β. Then $m_{l-1} \in \operatorname{dom}(f_{l-\frac{1}{2}})$. To get n_{l-1} into the range of f_l we reverse the roles of $\mathcal{M}$ and $\mathcal{N}$ above. This gives us f_l.

Now let $f := \bigcup_i f_i : \mathcal{M} \longrightarrow \mathcal{N}$. Then f is the required isomorphism. □

Remark. In the theorem we assumed that $\mathcal{M}$ and $\mathcal{N}$ are countable.

In general if $\mathcal{M}$ and $\mathcal{N}$ are ultrahomogeneous and $\operatorname{age}(\mathcal{M}) = \operatorname{age}(\mathcal{N})$ (with AP) then $\mathcal{M} \sim_\omega \mathcal{N}$ (i.e. $\mathcal{M} \equiv_{\infty,\omega} \mathcal{N}$.)

So if we can get Fraïssé limits which are countable ultrahomogeneous then they are unique. Now we prove the existence.

Proposition. *(Fraïssé, Existence) Let τ be a finite relational signature. If $\mathcal{K}$ is a countable set of τ-structures with HP, JEP and AP then there exists a τ-structure $\mathcal{M}$ which is countable, ultrahomogeneous and has age equal to $\mathcal{K}$.*

Remark. We will run through a similar construction as we used when we wanted to find an arbitrary limit for $\mathcal{K}$. This time, however, we must achieve ultrahomogeneity. So while we build $\mathcal{M}$ we will have certain instances of isomorphic finite substructures for which we should eventually extend to automorphisms of $\mathcal{M}$. What's more, as we do the construction new finite substructures will appear and so these will bring new constraints on the further construction. However there will only be finitely many new appearances at each stage and so we can dovetail the construction to take care of all the constraints.

Proof. We will build $\mathcal{M}$ as a union $\mathcal{M} = \bigcup_{n=0}^{\infty} \mathcal{M}_n$. We arrange that the following conditions hold.

(1) Each $\mathcal{M}_n$ is in $\mathcal{K}$.
(2) For every $\mathfrak{A} \in \mathcal{K}$ there is some n such that $\mathfrak{A} \subseteq \mathcal{M}_n$.
(3) For each n, if $f : \mathfrak{A} \longrightarrow \mathfrak{B}$ is an isomorphism with $\mathfrak{A}, \mathfrak{B} \subseteq \mathcal{M}_n$ substructures, and $a \in \mathcal{M}_n$ then there exists $N > n$ and $b \in \mathcal{M}_N$ such that $g : A \cup \{a\} \longrightarrow B \cup \{b\}$ (sending $a \mapsto b$ and $g|_A = f$) is an isomorphism.

Let us first check that if we succeed with the above conditions, then the union $\mathcal{M} = \bigcup_n \mathcal{M}_n$ will indeed we an ultrahomogeneous countable model with $\mathrm{age}(\mathcal{M}) = \mathcal{K}$.

- $\mathcal{M}$ is countable since it is a countable union of finite structures.
- $\mathrm{age}(\mathcal{M}) \subseteq \mathcal{K}$ since if $\mathfrak{A} \in \mathcal{K}$ then there is some n such that $\mathfrak{A} \subseteq \mathcal{M}_n \in \mathcal{K}$ and so $\mathfrak{A}$ is in $\mathcal{K}$ by the Hereditary Property.
- $\mathcal{K} \subseteq \mathrm{age}(\mathcal{M})$ since for $\mathfrak{A} \in \mathcal{K}$ we arrange that $\mathfrak{A} \subseteq \mathcal{M}_n \subseteq \mathcal{M}$.
- Ultrahomogeneity follows from the third condition: Let $\mathcal{M} = \{a_i : i \in \omega\}$. Suppose $f : \mathfrak{A} \longrightarrow \mathfrak{B}$ is an isomorphism between finite substructures of $\mathcal{M}$. We build $f \subseteq f_0 \subseteq f_1 \subseteq \cdots$ a sequences of isomorphisms between finite substructures such that $\mathrm{dom}(f_n) \supseteq \{a_i : i < n\}$ and $\mathrm{range}(f_n) \supseteq \{a_i : i < n\}$. We take $f_0 := f$. At stage l, take n so that $\mathrm{dom}(f_l) \cup \{a_l\} \subseteq \mathcal{M}_n$, then applying condition (3) with $\mathfrak{A} = \mathrm{dom}(f_l)$, $a = a_l$ and $f = f_l$ to obtain $f_{l+\frac{1}{2}} := g$. Then applying condition (3) again to $f^{-1}_{l+\frac{1}{2}}$ to get f_{l+1}. Then we get $\sigma = \bigcup_i f_i \in \mathrm{Aut}(\mathcal{M})$ is the desired extension of f.

So if we succeed in achieving (1), (2) and (3) then we are done. We now show how to do so.

There is some choice in how we index the construction. We will do the construction in ω-steps, but each step will include a finite number of "sub-steps" wherein we ensure that all isomorphisms that need extending (from the current and previous steps) get extended.

List $\mathcal{K} = \{\mathfrak{C}_i : i \in \omega\}$.

Let $\mathcal{M}_0 := \varnothing$.

At stage $n + 1$ we have $\mathcal{M}_n$ constructed. Let

$$\mathcal{I}_n := \{(f, \mathfrak{A}, \mathfrak{B}, a) : \mathfrak{A}, \mathfrak{B} \subseteq \mathcal{M}_n, f : \mathfrak{A} \longrightarrow \mathfrak{B} \text{ an isomorphism } a \in M_n\}$$

be the set of "extension problems" that must be solved before moving on to $\mathcal{M}_{n+1}$. Note that $\mathcal{I}_n$ is a finite set. Let us list it as $\mathcal{I}_n = \{(f_i, \mathfrak{A}_i, \mathfrak{B}_i, a_i) : 0i < J\}$. To build $\mathcal{M}_{n+1}$ we make a sequence $\mathcal{M}_{n,0} \subseteq \mathcal{M}_{n,1} \subseteq \cdots \mathcal{M}_{n,J} \in \mathcal{K}$ such that each step $\mathcal{M}_{n,i}$ takes care of the extension problems $(f_j, \mathfrak{A}_j, \mathfrak{B}_j, a_j)$ for $j < i$.

First we get $\mathcal{M}_{n,0} \in \mathcal{K}$ by the Joint Embedding Property:

$$\begin{tikzcd}
& \mathcal{M}_{n,0} & \\
\mathcal{M}_n \arrow[ur, dashed, hook] & & \mathfrak{C}_n \arrow[ul, dashed, hook']
\end{tikzcd}$$

At stage $n, j+1$ we must solve $(f_j, \mathfrak{A}_j, \mathfrak{B}_j, a_j)$. We use the Amalgamation Property.

First we use AP to get the following diagram in $\mathcal{K}$:

$$\begin{tikzcd}
& \mathfrak{D} & \\
\mathfrak{A}_j \cup \{a_j\} \arrow[ur, dashed, hook, "\iota"] & & \mathfrak{B}_j \arrow[ul, dashed, hook', "j"'] \\
& \mathfrak{A}_j \arrow[ul, hook] \arrow[ur, hook] &
\end{tikzcd}$$

Now we extend f_j to get a_i into the domain, by $f_j(a_i) = \iota(a_j)$. By the commutativity of the diagram this extends f to an isomorphism of $\mathfrak{A}_j \cup \{a_j\} \longrightarrow \mathfrak{B}_j \cup \{b_j\}$ where $b_j := j(a_i)$. By the Amalgamation Property a second time we get $\mathcal{M}_{n,j+1} \in \mathcal{K}$ such that the following diagram commutes:

$$\begin{tikzcd}
& \mathcal{M}_{n,j+1} & \\
\mathcal{M}_{n,j} \arrow[ur, dashed, hook, "\iota"] & & \mathfrak{B}_j \cup \{b_j\} \arrow[ul, dashed, hook', "j"'] \\
\mathfrak{A}_j \cup \{a_j\} \arrow[u, dashed, hook] \arrow[urr, "\sim"] & & \mathfrak{B}_j \arrow[u, dashed, hook] \\
& \mathfrak{A}_j \arrow[ul, hook] \arrow[ur, hook] &
\end{tikzcd}$$

Now letting $\mathcal{M}_{n+1} := \mathcal{M}_{n,J-1}$ we are done. $\square$

Example. (The Random Graph) If $\mathcal{K}$ is the class of finite graphs, then the countable ultrahomogeneous Fraïssé limit, $\mathcal{M}$ of $\mathcal{K}$ is called the random graph. It has the following zero-one property:

For any sentence φ, we have $\mathcal{M} \models \varphi$ if and only if $\lim_{n\to\infty} \Pr_n(\varphi) = 1$ where

$$\Pr_n(\varphi) = \frac{\#\{E \subseteq n \times n | (n, E) \models \varphi\}}{\#\{E \subseteq n \times n\}}$$

is the proportion of the number of graphs with n vertices satisfying φ relative to the total number of graphs with n vertices. In fact this limit always exists, and is always either zero or one.

Lecture 24

Recall that a theory T is said to be κ-categorical if there exists only one model of T (up to isomorphism) of cardinality κ.

Question. Let τ be a countable signature. For T a complete theory, what conditions on T are equivalent to T being $\aleph_0$-categorical?

We shall find several non-obvious conditions that will turn out to be equivalent to T being $\aleph_0$-categorical. For instance all $\aleph_0$-categorical theories can all be realized (possibly after moving to a definitional expansion) as Fraïssé limits. Conversely given a Fraïssé class (i.e. class of finite structures closed under isomorphism and with AP, JEP and HP) in a finite signature, then the limit of this class will be $\aleph_0$-categorical.

A different answer which we will give is one involving the types. For this we must take a detour through the Omitting Types Theorem.

24.1 Omitting Types

Question. Given a countable signature τ and a complete theory T, and a complete type $p \in S_n(\varnothing)$ relative to T. Under what conditions must p be realized in every model of T?

An obvious condition is that T be κ-categorical for all $\kappa \geqslant \aleph_0$. Then all types, $p \in S_n(\varnothing)$, are realized in every model since we have seen that there is always some elementary extension in which p is realized. Using Löwenheim-Skolem to move down again we can ensure that p is realized in the original model.

This is a bit too strong. We shall give a different more syntactic answer.

Definition. A type p is **principal** if there is some $\varphi \in p$ such that for all $\theta \in p$ we have $T \vdash \varphi \longrightarrow \theta$.

Note that if p is principal then the formula φ determines p.

Proposition. *Let* T *be a theory in a countable language and* p *a type. If* p *is principal then for all models* $\mathcal{M}$ *of* T, p *is realized in* $\mathcal{M}$.

Proof. T says that p is consistent. In particular $T \vdash \exists x\varphi$. Thus $\mathcal{M} \models \exists x\varphi$. Let $a \in M^n$ be a witness. Then for any $\theta \in p$ we have $\mathcal{M} \models \varphi \to \theta$ and so $\mathcal{M} \models \theta(a)$, so $\mathcal{M}$ realizes p with a. □

Remark. A type p is principal type if and only if $\{p\}$ is an open subset of $S_n(\varnothing)$, i.e. p is an isolated point.

So if p is principal then it is realized in every model of T. In fact a converse holds as well, as we now prove.

Theorem 24.1. *(Omitting Types) Let $\mathcal{L}(\tau)$ be a countable language. If* T *is a complete τ-theory and $\{p_i\}_{i\in\omega}$ is a sequence of non-principal types, then there exists $\mathcal{M} \models T$ such that $\mathcal{M}$ omits all the p_i's.*

In the proof we go back to the proof of the compactness theorem using the Henkin construction. Except now we must arrange that all the constants that we put in to the domain of our model avoid realizing the types $\{p_i\}$.

Proof. Fix a listing $\{\varphi_l(x)\}_{l\in\omega}$ of the language $\mathcal{L}(\tau_x)$. Let $\tau' \supset \tau$ be an extension of signatures by constants

$$\mathcal{C}_{\tau'} := \mathcal{C}_\tau \cup \{c_i : i < \omega\}.$$

We build a sequence $(T_n)_n$ of extensions of T satisfying

1) $T = T_0 \subseteq T_1 \subseteq T_2 \subseteq \cdots$.
2) Each T_n is a consistent $\mathcal{L}(\tau \cup \{c_j : j < n\})$-theory.
3) For each i, $|T_i \setminus T|$ is finite.
4) For $l, k < i$ there is some $\theta_{l,k} \in p_k$ such that $\neg\theta_{l,k}(c_l) \in T_i$.
5) The constant c_l is a Henkin constant for $\varphi_l(x)$ (in our listing $\{\varphi_l(x)\}_{l\in\omega}$ of the language $\mathcal{L}(\tau_x)$), i.e. $\exists x\varphi_l(x) \to \varphi_l(c_l) \in T_{l+1}$.

We start with $T_0 := T$. At stage $i+1$ we have T_i satisfying the constraints, and must construct T_{i+1}.

- Let $T_{i,0} := T_i \cup \{\exists x\varphi_i \to \varphi_i(c_i)$. This is indeed consistent, being a sub-theory of $T_i \cup T^{Hen}$.
- We now treat condition 4) one by one. List all pairs (l, k) with $l, k < i$. Suppose we fail to satisfy 4). Then it must be the case that

$$T_{i,*} \vdash \theta_{l,k}(c_l)$$

for some $l, k < i$ (where the $*$-symbol indicates the relevant level of the construction). But $T_{i,*} = T \cup \{\eta_1(\bar{c}), \ldots, \eta_t(\bar{c})\}$, so it follows that

$$T \vdash \bigwedge \eta_s(\bar{c}) \longrightarrow \theta_{l,k}(c_l)$$

and so since all the constants $\bar{c}$ are new we get

$$T \vdash \forall x_0, \ldots, x_i \bigwedge \eta_s(\bar{x}) \longrightarrow \theta_{l,k}(x_l).$$

Thus since the $\bar{x}$ actually can be realized we have,

$$T \vdash \forall x_l \left(\exists x_0, \ldots, x_{l-1}, x_{l+1}, \ldots, x_i \bigwedge \eta_s(\bar{x}) \longrightarrow \theta_{l,k}(x_l) \right).$$

But then p_k is isolated by the formula $\exists x_0, \ldots, x_{l-1}, x_{l+1}, \ldots, x_i \bigwedge \eta_s(\bar{x})$ since $T \vdash \exists x_0, \ldots, x_{l-1}, x_{l+1}, \ldots, x_i \bigwedge \eta_s(\bar{x})$. This is a contradiction with the assumption that all the p_i's are non-principal.

- So there is some $\theta_{l,k} \in p_k$ such that $T_{i,*} \nvdash \theta_{l,k}(c_l)$. We define $T_{i,*+1}$ to be $T_{i,*} \cup \{\neg\theta_{l,k}(c_l)\}$.
- We get T_{i+1} as the union $\bigcup_j T_{i,j}$.

Now given the sequence of T_i's satisfying the five requirements we let,

$$\tilde{T} := \bigcup_n T_n.$$

This is a consistent theory with Henkin constants $\{c_i\}_{i\in\omega}$. As done in the proof of the compactness theorem, we can complete this theory, say by an extension $\hat{T} \supseteq \tilde{T}$. Then the Henkin model $\mathcal{M}$ of $\hat{T}$ is also a model of $\tilde{T}$ and every element of $\mathcal{M}$ is an instantiation of $c_i^{\mathcal{M}}$ for some i. By construction none of these elements realize any of the p_k's. The desired model is then $\mathcal{M}|_\tau$. □

Question. There exists a complete theory T in some countable language such that for all $\mathcal{M} \models T$ there is some non-principal type p which is realized in $\mathcal{M}$. Can you find an example of such a theory?

Looking ahead to our study of $\aleph_0$-categoricity we have the following corollaries.

Corollary. *If T is $\aleph_0$-categorical, then every type is principal.*

Proof. Every type is realized. If a type is not principal then we have just seen that we could omit it. □

Corollary. *If T is $\aleph_0$-categorical then for all n the type space $S_n(\varnothing)$ is finite.*

Proof. $S_n(\varnothing)$ is compact. It is covered by the singletons which are open (since they are principal by the above corollary), and so there is some finite sub cover. □

24.2 Atomic and Prime Models

Definition. We say that a model $\mathcal{M}$ is **atomic** if every n-type realized in $\mathcal{M}$ is principal. I.e. for all n and all $a \in M^n$ we have that $\mathrm{tp}(a)$ is principal.

Definition. We say that a model $\mathcal{M}$ is **prime** if for all $\mathcal{N}$ such that $\mathcal{N} \equiv \mathcal{M}$ there is an elementary embedding $\mathcal{M} \longrightarrow \mathcal{N}$.

Example. An ultrahomogeneous model in a finite relational language is both atomic and prime.

Example. The algebraic numbers in the language of fields, are both atomic and prime.

Example. Any uncountable set is atomic. They are however *not* prime (by Löwenheim-Skolem).

Proposition. *If $\mathcal{M}$ is a prime model in a countable language then $\mathcal{M}$ is atomic.*

Proof. If not then there is some $a \in M^n$ such that $p = \mathrm{tp}(a)$ is not principal. The theory $\mathrm{Th}(\mathcal{M})$ is complete. By the omitting types theorem there is a model $\mathcal{N} \models \mathrm{Th}(\mathcal{M})$ such that $\mathcal{N}$ omits p. But since $\mathcal{N} \equiv \mathcal{M}$, there exists an elementary embedding $\iota : \mathcal{M} \longrightarrow \mathcal{N}$, but then $\iota(a)$ realizes p inside $\mathcal{N}$, a contradiction. $\square$

The converse of this proposition is not true as is seen by the example of an uncountable set. However, if the model in question is countable then the converse is true.

Proposition. *If $\mathcal{M}$ is a countable atomic model in a countable language, then $\mathcal{M}$ is prime.*

Proof. Let $\mathcal{N}$ be elementarily equivalent to $\mathcal{M}$. We must find an elementary map $\iota : \mathcal{M} \longrightarrow \mathcal{N}$. The domain M of $\mathcal{M}$ is countable, so let $\{m_i : i < \omega\}$ be a listing of its elements. We will build a sequence $(f_k)_{k\in\omega}$ of partial maps from $\mathcal{M}$ to $\mathcal{N}$ satisfying the following conditions:

1) $f_0 \subseteq f_1 \subseteq f_2 \subseteq \cdots$
2) $\mathrm{dom}(f_i) = \{m_j : j < i\}$
3) $(\mathcal{M}, m_0, \ldots, m_{i-1}) \equiv (\mathcal{N}, f_i(m_0), \ldots, f_i(m_{i-1}))$

Since $\mathcal{M} \equiv \mathcal{N}$ we can take $f_0 = \varnothing$ which does satisfy the requirements. At stage $i+1$ we are given f_i. Now consider the type $p := \mathrm{tp}(m_0, \ldots, m_i)$ given by the $i+1$ first elements of $\mathcal{M}$. Since $\mathcal{M}$ is atomic p is principal, say isolated by $\theta(x_0, \ldots, x_i)$. Also $\mathcal{M} \models \theta(m_0, \ldots, m_i)$ so in particular $\mathcal{M} \models \exists x \theta(m_0, \ldots, m_{i-1}, x)$. By the induction hypothesis f_i is an elementary map, so we have that $\mathcal{N} \models \exists x \theta(f_i(m_0), \ldots, f_i(m_{i-1}), x)$. Let b be a witness, so that $\mathcal{N} \models \theta(f_i(m_0), \ldots, f_i(m_{i-1}), b)$.

We define f_{i+1} on m_i to be b. This will work, since if $\mathcal{M} \models \psi(m_0, \ldots, m_i)$ then $\psi \in p$ and so $\mathcal{M} \models \theta \to \psi$, and so $\mathcal{N} \models \theta \to \psi$ which means that $\mathcal{N} \models \psi(f_{i+1}(m_0), \ldots, f_{i+1}(m_i))$. □

Question. Can you show that a prime model in a countable language is unique (up to isomorphism)?

When do atomic models exist? The next proposition gives a sufficient condition.

Proposition. *If* T *is a complete theory in a countable language such that for all* n *the type space* $S_n(\varnothing)$ *is at most countable, then* T *has an atomic model.*

Proof. List all the non-principal types (of which there can be at most countable many). By the Omitting Types theorem we can find a model, $\mathcal{M}$, of T which omits all these types. Then $\mathcal{M}$ is atomic. □

Remark. It may still happen that T has an atomic model, even if $S_n(\varnothing)$ has size $2^{\aleph_0}$. For instance the structure $(\mathbb{Q}, <, \{q\}_{q \in \mathbb{Q}})$ has $2^{\aleph_0}$-many 1-types and yet is atomic. In general, given any countable structure it becomes atomic if we name all elements.

Definition. Let $\mathfrak{A}$ and $\mathfrak{B}$ be τ-structures and $X \subseteq A$, $Y \subseteq B$. We say that a map $f : X \to Y$ is **elementary** if, for all $\bar{a}$ from X and all formulae φ we have

$$\mathfrak{A} \models \varphi(\bar{a}) \qquad \Leftrightarrow \qquad \mathfrak{B} \models \varphi(f\bar{a}).$$

Theorem 24.2. *Let* $\mathfrak{A}$ *and* $\mathfrak{B}$ *be countable structures in any (possibly uncountable) language* $\mathcal{L}(\tau)$. *If* $\mathfrak{A} \equiv \mathfrak{B}$ *and both are atomic, then* $\mathfrak{A} \cong \mathfrak{B}$.

Proof. We use a Back-and-Forth Argument. List $\mathfrak{A}$ and $\mathfrak{B}$ as $\{a_n : n \in \omega$ and $\{b_n : n \in \omega\}$, respectively. We build, by recursion, a sequence of maps $(f_n)_{n \in \omega}$ such that

1) $f_0 \subseteq f_1 \subseteq f_2 \subseteq \cdots$.
2) $\mathrm{dom}(f_n) \supseteq \{a_m : m < n\}$.
3) $\mathrm{range}(f_n) \supseteq \{b_m : m < n\}$.
4) Each f_n is elementary.
5) Each f_n is finite.

Since $\mathfrak{A} \equiv \mathfrak{B}$ we may let $f_0 = \varnothing$.

At stage $n+1$ we are given f_n and must construct f_{n+1}. We wish to include a_n in the domain of f_{n+1}. List the domain of f_n as $\{c_1, \ldots, c_l\}$. By assumption, the type $p := \mathrm{tp}(c_1, \ldots, c_l, a_n)$ is principal, say isolated by θ. Then $\theta(c_1, \ldots, c_l, x)$ will isolate the 1-type $\mathrm{tp}(a_n/\bar{c})$. So

$$(\mathfrak{A}, \bar{c}) \models \exists x \theta(\bar{c}, x).$$

By our induction hypothesis on f_n, this implies that

$$\mathfrak{B} \models \exists x \theta(f_n(\bar{c}), x).$$

Let us witness the existential by some $b \in B$. We then define $f_{n+1}(a_n) := b$. Now for any formula $\varphi(x_1, \ldots, x_l, y)$ we have

$$\begin{aligned} \mathfrak{A} \models \varphi(c_1, \ldots, c_l, a_n) &\Leftrightarrow \varphi \in \mathrm{tp}(c_1, \ldots, c_l, a_n) \\ &\Leftrightarrow \mathfrak{A} \models \theta \rightarrow \varphi \\ &\Leftrightarrow \mathfrak{B} \models \theta \rightarrow \varphi \\ &\Rightarrow \mathfrak{B} \models \varphi(f_{n+1}(\bar{c}), f_{n+1}(a_n)). \end{aligned}$$

Now reversing the rôles of $\mathfrak{A}$ and $\mathfrak{B}$ we can ensure that b_n comes into the domain of f_{n+1}^{-1}, i.e. into the range of f_{n+1}.

Letting $f := \bigcup_n f_n$ we have that $f : \mathfrak{A} \longrightarrow \mathfrak{B}$ is an isomorphism, since $\mathfrak{A}$ and $\mathfrak{B}$ are both countable. $\square$

In particular is T is a complete theory such that all countable models are atomic, then T is $\aleph_0$-categorical.

Remark. If we drop the assumption that $\mathfrak{A}$ and $\mathfrak{B}$ are countable, then the above back-and-forth argument gives us that $\mathfrak{A} \equiv_{\infty,\omega} \mathfrak{B}$.

Question. What would it mean about a theory T is it has the "Shroeder-Bernstein property": given $\mathfrak{A}$ and $\mathfrak{B}$ with elementary embeddings in both directions, then $\mathfrak{A}$ and $\mathfrak{B}$ are in fact isomorphic?

We also note the following useful lemma.

Lemma. *If $\mathfrak{A}$ is atomic and $\bar{b}$ is a tuple from A then $(\mathfrak{A}, \bar{b})$ is also atomic.*

Proof. If $\bar{c}$ is another tuple (possibly of different length) then $\bar{b}^\frown\bar{c}$ must realize a principal type, isolated by some formula $\theta(\bar{x}; \bar{y})$. But then $\theta(\bar{b}; y)$ will isolate the type $\mathrm{tp}(\bar{c}/\bar{b})$. So $(\mathfrak{A}, \bar{b})$ is atomic. $\square$

Lecture 25

25.1 Countable Categoricity

We have developed enough tools to give a proof of a list of characterisations of $\aleph_0$-categoricty. The theorem is due to Ryll-Nardzewski, Engler and Svenonius.

Definition. For $\mathfrak{A}$ a τ-structure, $\mathrm{Aut}(\mathfrak{A})$ acts naturally on A^n where A is the domain of $\mathfrak{A}$. $\mathrm{Aut}(\mathfrak{A})$ is said to **act oligomorphically** if, for all n there are only finitely many orbits of the action of $\mathrm{Aut}(\mathfrak{A})$ on A^n.

Theorem 25.1. *(Ryll-Nardzewski, Engler and Svenonius.) Let $\mathcal{L}(\tau)$ be a countable language and T a complete τ-theory which has infinite models. The following are equivalent.*

1) T *is $\aleph_0$-categorical.*
2) *Every model of* T *is atomic.*
3) *Every countable model of* T *is atomic.*
4) *For all countable models $\mathfrak{A}$ of* T, $\mathrm{Aut}(\mathfrak{A})$ *acts oligomorphically*
5) *There exists a model $\mathfrak{B}$ of* T *such that* $\mathrm{Aut}(\mathfrak{B})$ *acts oligomorphically.*
6) *For each n there are only finitely many n-types.*

Proof. We have already proved several of the implications.

1) $\Rightarrow$ 2): Proved last time as a corollary to the omitting types theorem.

2) $\Rightarrow$ 3): Trivial.

1) $\Rightarrow$ 6): Also proved last time as a corollary to the omitting types theorem.

3) $\Rightarrow$ 1): Proved last time as well (see theorem at the end of lecture 24).

3) $\Rightarrow$ 4): Let $\mathfrak{A}$ be a countable model of T. We must show that $\mathrm{Aut}(\mathfrak{A})$ acts oligomorphically. By 3) we know that $\mathfrak{A}$ is atomic. By 3) $\Rightarrow$ 1) $\Rightarrow$ 6) we know that $S_n(\varnothing)$ is finite for each n. So there is a finite list $\theta_1, \dots, \theta_l$ such that

$$A^n = \theta_1(\mathfrak{A}) \cup \cdots \cup \theta_l(\mathfrak{A})$$

is a disjoint union, and such that each θ_i is a complete formula (i.e. cannot be split into proper definable subsets). Let $a, b \in A^n$ be such that $\mathfrak{A} \models \theta_i(a)$ and $\mathfrak{A} \models \theta_i(b)$ for some i. Then $(\mathfrak{A}, a) \equiv (\mathfrak{A}, b)$ since θ_i isolates a type. But $\mathfrak{A}$ is atomic.

By the lemma from the end of lecture 24 we know that both $(\mathfrak{A}, a)$ and $(\mathfrak{A}, b)$ are atomic. They are also countable and so by the last theorem of lecture 24, $(\mathfrak{A}, a)$ and $(\mathfrak{A}, b)$ are isomorphic. I.e. there is some $\sigma \in \mathrm{Aut}(\mathfrak{A})$ such that $\sigma(a) = b$. Thus $|\mathrm{Aut}(\mathfrak{A})/A^n| = l$, i.e. there are only l orbits.

4) $\Rightarrow$ 5): Straight forward since we have assumed that T has infinite models.

5) $\Rightarrow$ 6): Let $\mathfrak{B}$ be a model of T such that $\mathrm{Aut}(\mathfrak{B})$ acts oligomorphically. Let $n \in \omega$. We wish to show that $S_n(T)$ is finite. If $a, b \in B^n$ and if there is $\sigma \in \mathrm{Aut}(\mathfrak{B})$ such that $\sigma(a) = b$, then $\mathrm{tp}(a) = \mathrm{tp}(b)$. Thus, since $\mathrm{Aut}(\mathfrak{B})/B^n$ is finite, the number of *realized* types is finite. For each orbit $X \subseteq B^n$ of $\mathrm{Aut}(\mathfrak{B})$ there is some formula θ_X such that θ_X isolates the type of an element of X. This is because the number of types is bounded by the number of orbits. Now

$$\mathfrak{B} \models \forall x_1, \ldots, x_n \left(\bigvee_{X \text{ an orbit}} \theta_X(\bar{x}) \right)$$

and for each (X, φ) we have

$$\mathfrak{B} \models \forall \bar{x}\, (\theta_X \longrightarrow \varphi)$$

or

$$\mathfrak{B} \models \forall \bar{x}\, (\theta_X \longrightarrow \neg\varphi)$$

Since T is complete we have that $\theta_X \to \varphi$ or $\theta_X \to \neg\varphi$ belongs to T. Thus the number of n-types is finite.

6) $\Rightarrow$ 1): We proved 6) $\Rightarrow$ 2) in lecture 24. Now we have seen above that 2) $\Rightarrow$ 3) $\Rightarrow$ 1). □

Remark. There are further equivalent conditions. For example,

7) Every countable model is saturated (see next lecture for definition)
8) Every countable model is bi-definable[1] with a Fraïssé limit
9) Every countable model is bi-interpretable with a Fraïssé limit.

Proposition. *If $\mathcal{K}$ is a class of finite structures in a finite relational signature, with AP, JEP and HP, and $\mathfrak{D}$ is its Fraïssé limit, then* $\mathrm{Th}(\mathfrak{D})$ *is $\aleph_0$-categorical.*

Remark. We have proved that there is a unique *ultrahomogeneous* countable model whose age is $\mathcal{K}$. The proposition states that in fact the theory of the countable limit is already enough to get uniqueness.

The proof goes along the following lines. We basically just *say* what the theory $\mathrm{Th}(\mathfrak{D})$ is, i.e. we axiomatize $\mathrm{Th}(\mathfrak{D})$. We must say that the age of a model of $\mathrm{Th}(\mathfrak{D})$ is contained in $\mathcal{K}$, and that the reverse inclusion holds as well.

[1]Two structures are bi-definable if they have the same class of definable sets.

Proof. We construct a theory T which we claim is equal to $\mathrm{Th}(\mathfrak{D})$. For each n let φ_n be the sentence

$$\forall x_1, \ldots, x_n \bigvee_{\substack{\mathfrak{B} \in \mathcal{K} \\ |B|=n}} \bigwedge \mathrm{Diag}(\mathfrak{B})(\bar{x}).$$

Each φ_n is actually a first-order formula since the language is finite relational so there are only finitely many $\mathfrak{B} \in \mathcal{K}$ of size equal to n, and each diagram is also finite. If $\mathcal{M}$ models all the φ_n's then the age of $\mathcal{M}$ must be contained in $\mathcal{K}$.

The second axiom schema will force the age of a model of T to contain $\mathcal{K}$, swell as making such models ultrahomogeneous. For each pair $\mathfrak{A} \subseteq \mathfrak{B} \in \mathcal{K}$ we define $\psi_{\mathfrak{A},\mathfrak{B}}$ as follows: list $A = \{a_1, \ldots, a_n\}$ and $B = \{a_1, \ldots, a_n, b_1, \ldots, b_m\}$ and let $\psi_{\mathfrak{A},\mathfrak{B}}$ be

$$\forall x_1, \ldots, x_n \left[\bigwedge \mathrm{Diag}(\mathfrak{A})(\bar{x}) \longrightarrow \exists y_1, \ldots, y_m \bigwedge \mathrm{Diag}(\mathfrak{B})(\bar{x}, \bar{y})\right].$$

The sentence $\psi_{\mathfrak{A},\mathfrak{B}}$ expresses that whenever one has a copy of $\mathfrak{A}$ it can be extended to a copy of $\mathfrak{B}$.

Now we let T be given by $\{\varphi_n : n \in \omega\} \cup \{\psi_{\mathfrak{A},\mathfrak{B}} : \mathfrak{A} \subseteq \mathfrak{B} \in \mathcal{K}\}$. We claim that T axiomatizes $\mathrm{Th}(\mathfrak{D})$.

First we check that $\mathfrak{D} \models T$; Since $\mathrm{age}(\mathfrak{D}) = \mathcal{K}$, each n-element substructure of $\mathfrak{D}$ is isomorphic to some $\mathfrak{B} \in \mathcal{K}$, so $\mathfrak{D} \models \varphi_n$. Furthermore the ultrahomogeneity will show that $\mathfrak{D}$ models $\psi_{\mathfrak{A},\mathfrak{B}}$ for each $\mathfrak{A} \subseteq \mathfrak{B} \in \mathcal{K}$. More precisely, suppose we have $\mathfrak{A}' \subseteq \mathfrak{D}$ with $\mathfrak{A}' \cong \mathfrak{A} \in \mathcal{K}$, and suppose there is some extension $\mathfrak{A} \subseteq \mathfrak{B} \in \mathcal{K}$. Since $\mathrm{age}(\mathfrak{D}) = \mathcal{K}$, there is some $\mathfrak{B}'' \subseteq \mathfrak{D}$ which is isomorphic to $\mathfrak{B}$. Then using this isomorphism to pull back $\mathfrak{A}$ we get some $\mathfrak{A}'' \subseteq \mathfrak{B}'' \subseteq \mathfrak{D}$ with $\mathfrak{A}'' \cong \mathfrak{A}$. By ultrahomogeneity there is some $\sigma \in \mathrm{Aut}(\mathfrak{D})$ such that σ restricted to $\mathfrak{A}''$ agrees with the isomorphism $\mathfrak{A}'' \cong \mathfrak{A}$. Now we set $\mathfrak{B}' = \sigma(\mathfrak{B}'')$. Then $\mathfrak{D} \models \psi_{\mathfrak{A},\mathfrak{B}}$. This proves that $\mathrm{Mod}(\mathrm{Th}(\mathfrak{D})) \subseteq \mathrm{Mod}(T)$.

Now suppose that $\mathfrak{D}' \models T$ and that $\mathfrak{D}'$ is countable. Then $\mathrm{age}(\mathfrak{D}') \subseteq \mathcal{K}$ by the axiom schema φ_n. Further, taking $\mathfrak{A} = \varnothing$ (which is in $\mathcal{K}$ since the language is relational), we see that since $\mathfrak{D}' \models \psi_{\mathfrak{A},\mathfrak{B}}$ for all $\mathfrak{B}$ it follows that $\mathcal{K} \subseteq \mathrm{age}(\mathfrak{D}')$. So $\mathrm{age}(\mathfrak{D}') = \mathcal{K}$. Finally we must show that $\mathfrak{D}'$ is ultrahomogeneous in order to get that $\mathfrak{D}' \cong \mathfrak{D}$.

Suppose $f : \mathfrak{A} \longrightarrow \mathfrak{A}'$ is an isomorphism for some finite substructures $\mathfrak{A}, \mathfrak{A}' \subseteq \mathfrak{D}'$. By the back-and-forth argument it is sufficient to show that we can extend f by just one element. Since $\mathrm{age}(\mathfrak{D}') = \mathcal{K}$ we know that $\mathfrak{A}$ and $\mathfrak{A}'$ are isomorphic to some $\mathfrak{A}_0$ and $\mathfrak{A}_0'$ in $\mathcal{K}$. Taking $\mathfrak{B} \in \mathcal{K}$ so that $\mathfrak{A}_0 \subseteq \mathfrak{B}$ and using the fact that $\mathfrak{D}' \models \psi_{\mathfrak{A}_0,\mathfrak{B}}$ it follows that there is some $\mathfrak{B}' \subseteq \mathfrak{D}'$ and an extension of f to $\mathfrak{B}'$.

Thus $\mathfrak{D}'$ is ultrahomogeneous, countable and has the same age as $\mathfrak{D}$, so by Fraïssé's Theorem, $\mathfrak{D}'$ is isomorphic to $\mathfrak{D}$. Thus, $\mathrm{Th}(\mathfrak{D})$ is $\aleph_0$-categorical. $\square$

In the coming lecture we will discuss various notions of saturation. In anticipation of this we prove the following proposition.

Proposition. *Let* T *be an* $\aleph_0$*-categorical theory,* $\mathfrak{A}$ *is a model of* T *and* $B \subseteq A$ *is finite. If* p *is a* 1*-type over* B *then* p *is realized in* $\mathfrak{A}$.

Proof. By the Theorem we know that $\mathfrak{A}$ must be atomic. By the lemma from lecture 24 it follows that $(\mathfrak{A}, B)$ is also atomic. Thus $\mathrm{Th}(\mathfrak{A}_B)$ is atomic, so every type is principal, hence is realized. □

Remark. In terminology to be defined later, this proposition states that $\mathfrak{A}$ is $\aleph_0$-*saturated*.

Lecture 26

26.1 Saturation

Definition. Let λ be an infinite cardinal and $\mathfrak{A}$ and τ-structure.

1) $\mathfrak{A}$ is λ**-compact** if, for each $\Sigma(x) \subseteq \mathcal{L}(\tau_{A,x})$ with $|\Sigma| < \lambda$ and where Σ is finitely satisfiable, then Σ is realized in $\mathfrak{A}$.
2) $\mathfrak{A}$ is λ**-saturated** if, for all $B \subseteq A$ with $|B| < \lambda$ then for all 1-types p over B, then p is realized.
3) $\mathfrak{A}$ is λ**-universal** if, for all $\mathfrak{B}$ such that $\mathfrak{B} \equiv \mathfrak{A}$ and $|\mathfrak{B}| \leqslant \lambda$ then there exists an elementary embedding $\iota : \mathfrak{B} \longrightarrow \mathfrak{A}$.
4) $\mathfrak{A}$ is λ**-homogeneous** if, whenever we have $B, C \subseteq A$ such that $\mathfrak{A}_B \equiv \mathfrak{A}_C$, and given $b \in A$ then there is some $c \in A$ such that $\mathfrak{A}_{B,b} \equiv \mathfrak{A}_{C,c}$.
5) $\mathfrak{A}$ is λ**-strongly homogeneous** if, for all $B, C, \subseteq A$ with $|B| < \lambda$ and $f : B \longrightarrow C$ and elementary map, then there exists $\sigma \in \mathrm{Aut}(\mathfrak{A})$ such that $\sigma|_B = f$.

Remark. Note that in the definition of λ-universality we allow that the cardinality of $\mathfrak{B}$ is *equal* to λ.

Remark. In the definition of λ-strongly homogeneous the map $f : B \longrightarrow C$ is a bijection, so in particular $|C| < \lambda$.

Notation. If $\mathcal{M}$ is λ-homogeneous but not necessarily λ-strongly homogenous then we can stress this by saying that $\mathcal{M}$ is λ-weakly-homogeneous.

Definition. Given a 1-type $p(x)$ over A and a map $f : A \longrightarrow B$ we define the **push-forward of** p **by** f to be

$$f_*(p) := \{\varphi(x, f(\bar{a})) : \varphi(x, \bar{a}) \in p\}.$$

Theorem 26.1. *Let τ be a signature and λ a cardinal $\geqslant |\mathcal{L}(\tau)|$. Let $\mathcal{M}$ be a τ-structure. The following are equivalent.*

1) $\mathcal{M}$ *is λ-saturated.*
2) $\mathcal{M}$ *is λ-homogeneous and λ-universal.*

Proof. 1) $\Longrightarrow$ 2): Suppose first that $\mathcal{M}$ is λ-saturated.

- First we show that $\mathcal{M}$ is λ-universal: Let $\mathcal{N}$ be a τ-structure of cardinality $< \lambda$ and such that $\mathcal{N} \equiv \mathcal{M}$. We must find an elementary embedding of $\mathcal{N}$ into $\mathcal{M}$. Let $N = \{b_\alpha : \alpha < \lambda\}$ (possibly using repetitions). We construct a sequence of maps (f_α) such that $\mathrm{dom}(f_\alpha) = \langle b_\beta : \beta < \alpha \rangle$ and such that $f_\alpha : N \to M$ is a partial elementary map. Furthermore we will ensure that $f_\alpha \subseteq f_\beta$ for $\alpha < \beta$.

  At stage 0 we simply let $f_0 = \varnothing$. This works since $\mathcal{M} \equiv \mathcal{N}$.

  At stage $\alpha + 1$ we have defined f_α and must put b_α into the domain of $f_{\alpha+1}$. Consider $q = \mathrm{tp}(a_\alpha/\mathrm{dom}(f_\alpha))$. Set

  $$p := (f_\alpha)_*(q) = \{\varphi(x, f_\alpha(\bar{c})) \mid \varphi(x, \bar{c}) \in q\}.$$

  p is a consistent type since f_α is elementary. For any finite subset of q, $\varphi_1(x, c^{(1)}), \ldots, \varphi_n(x_n, c^{(n)})$ we have

  $$\mathcal{N} \models \exists x \bigwedge_i \varphi_i(x, c^{(i)})$$

  since q is by definition realized. Thus,

  $$\mathcal{M} \models \exists x \bigwedge_i \varphi_i(x, f_\alpha(c^{(i)}))$$

  since f_α is (by the induction hypothesis) elementary. Thus p is a consistent type over $\mathrm{range}(f_\alpha)$. I.e. $p \in S_1(\mathrm{range}(f_\alpha))$. Now since $|\mathrm{range}(f_\alpha)| < \lambda$ we have, by λ-saturation of $\mathcal{M}$, that p is realized. So there is some $c \in \mathcal{M}$ such that $p = \mathrm{tp}(c/\mathrm{range}(f_\alpha))$. We define $f_{\alpha+1}(b_\alpha) := c$.

  Finally for $\alpha = \lambda$ a limit ordinal we set $f_\lambda = \bigcup_{\alpha<\lambda} f_\alpha$. Now we let $f := f_\lambda$. By construction we have that $f : \mathcal{N} \longrightarrow \mathcal{M}$ is an elementary embedding.

- Now we show that $\mathcal{M}$ is λ-homogeneous. So let $A, B \subseteq M$ be given, such that $|A| = |B| < \lambda$ and such that there is a bijection $f : A \longrightarrow B$ which is partial elementary, i.e.

  $$(\mathcal{M}, \{a\}_{a\in A}) \equiv (\mathcal{M}, \{f(a)\}_{a\in A}).$$

  We must show that given $c \in \mathcal{M}$ we can extend f to c while still being elementary.

  Consider the 1-type $p := \mathrm{tp}(c/A)$ of c over A. Define q to be the push-forward p by f, i.e.

  $$q := f_*(p).$$

We claim that q is a type over B. Consistency: Let $\Phi \subseteq q$ be a finite subset of q, say

$$\Phi = \{\varphi_1(x, f(a_1)), \dots, \varphi_n(x, f(a_n))\}$$

(where the a_i's are actually tuples). By definition of q there is a corresponding set

$$\{\varphi_1(x, a_1), \dots, \varphi_n(x, a_n)\} \subseteq p$$

which is satisfied by some element. I.e.

$$\mathcal{M} \models \exists x \bigwedge \varphi_i(x, a_i).$$

Since f is elementary we have

$$\mathcal{M} \models \exists x \bigwedge \varphi_i(x, f(a_i)).$$

which is what we wanted. Thus q is consistent by compactness. Furthermore to show that q is a type we must show that it is complete: Let $\psi \in \mathcal{L}(\tau_{B,x})$. We may write ψ as $\tilde{\psi}(x; \bar{b})$ where $\psi(x; \bar{y})$ is a formula without parameters from B. Now since p is complete either $\tilde{\psi}(x, f^{-1}(b))$ or $\neg\tilde{\psi}(x, f^{-1}(b))$ is in p. Pushing forward by f we see that either ψ or $\neg\psi$ is in q. Thus q is complete and so is indeed a type.

So q is a 1-type over B. Since $|B| < \lambda$ and $\mathcal{M}$ is λ-saturated, it follows that q is realized, i.e. there is some $d \in M$ such that $q = \mathrm{tp}(d/B)$. Extending f to c by $f(c) := d$ will thus work.

2) $\Longrightarrow$ 1): We now assume that $\mathcal{M}$ is λ-homogeneous and λ-universal. Let p be a 1-type over A for some subset A of M with $|A| < \lambda$.

So by definition p is a complete and consistent theory in $\mathcal{L}(\tau_{A,x})$ extending $\mathrm{Th}(\mathcal{M}_A)$. Note that

$$|p| \leqslant |\mathcal{L}(\tau_{A,x})| \leqslant \lambda$$

since we assumed $|\mathcal{L}(\tau)| \leqslant \lambda$. By Löwenheim-Skolem there is some $(\mathcal{N}, b) \models p$ such that $|N| \leqslant \lambda$ and b realizes p. Consider the reduct of $\mathcal{N}$ back to τ. We have

$$\mathcal{N}|_\tau \models \mathrm{Th}(\mathcal{M})$$

since $\mathcal{N} \models p$. So $\mathcal{N}|_\tau \equiv \mathcal{M}$. By λ-universality of $\mathcal{M}$, there exists an elementary embedding

$$\iota : \mathcal{N}|_\tau \longrightarrow \mathcal{M}$$

Define

$$B := \{\iota(a^{\mathcal{N}}) : a \in A\} \subseteq M.$$

Let $g : B \longrightarrow A$ be given by $\iota(a^{\mathcal{N}}) \longmapsto a^{\mathcal{N}}$. Then g is a partial elementary bijection since ι is an elementary embedding.

Let $c := \iota(x^{\mathcal{N}}) = \iota(b)$. By λ-homogeneity applied to $g : B \longrightarrow \mathfrak{A}$ and to the element c, there is an element $d \in M$ such that g extends to $B \cup \{c\}$ via. $f(c) = d$.

Now we claim that since g is elementary, d realizes p: suppose $\varphi(x) \in p$. Write φ as $\tilde{\varphi}(x; \bar{a})$ where $\tilde{\varphi}(x, \bar{y})$ does not use parameters from A. Now $\mathcal{N} \models \tilde{\varphi}(b, \bar{a}^{\mathcal{N}})$ and so

$$\mathcal{M} \models \tilde{\varphi}(c, g^{-1}(\bar{a}))$$

so applying g we see that

$$\mathcal{M} \models \tilde{\varphi}(d, \bar{a})$$

so $\mathcal{M} \models \varphi(d)$. □

Remark. Hodges' defines λ-universality using a strict inequality. In this case the above statement should be modified to "λ^+-universal".

The theorem shows that saturation is a very strong property. If $\mathcal{M}$ is λ-saturated for $\lambda = |\mathcal{M}|$ then $\mathcal{M}$ is $|\mathcal{M}|$-homogeneous which in fact (see proposition below for proof) implies that $\mathcal{M}$ is $|\mathcal{M}|$-*strongly* homogeneous since we can carry out the back-and-forth argument all the way up to $|\mathcal{M}|$. So the existence of models that are saturated in their own cardinality is definitely to be desired!

Definition. A model $\mathcal{M}$ is said to be **saturated** if $\mathcal{M}$ is $(|\mathcal{M}| + |\mathcal{L}(\tau)|)$-saturated.

Remark. Note that we do not pursue λ-saturation for $\lambda > |\mathcal{M}|$ since if $\mathcal{M}$ is infinite then it is never $|\mathcal{M}|^+$-saturated: In this case we can take the partial type

$$\Sigma(x) := \{\text{“}x \neq m\text{"} : m \in M\}.$$

Σ is consistent by compactness, but there can be no complete extension of Σ which is realized in $\mathcal{M}$.

By definition we have the following corollary.

Corollary. *If $\mathcal{M}$ is saturated then $\mathcal{M}$ is $|\mathcal{M}|$-universal and $\mathcal{M}$ is $|\mathcal{M}|$-homogeneous.*

Proposition. *If $\mathcal{M}$ is $|\mathcal{M}|$-weakly-homogeneous then $\mathcal{M}$ is $|\mathcal{M}|$-strongly-homogeneous.*

Proof. We start with some $f : A \longrightarrow B$ for some $A, B \subseteq M$ and $|A| = |B| < |\mathcal{M}|$, and assume that f is a partial elementary map. We must find some $\sigma \in \mathrm{Aut}(\mathcal{M})$ that extends f.

We use a back-and-forth argument. List $M = (m_\alpha : \alpha < |\mathcal{M}|)$. We build a chain $(\sigma_\alpha)_{\alpha < |\mathcal{M}|}$ such that

1) σ_α is a partial elementary map.

2) $|\sigma_\alpha| < |\alpha| + |\mathcal{L}(\tau_A)|$
3) For all α, σ_α extends f.
4) For $\alpha < \beta$ we have $\sigma_\alpha \subseteq \sigma_\beta$
5) For all $\alpha < |\mathcal{M}|$ we have $\mathrm{dom}(\sigma_\alpha) \supseteq \{m_\beta : \beta < \alpha\}$ and $\mathrm{range}(\sigma_\alpha) \supseteq \{m_\beta : \beta < \alpha\}$.

- For $\alpha = 0$, define $\sigma_0 := f$. Since f is partial elementary this satisfies the conditions.
- For α a limit we let $\sigma_\alpha = \bigcup_{\beta<\alpha} \sigma_\beta$. This works since at each step we will in fact only add 2 elements to $\sigma_{\beta+1}$.
- For $\alpha + 1$ we need a bit more work. By (weak)-homogeneity applied to σ_α and m_α, there is a map $g \supseteq \sigma_\alpha$ such that $g(m_\alpha)$ is defined and such that g is still a partial elementary map. Similarly by weak-homogeneity applies to g^{-1} and m_α we can extend g to some h which is defined on m_α. Thus letting $\sigma_{\alpha+1}(m_\alpha) = g(m_\alpha)$ and $\sigma_{\alpha+1}(h(m_\alpha)) = m_\alpha$ so $\sigma_{\alpha+1}$ works.

Then the union $\sigma = \bigcup_\alpha \sigma_\alpha$ will be the required automorphism. □

Corollary. *Suppose $\mathcal{M}$ if saturated. If α is some ordinal such that $\alpha < |\mathcal{M}|$ and $a, b \in M^\alpha$ then $\mathrm{tp}(a) = \mathrm{tp}(b)$ if and only if there is some automorphism $\sigma \in \mathrm{Aut}(\mathcal{M})$ such that $\sigma(a) = b$.*

Proof. Since $\mathrm{tp}(a) = \mathrm{tp}(b)$ the map $f : a_\beta \longmapsto b_\beta$ is a partial elementary bijection. By strong homogeneity f extends to an automorphism of $\mathcal{M}$. □

So in a saturated model, two small tuples "look the same" if and only if there is some automorphism that exchanges them.

Notation. We defined n-types. This can be extended to arbitrary ordinals. Let α be an ordinal. Then an α-type is a complete extension of $\mathrm{Th}(\mathcal{M})$ in variable $\bar{x}$ where $\bar{x}$ has length $\leqslant \alpha$.

Corollary. *Let $\mathcal{M}$ be saturated and let $\alpha < |\mathcal{M}|$. There is a map*

$$S_\alpha(\mathrm{Th}(\mathcal{M})) \longrightarrow \mathrm{Aut}(\mathcal{M})/M^\alpha$$

taking p to the class $[a]$ where a realizes p. This is a bijection.

Proof. We have already proven that $\bar{a}$ and $\bar{b}$ (both α-tuples) are in the same $\mathrm{Aut}(\mathcal{M})$-orbit if and only they have the same type. Thus *if* the map is defined it is a bisection.

But to see that the map is in fact defined we must note that every type in $\leqslant |\mathcal{M}|$ variables is actually realized in $\mathcal{M}$. This follows from the fact that $\mathcal{M}$ is $|\mathcal{M}|$-universal (and using Downward Löwneheim-Skolem). □

26.2 Existence of Saturated Models

We have now established that saturated models are wonderful! But do they exist?

In general, assuming that ZFC is consistent, it is not provable that it is consistent to assume that they exist. In fact the statement that "every theory has a saturated model" is consistency-equivalent to the statement that "strongly inaccessible cardinals exists".

Proposition. *Given any τ-structure $\mathcal{M}$ and cardinal $\lambda \geqslant |\mathcal{L}(\tau)|$ then there exists an elementary extension $\mathcal{N}$ of $\mathcal{M}$ such that $\mathcal{N}$ is λ-saturated.*

Proof. Without loss of generality we may assume that λ is *regular*. If λ is not regular then we carry out the construction for λ^+ which is regular. If $\mathcal{N}$ is λ^+-saturated then it is also λ-saturated.

We will build by recursion on α, a chain $(\mathcal{M}_\alpha)_{\alpha<\lambda}$ such that

$$\mathcal{M} = \mathcal{M}_0 \preccurlyeq \mathcal{M}_1 \preccurlyeq \cdots$$

and such that every type in $S_1(M_\alpha)$ is realized in $\mathcal{M}_{\alpha+1}$.

- We start with $\mathcal{M}_0 := \mathcal{M}$.
- If $\mu < \lambda$ is a limit, then we define $\mathcal{M}_\mu := \bigcup_{\alpha<\mu} \mathcal{M}_\alpha$. By the Tarski-Vaught Theorem on Unions of Elementary Chains we know that $\mathcal{M}_\alpha \preccurlyeq \mathcal{M}_\mu$ for all $\alpha < \mu$.
- At stage $\alpha + 1$ we have built $\mathcal{M}_\alpha$. Consider the theory

$$T := \mathrm{Th}((\mathcal{M}_\alpha)_{M_\alpha}) \cup \bigcup_{p \in S_1(M_\alpha)} p(x_p)$$

in the language $\mathcal{L}(\tau_{M_\alpha,\{x_p : p\in S_1(M_\alpha)\}})$. To see that T is consistent let $T_0 \subseteq T$ be finite. Say

$$T_0 \subseteq \mathrm{Th}((\mathcal{M}_\alpha)_{M_\alpha}) \cup p_0^1(x_{p^1}) \cup \cdots \cup p_0^n(x_{p^n})$$

where $p_0^i(x^{p^i})$ is a finite subsets of the type p^i. For each finite $p_0^i \subseteq p^i$ we know that it can be satisfied. By a Joint Consistency argument, we can satisfy the p^i's all together to finally satisfy T_0. By compactness T is consistent. Let $\mathcal{N}$ be a model of T. Now $\mathcal{M}_\alpha$ can be elementarily embedded into $\mathcal{N}|_\tau$.

Remark. We may assume that

$$|\mathcal{N}| \leqslant |\mathcal{L}(\tau)| + |M_\alpha| + |S_1(M_\alpha)| \leqslant |\mathcal{L}(\tau)| + 2^{|M_\alpha|}.$$

Now let $\widetilde{\mathcal{M}} := \mathcal{M}_\lambda$ be the union of the chain. Then $\mathcal{M} \preccurlyeq \widetilde{\mathcal{M}}$. Let $A \subseteq \widetilde{M}$ have cardinality $< \lambda$ and $p \in S_1(A)$. Since λ is regular, there is some $\alpha < \lambda$ such that $A \subseteq M_\alpha$. Let $q \in S_1(M_\alpha)$ be an extension of p. By construction x_q realizes q is $\mathcal{M}_{\alpha+1}$. So the x_q also realizes p in $\widetilde{\mathcal{M}}$. Thus $\widetilde{\mathcal{M}}$ is λ-saturated.

□

Taking care of the cardinals in the above construction one can get a bit more.

Definition. A regular cardinal λ is **strongly inaccessible** if $\lambda > \aleph_0$, and for every $\mu < \lambda$ one has $2^\mu < \lambda$.

Notation. For any cardinal μ we define the $\beth_{()}$ operation:

- $\beth_0(\mu) := \mu$.
- $\beth_{\alpha+1}(\mu) := 2^{\beth_\alpha(\mu)}$.
- For δ a limit let $\beth_\delta(\mu) := \bigcup_{\alpha<\delta} \beth_\alpha(\mu)$.

Proposition. *If $\lambda > |\mathcal{M}| + |\mathcal{L}(\tau)|$ is strongly inaccessible, then there exists $\mathcal{N}$ an elementary extension of $\mathcal{M}$ such that $|\mathcal{N}| = \lambda$ and $\mathcal{N}$ is λ-saturated, i.e. $\mathcal{N}$ is saturated.*

Proof. In our construction one may arrange that

$$|\mathcal{M}_\alpha| \leqslant \beth_\alpha(|\mathcal{M}|)$$

To see this we split into the three different cases.

- $\mathcal{M}_0 = \mathcal{M}$ so $|\mathcal{M}_0| = |\mathcal{M}| = \beth_0(|\mathcal{M}|)$.

- For δ a limit,

$$|\mathcal{M}_\delta| = |\bigcup_{\alpha<\delta} \mathcal{M}_\alpha| = \lim_{\alpha<\delta} |\mathcal{M}_\alpha| \leqslant \lim_{\alpha<\delta} \beth_\alpha(|\mathcal{M}|) \leqslant \beth_\delta(|\mathcal{M}|).$$

- For $\alpha + 1$ we use the remark that we made during the proof. I.e.

$$|\mathcal{M}_{\alpha+1}| \leqslant 2^{|\mathcal{M}_\alpha|} \leqslant 2^{\beth_\alpha(|\mathcal{M}|)} = \beth_{\alpha+1}(|\mathcal{M}|)$$

Now for λ a strongly inaccessible cardinal and $\mu, \alpha < \lambda$ we have $\beth_\alpha(\mu) < \lambda$. But $|\mathcal{M}_\alpha| \geqslant |\alpha|$ and so in the limit we get $|\mathcal{M}_\lambda| = |\mathcal{N}| = \lambda$. So $\mathcal{N}$ is indeed saturated. □

So under the set-theoretic assumption that there are arbitrarily large inaccessible cardinals, if follows that any structure can be elementarily embedded into a saturated structure.

Remark. If one assumes the Generalized Continuum Hypothesis then saturated models do exist.

26.3 Categoricity of Saturated Models

In general it is not true that elementarily equivalence implies isomorphism. But if the structures in question are saturated, then this does hold.

Theorem 26.2. *If $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent and $|\mathfrak{A}| = |\mathfrak{B}|$ and $\mathfrak{A}$ and $\mathfrak{B}$ are saturated, then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. We use the back-and-forth method.

We list $A = (a_\alpha : \alpha < |\mathfrak{A}|)$ and $B = (b_\alpha : \alpha < |\mathfrak{B}|)$. We'll build, by recursion, a chain $(f_\alpha)_{\alpha<|\mathfrak{A}|}$ of partial isomorphisms such that

- $|f_\alpha| \leqslant |\alpha|.2$ (multiplication in the sense of cardinals).
- For all β we'll have $\mathrm{dom}(f_\beta) \supseteq \{a_\alpha : \alpha < \beta\}$.
- For all β we'll have $\mathrm{range}(f_\beta) \supseteq \{b_\alpha : \alpha < \beta\}$.

The construction is as follows.

- For $\alpha = 0$ let $f_0 := \varnothing$. This is a partial isomorphism since we have assumed that $\mathfrak{A} \equiv \mathfrak{B}$.
- For α a limit we let $f_\alpha := \bigcup_{\beta<\alpha} f_\beta$.
- For $\alpha + 1$ we are given f_α and charged with putting a_α into the domain and b_α into the range. Consider the type $p := \mathrm{tp}(a_\alpha/\mathrm{dom}(f_\alpha))$. Let q be the push-forward
$$q := (f_\alpha)_*(p).$$
Since f_α is a partial isomorphism q is actually a 1-type over $\mathrm{range}(f_\alpha)$. By induction $|f_\alpha| < |\alpha|.2 < |\mathfrak{B}|$ so since $\mathfrak{B}$ is saturated q is realized by some $c \in B$. We define $f_{\alpha+\frac{1}{2}}(a_\alpha) := c$. Now we consider $(f_{\alpha+\frac{1}{2}})^{-1}$ to get b_α into the range of $f_{\alpha+\frac{1}{2}}$. The resulting map is $f_{\alpha+1}$, which clearly satisfies the conditions.

Now let $f := \bigcup_{\alpha<|\mathfrak{A}|} f_\alpha$. Since (f_α) is a chain, f is a function. By construction $\mathrm{dom}(f) = A$ and $\mathrm{range}(f) = B$. Since each f_α is a partial isomorphism, f is an actual isomorphism. □

So while saturated models might not exists (for a given cardinality) if there is one, then it is unique.

We have seen that λ-saturation is equivalent to λ-homogeneous and λ-universal. In fact we can weaken the above assumptions slightly.

Theorem 26.3. *If $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent, $|\mathfrak{A}| = |\mathfrak{B}| =: \lambda$ and $\mathfrak{A}$ and $\mathfrak{B}$ are λ-homogeneous and if they realize the same n-types over $\varnothing$ (for all n), then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. We use the back-and-forth argument, but this time we will need a bit more work. First we prove a lemma.

Lemma. *For all* $\beta < \lambda$ *and all* $(a_\alpha)_{\alpha<\beta} \in A^\beta$ *there exists* $(b_\alpha)_{\alpha<\beta} \in B^\beta$ *such that* $(\mathfrak{A}, (a_\alpha)_{\alpha<\beta}) \equiv (\mathfrak{B}, (b_\alpha)_{\alpha<\beta})$ *(likewise with the roles of* $\mathfrak{A}$ *and* $\mathfrak{B}$ *reversed).*

Proof. (of lemma) For β finite this follows from our hypothesis that $\mathfrak{A}$ and $\mathfrak{B}$ realize the same n-types for all finite n. For $\beta \geqslant \omega$ there are two cases.

1) The first case is when β is *not* a cardinal. Then the conclusion follows by induction. Let γ be any ordinal that is smaller (in the sense of order), but has the same cardinality as β. Fix a bijection $g : \beta \longrightarrow \gamma$. Now use g to reorder β: i.e. $a'_\delta := \alpha_{g^{-1}(\delta)}$ for $\delta < \gamma$. By induction there is some $(b'_\delta)_{\delta<\gamma}$ such that $(\mathfrak{A}, (a'_\delta)_{\delta<\gamma}) \equiv (\mathfrak{B}, (b'_\delta)_{\delta<\gamma})$ but this is just saying that $(\mathfrak{A}, (a_\alpha)_{\alpha<\beta}) \equiv (\mathfrak{B}, (b_\alpha)_{\alpha<\beta})$.
2) Now suppose β is a cardinal. We will construct the sequence $(b_\alpha)_{\alpha<\beta}$ be recursion on α. At a given stage $\gamma < \beta$ we have constructed $(b_\alpha)_{\alpha<\gamma}$. Since $\gamma < \beta$ there exists (by the induction hypothesis) $(c_\alpha)_{\alpha\leqslant\gamma}$ such that

$$(\mathfrak{A}, (a_\alpha)_{\alpha\leqslant\gamma}) \equiv (\mathfrak{B}, (c_\alpha)_{\alpha\leqslant\gamma}).$$

Now since

$$(\mathfrak{A}, (a_\alpha)_{\alpha<\gamma}) \equiv (\mathfrak{B}, (b_\alpha)_{\alpha<\gamma}).$$

by transitivity we have

$$(\mathfrak{B}, (b_\alpha)_{\alpha<\gamma}) \equiv (\mathfrak{B}, (c_\alpha)_{\alpha<\gamma}).$$

By the homogeneity of $\mathfrak{B}$ used on c_γ there is some $b_\gamma \in \mathfrak{B}$ such that

$$(\mathfrak{B}, (b_\alpha)_{\alpha<\gamma}, b_\gamma) \equiv (\mathfrak{B}, (c_\alpha)_{\alpha<\gamma}, c_\gamma) \equiv (\mathfrak{A}, (a_\alpha)_{\alpha\leqslant\gamma}).$$

So we can handle this case as well.

This proves the lemma. □

We complete the proof of the theorem using the back-and-forth argument. List $A = (a_\alpha : \alpha < \lambda)$ and $B = (b_\alpha : \alpha < \lambda)$. We will construct a chain (f_α) with the same requirements as in the proof of Theorem 26.2. For $\alpha = 0$ and for α a limit ordinal, things work like they did above.

In the successor case $(\alpha + 1)$ we get, by the lemma, that there is some $(c_\beta)_{\beta\leqslant\alpha} \in B^\alpha$ such that

$$(\mathfrak{A}, (a_\beta)_{\beta\leqslant\alpha}) \equiv (\mathfrak{B}, (c_\beta)_{\beta\leqslant\alpha}).$$

Using homogeneity and that

$$(\mathfrak{B}, (c_\beta)_{\beta<\alpha}) \equiv (\mathfrak{B}, (f_\alpha(a_\beta))_{\beta<\alpha}),$$

there is some d_α such that

$$(\mathfrak{A}, (a_\beta)_{\beta\leqslant\alpha}) \equiv (\mathfrak{B}, (c_\beta)_{\beta<\alpha}, c_\alpha) \equiv (\mathfrak{B}, (f_\alpha(a_\beta))_{\beta<\alpha}, d_\alpha).$$

We extend f_α by $f_{\alpha+1}(a_\alpha) := d_\alpha$. Likewise going backwards.

Taking the union of the chain of partial isomorphisms gives the required isomorphism $f : \mathfrak{A} \longrightarrow \mathfrak{B}$. $\square$

Lecture 27

27.1 Quantifier Elimination

We can now use the machinery of saturated models to find more semantic (algebraic) methods of proving quantifier elimination. Methods that allow us to use structural information about models of a theory rather than brute force syntactic arguments.

Proposition. *Let $\mathcal{M}$ be a τ-structure and assume that $\mathcal{M}$ is saturated. The following are equivalent.*

1) *For any $\mathfrak{A}, \mathfrak{B} \subseteq \mathcal{M}$ such that $|\mathfrak{A}| < |\mathcal{M}|$ and such that there is an isomorphism $f : \mathfrak{A} \longrightarrow \mathfrak{B}$, then there exists and extension of f to an automorphism of $\mathcal{M}$.*
2) $\mathrm{Th}(\mathcal{M})$ *eliminates quantifiers.*

Proof. "$\Longleftarrow$": Suppose $\mathrm{Th}(\mathcal{M})$ eliminates quantifiers. Then every isomorphism of substructures of $\mathcal{M}$ is a partial elementary map since such isomorphisms preserve quantifier-free formulae, which is suffices by quantifier elimination. Now since $\mathcal{M}$ is saturated, every partial map between small substructures extends to an automorphism of $\mathcal{M}$.

"$\Longrightarrow$": Now we suppose that $\mathcal{M}$ satisfies property 1). We first prove a lemma:

Lemma. *For any $\mathcal{N} \models \mathrm{Th}(\mathcal{M})$ and any $a \in N^n$ (for $n \in \omega$) then $\mathrm{qf.tp}(a) \vdash \mathrm{tp}(a)$.*

Proof. (of lemma) If the claim were false then there would exist some $\mathcal{N}$ which is elementarily equivalent to $\mathcal{M}$, and some $a \in N^n$ such that $\mathrm{qf.tp}(a) \nvdash \mathrm{tp}(a)$. In particular there is some elementary extension $\mathcal{N}'$ of $\mathcal{N}$ and some $b \in (N')^n$ such that $\mathrm{qf.tp}(a) = \mathrm{qf.tp}(b)$ but $\mathrm{tp}(a) \neq \mathrm{tp}(b)$. Now by the Downward Löwenheim-Skolem there is some elementary substructure $\mathcal{N}'' \preccurlyeq \mathcal{N}'$ such that $a, b \in (N'')^n$ and such that $|\mathcal{N}''| \leqslant |\mathcal{L}(\tau)| \leqslant |\mathcal{M}|$. Since $\mathcal{M}$ is saturated, it is $|\mathcal{M}|$-universal, and so there is some embedding $\iota : \mathcal{N}'' \longrightarrow \mathcal{M}$. But now consider the substructures $\mathfrak{A} = \langle \iota(a) \rangle_{\mathcal{M}}$ and $\mathfrak{B} = \langle \iota(b) \rangle_{\mathcal{M}}$. Since $\mathrm{qf.tp}(a) = \mathrm{qf.tp}(b)$ the map $a \longmapsto b$ induces an $\mathcal{L}(\tau)$-isomorphism from $f : \mathfrak{A} \longrightarrow \mathfrak{B}$. By hypothesis 1) there is some automorphism

$\sigma \in \mathrm{Aut}(\mathcal{M})$ which extends f. But then

$$\mathrm{tp}(a) = \mathrm{tp}(\iota(a) = \mathrm{tp}(\sigma(\iota(a)))_) = \mathrm{tp}(\iota(b)) = \mathrm{tp}(b)$$

contradicting that $\mathrm{tp}(a) \neq \mathrm{tp}(b)$. $\square$

By the lemma we know that for all models of $\mathrm{Th}(\mathcal{M})$ the quantifier-free types determine the types. We finish the proof by showing that this implies that $\mathrm{Th}(\mathcal{M})$ has quantifier-elimination.

Let $\theta(x)$ be some formula (x is a tuple). Consider the partial type over $\varnothing$ given by

$$\{\theta(a) \wedge \neg\theta(b)\} \cup \{\psi(a) \leftrightarrow \psi(b) : \psi \in \mathcal{L}(\tau_x) \text{ and } \psi \text{ is quantifier-free}\}.$$

If this theory is consistent, then we can find $a, b \in M^{l(x)}$ such that $\mathrm{qf.tp}(a) = \mathrm{qf.tp}(b)$ and $\mathrm{tp}(a) \neq \mathrm{tp}(b)$. But by the lemma we know that this cannot happen. So the theory must be inconsistent. By compactness there is some finite set Ψ of quantifier-free formulae such that if a and b agree on them, then a and b agree on θ as well. I.e.

$$\mathrm{Th}(\mathcal{M}) \vdash \forall x, y \left[\bigwedge_{\psi \in \Psi} (\psi(x) \leftrightarrow \psi(y)) \longrightarrow [\theta(x) \leftrightarrow \theta(y)] \right].$$

Define

$$\Theta := \left\{ s \subseteq \Psi : \mathcal{M} \models \exists x \left[\theta(x) \wedge \bigwedge_{\Psi \in s} \psi(x) \wedge \bigwedge_{\psi \in \Psi \setminus s} \neg\psi(x) \right] \right\}$$

and let η be the formula

$$\eta(x) := \bigvee_{s \in \Theta} \left[\bigwedge_{\psi \in s} \psi(x) \wedge \bigwedge_{\psi \notin s} \neg\psi \right].$$

Then η is quantifier-free and first-order since Ψ is finite. Now we claim that η and θ are equivalent.

Suppose first that $\mathcal{M} \models \theta(a)$. Then let $s := \{\psi \in \Psi : \mathcal{M} \models \psi(a)\}$, so $s \in \Theta$ is witnessed by a. Thus

$$\mathcal{M} \models \bigwedge_{\psi \in s} \psi(a) \wedge \bigwedge_{\psi \notin s} \neg\psi(a)$$

so $\mathcal{M} \models \eta(a)$. Thus $\mathrm{Th}(\mathcal{M}) \vdash \theta \longrightarrow \eta$.

Now suppose $\mathcal{M} \models \eta(b)$. Let $s \in \Theta$ be such that

$$\mathcal{M} \models \bigwedge_{\psi \in s} \psi(b) \wedge \bigwedge_{\psi \notin s} \neg\psi(b)$$

then

$$\mathcal{M} \models \exists x \left[\theta(x) \wedge \bigwedge_{\psi \in s} \psi(x) \wedge \bigwedge_{\psi \notin s} \neg\psi(x) \right].$$

Let c from M witness this existential. Then b and c agree on all formulae ψ from Ψ, and further more $\mathcal{M} \models \theta(c)$. But this implies by assumption on Ψ then $\mathcal{M} \models \theta(b)$. Thus $\mathcal{M} \models \eta(b) \longrightarrow \theta(b)$.

All in all we now have that $\mathrm{Th}(\mathcal{M}) \vdash \theta \longleftrightarrow \eta$, so $\mathrm{Th}(\mathcal{M})$ eliminates quantifiers. □

27.2 Examples of Quantifier Elimination

Proposition. $(\mathbb{Q}, <)$ *eliminates quantifiers.*

Proof. Let $\mathcal{M} = (\mathbb{Q}, <)$. $\mathcal{M}$ is $\aleph_0$-categorical and thus $\aleph_0$-saturated (since every finitary type over a finite set of parameters is realized). Suppose $\mathfrak{A}, \mathfrak{B} \subseteq \mathcal{M}$ are given such that $f : \mathfrak{A} \longrightarrow \mathfrak{B}$ is an isomorphism and $|\mathfrak{A}| < |\mathcal{M}| = \aleph_0$.

List $\mathfrak{A}$ and $\mathfrak{B}$ as $\mathfrak{A} = \{a_1 < \cdots < a_n\}$ and $\mathfrak{B} = \{b_1 < \cdots < b_n\}$. Cantor's back-and-forth argument extends f to an automorphism $\sigma \in \mathrm{Aut}(\mathcal{M})$. By the theorem, $\mathrm{Th}(\mathcal{M})$ eliminates quantifiers. □

Now fix a field F. Let $\tau = \{0, +, -, (\lambda)_{\lambda \in F}\}$ be the signature of vector spaces over F. So for V a vector space over F, V becomes a τ-structure under the obvious interpretations of the symbols.

Proposition. *If* V *a vector space over* F*, then* $\mathrm{Th}(V)$ *eliminates quantifiers.*

Proof. We will actually show that every partial isomorphism between substructures of *any* model of $\mathrm{Th}(V)$ extends to an automorphism of that model.

Let $W \models \mathrm{Th}(V)$. Let $\mathfrak{A}, \mathfrak{B} \subseteq W$ be isomorphic say by $f : \mathfrak{A} \longrightarrow \mathfrak{B}$ and such that $|\mathfrak{A}| < |W|$. We claim that f extends to an automorphism of W.

Let X be a maximal subset of W whose image in $W/\mathfrak{A}$ is linearly independent, and let Y be a maximal subset of W whose image in $W/\mathfrak{B}$ is linearly independent.

By general linear algebra we have

$$\mathfrak{B} \oplus W/\mathfrak{B} \cong W \cong \mathfrak{A} \oplus W/\mathfrak{A}.$$

Since $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic, there is a bijection $X \longrightarrow Y$. Using this bijection we see that $W/\mathfrak{B}$ and $W/\mathfrak{A}$ are isomorphic. The direct sum of two isomorphisms of two vector spaces is again an isomorphism which extends each summand. Thus we can extend f to an isomorphism

$$W \cong \mathfrak{A} \oplus W/\mathfrak{A} \cong \mathfrak{B} \oplus W/\mathfrak{B}.$$

□

The following is a very important example.

Proposition. *If* k *is an algebraically closed field then* $\mathrm{Th}(k)$ *eliminates quantifiers.*

Proof. Suppose $L \equiv k$. Let $\mathfrak{A}, \mathfrak{B} \subseteq L$ be small substructures of L and let $f : \mathfrak{A} \longrightarrow \mathfrak{B}$ be an isomorphism. We are working in the language of rings and so $\mathfrak{A}$ and $\mathfrak{B}$ are subrings. However, we may assume that $\mathfrak{A}$ and $\mathfrak{B}$ are fields since we can extend f to the fraction fields of $\mathfrak{A}$ and $\mathfrak{B}$. Further more, by a theorem of algebra an isomorphism of fields extends (non-uniquely) to an isomorphism of the algebraic closures of the fields. So we may assume that $\mathfrak{A}$ and $\mathfrak{B}$ are algebraically closed. Now by a back-and-forth argument we can extend f to a transcendence bases of L over $\mathfrak{A}$ and L over $\mathfrak{B}$. Thus, f extends all the way to an automorphism of L. □

Bibliography

[1] Wilfrid Hodges, *a shorter model theory*. Cambridge University Press, 1997.

[2] C.C. Chang & Jerome Keisler, *Model Theory*. Dover Books on Mathematics, 3rd. Ed., 2012.

[3] Matatyahu Rubin, *The Reconstruction of Trees from Their Automorphism Groups*. American Mathematical Society, 1991.

[4] Peter J. Cameron, *Oligomorphic Permutation Groups*. Cambridge University Press, 1990.