Achim Blumensath
blumens@fi.muni.cz

This document was last updated 2024-04-09.
The latest version can be found at

www.fi.muni.cz/~blumens

Contents

Part E.

Classical Model Theory

E1. *Saturation*

1. *Homogeneous structures*

Recall the relations $\sqsubseteq^{\kappa}_{\mathrm{FO}}$ introduced in Section C4.4. We have seen that, in general, they are not reflexive. In this section we will take a closer look at those structures $\mathfrak{A}$ that satisfy $\mathfrak{A} \cong^{\kappa}_{\mathrm{FO}} \mathfrak{A}$.

Definition 1.1. Let $\mathfrak{A}$ be a Σ-structure and κ a cardinal.

(a) $\mathfrak{A}$ is *κ-homogeneous* if $\mathfrak{A} \cong^{\kappa}_{\mathrm{FO}} \mathfrak{A}$, that is, whenever $\bar{a}, \bar{b} \in A^{<\kappa}$ are sequences of length less than κ with $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle$ and $c \in A$ is another element, then there exists an element $d \in A$ such that $\langle \mathfrak{A}, \bar{a}c \rangle \equiv \langle \mathfrak{A}, \bar{b}d \rangle$. We call $\mathfrak{A}$ *homogeneous* if it is $|A|$-homogeneous.

(b) $\mathfrak{A}$ is *strongly κ-homogeneous* if, whenever $\bar{a}, \bar{b} \in A^{<\kappa}$ are sequences of length less than κ with $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle$ then there exists an automorphism π of $\mathfrak{A}$ such that $\pi(\bar{a}) = \bar{b}$. We call $\mathfrak{A}$ *strongly homogeneous* if it is strongly $|A|$-homogeneous.

Example. (a) The structures $\langle \mathbb{Z}, < \rangle$ and $\langle \mathbb{Q}, < \rangle$ are strongly homogeneous.

(b) The theory of $\langle \omega, \leq \rangle$ has exactly three countable (strongly) homogeneous models whose order types are ω, $\omega + \zeta$, and $\omega + \zeta \cdot \eta$, respectively, where ζ is the order type of the integers and η is the order type of the rationals.

Exercise 1.1. Show that $\langle \mathbb{R}, + \rangle$ is strongly $\aleph_0$-homogeneous.

Lemma 1.2. *Every strongly κ-homogeneous structure is κ-homogeneous.*

Proof. Let $\mathfrak{A}$ be strongly κ-homogeneous. Suppose that $\bar{a}, \bar{b} \in A^{<\kappa}$ are sequences with $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle$ and let $c \in A$. By assumption, there exists an isomorphism $\pi : \langle \mathfrak{A}, \bar{a} \rangle \to \langle \mathfrak{A}, \bar{b} \rangle$. If we set $d := \pi(c)$ then we have

$$\pi : \langle \mathfrak{A}, \bar{a}c \rangle \cong \langle \mathfrak{A}, \bar{b}d \rangle .$$

This implies $\langle \mathfrak{A}, \bar{a}c \rangle \equiv \langle \mathfrak{A}, \bar{b}d \rangle$. □

Lemma 1.3. *Every homogeneous structure is strongly homogeneous.*

Proof. Let $\mathfrak{A}$ be a homogeneous structure of size $\kappa := |A|$. If $\bar{a}, \bar{b} \in A^{<\kappa}$ are sequences with $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle$ then $\mathfrak{A} \cong^{\kappa}_{\mathrm{FO}} \mathfrak{A}$ implies, by definition of $\cong^{\kappa}_{\mathrm{FO}}$, that

$$\langle \mathfrak{A}, \bar{a} \rangle \cong^{\kappa}_{\mathrm{FO}} \langle \mathfrak{A}, \bar{b} \rangle .$$

By Lemma C4.4.10, it follows that $\langle \mathfrak{A}, \bar{a} \rangle \cong \langle \mathfrak{A}, \bar{b} \rangle$. □

Lemma 1.4. *Let T be a first-order theory that admits quantifier elimination for $\mathrm{FO}_{\infty\aleph_0}$. Every model of T is $\aleph_0$-homogeneous.*

Proof. If $\mathfrak{A}$ is a model of T then we have $\mathfrak{A} \cong^{\aleph_0}_{0} \mathfrak{A}$, by Theorem D1.2.9. This implies that $\mathfrak{A} \cong^{\aleph_0}_{\mathrm{FO}} \mathfrak{A}$. □

We have shown in Section C4.4 that $\cong^{\kappa}_{\mathrm{FO}}$ is an equivalence relation on the class of all κ-homogeneous structures. In the following lemmas we will study the corresponding equivalence classes. We will show that we have $\mathfrak{A} \cong^{\kappa}_{\mathrm{FO}} \mathfrak{B}$ if and only if both structures realise the same types.

Lemma 1.5. *Let $\mathfrak{B}$ be κ-homogeneous and suppose that $\mathfrak{A}$ is a structure such that, for all $n < \omega$, every n-type realised in $\mathfrak{A}$ is also realised in $\mathfrak{B}$. For each $\bar{a} \in A^{<\kappa}$, there exists a sequence $\bar{b} \in B^{<\kappa}$ such that*

$$\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{B}, \bar{b} \rangle .$$

Proof. Let $\bar{a} \in A^{\alpha}$, for $\alpha < \kappa$. We prove the statement by induction on α. If $\alpha < \omega$ then, since $\mathfrak{A}$ and $\mathfrak{B}$ realise the same α-types, we can find some tuple $\bar{b}$ with $\mathrm{tp}(\bar{b}/\mathfrak{B}) = \mathrm{tp}(\bar{a}/\mathfrak{A})$. If $\lambda := |\alpha| < \alpha$ then we can fix a bijection $g : \lambda \to \alpha$ and the claim follows if we apply the inductive hypothesis to the reordered sequence $(a_{g(i)})_{i<\lambda}$.

It therefore remains to consider the case that α is an infinite cardinal. We construct $(b_i)_{i<\alpha}$ by induction on i such that, at every step $\beta \leq \alpha$ we have

$$\langle \mathfrak{A}, (a_i)_{i<\beta} \rangle \equiv \langle \mathfrak{B}, (b_i)_{i<\beta} \rangle .$$

For $\beta = 0$, we have $\mathfrak{A} \equiv \mathfrak{B}$ since the unique complete 0-type $\mathrm{Th}(\mathfrak{A})$ realised in $\mathfrak{A}$ is also realised in $\mathfrak{B}$. If β is a limit ordinal then there is nothing to do. Suppose that $\beta = \gamma + 1$ is a successor and we have already defined $(b_i)_{i<\gamma}$. Since α is a limit we have $\beta < \alpha$. Therefore, we can apply the inductive hypothesis for α and it follows that there is some sequence $(c_i)_{i<\beta}$ such that

$$\langle \mathfrak{A}, (a_i)_{i<\beta} \rangle \equiv \langle \mathfrak{B}, (c_i)_{i<\beta} \rangle .$$

In particular, we have

$$\langle \mathfrak{B}, (b_i)_{i<\gamma} \rangle \equiv \langle \mathfrak{A}, (a_i)_{i<\gamma} \rangle \equiv \langle \mathfrak{B}, (c_i)_{i<\gamma} \rangle ,$$

and, since $\mathfrak{B}$ is κ-homogeneous, we can find some element $b_\gamma \in B$ such that

$$\langle \mathfrak{B}, (b_i)_{i<\gamma}, b_\gamma \rangle \equiv \langle \mathfrak{B}, (c_i)_{i<\gamma}, c_\gamma \rangle \equiv \langle \mathfrak{A}, (a_i)_{i<\gamma}, a_\gamma \rangle . \qquad \square$$

Proposition 1.6. *Let $\mathfrak{B}$ be κ-homogeneous and suppose that $\mathfrak{A}$ is a structure such that, for all $n < \omega$, every n-type realised in $\mathfrak{A}$ is also realised in $\mathfrak{B}$. Then $\mathfrak{A} \sqsubseteq^\kappa_{\mathrm{FO}} \mathfrak{B}$.*

Proof. Since $I^\kappa_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ is always κ-complete we only need to prove the forth property. Let $\bar{a} \mapsto \bar{b} \in I^\kappa_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ and $c \in A$. By the preceding lemma, we can find a sequence $\bar{b}'d' \subseteq B$ such that

$$\langle \mathfrak{A}, \bar{a}c \rangle \equiv \langle \mathfrak{B}, \bar{b}'d' \rangle .$$

In particular, we have $\langle \mathfrak{B}, \bar{b} \rangle \equiv \langle \mathfrak{B}, \bar{b}' \rangle$. Since $\mathfrak{B}$ is κ-homogeneous we can therefore find some element $d \in B$ such that

$$\langle \mathfrak{B}, \bar{b}d \rangle \equiv \langle \mathfrak{B}, \bar{b}'d' \rangle \equiv \langle \mathfrak{B}, \bar{b} \rangle .$$

Hence, $\bar{a}c \mapsto \bar{b}d \in I^{\kappa}_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$. □

Corollary 1.7. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be κ-homogeneous structures. We have*

$$\mathfrak{A} \cong^{\kappa}_{\mathrm{FO}} \mathfrak{B} \quad \text{iff} \quad \mathfrak{A} \text{ and } \mathfrak{B} \text{ realise the same } n\text{-types, for all } n < \omega.$$

Corollary 1.8. *If $\mathfrak{A}$ and $\mathfrak{B}$ are $\aleph_0$-homogeneous structures that realise the same n-types, for all $n < \omega$, and $\bar{a} \in A^{<\omega}$, $\bar{b} \in B^{<\omega}$ are finite tuples then*

$$\langle \mathfrak{A}, \bar{a} \rangle \equiv_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b} \rangle \quad \text{implies} \quad \langle \mathfrak{A}, \bar{a} \rangle \equiv_{\infty} \langle \mathfrak{B}, \bar{b} \rangle .$$

Proof. This follows by Proposition 1.6 and Theorem D1.2.13. □

Theorem 1.9. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be homogeneous structures of the same size $|A| = |B|$. If, for every $n < \omega$, $\mathfrak{A}$ and $\mathfrak{B}$ realise the same n-types then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. Let $\kappa := |A| = |B|$. By Proposition 1.6, we have $\mathfrak{A} \sqsubseteq^{\kappa}_{\mathrm{FO}} \mathfrak{B}$ and $\mathfrak{A} \sqsupseteq^{\kappa}_{\mathrm{FO}} \mathfrak{B}$. Hence, the claim follows from Lemma C4.4.10 (a). □

Corollary 1.10. *A complete first-order theory T has, up to isomorphism, for every cardinal κ at most $2^{2^{|T|}}$ homogeneous models of size κ.*

Proof. For every set $X \subseteq S^{<\omega}(T)$, there is, according to the preceding theorem, at most one homogeneous model of size κ that realises exactly the types in X. Since $|S^{<\omega}(T)| \leq 2^{|T|}$ the claim follows. □

To build κ-homogeneous structures we can use the following lemma. We will defer the proof of the fact that every structure has a κ-homogeneous elementary extension to Section 3 where it will follow from a much stronger result.

Lemma 1.11. *Let $\mathfrak{A}$ be a Σ-structure and $\bar{a}, \bar{b} \in A^{\alpha}$ tuples with $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle$.*

(a) *There exists an elementary extension* $\mathfrak{B} \succeq \mathfrak{A}$ *such that*

$$\langle \mathfrak{B}, \bar{a} \rangle \equiv_\infty \langle \mathfrak{B}, \bar{b} \rangle \quad \textit{and} \quad |B| \leq |A| \oplus |\Sigma| \oplus |\alpha| \oplus \aleph_0 \,.$$

(b) *There exists an elementary extension* $\mathfrak{B} \succeq \mathfrak{A}$ *and an automorphism* $\pi \in \operatorname{Aut} \mathfrak{B}$ *with* $\pi(\bar{a}) = \bar{b}$.

Proof. (a) For $0 \leq k < \omega$, let I_k be a new $2k$-ary relation symbol and set

$$\beta_k := (\forall \bar{x}\bar{y}.I_k\bar{x}\bar{y})[\forall u \exists v I_{k+1}\bar{x}u\bar{y}v \wedge \forall v \exists u I_{k+1}\bar{x}u\bar{y}v] \,,$$

and $\quad \psi_k^\varphi := (\forall \bar{x}\bar{y}.I_k\bar{x}\bar{y})[\varphi(\bar{a},\bar{x}) \leftrightarrow \varphi(\bar{b},\bar{y})] \,.$

The formula β_k says that I_k has the back-and-forth property with respect to I_{k+1}, and the ψ_k^φ hold if every tuple $\langle \bar{c}, \bar{d} \rangle \in I_k$ corresponds to a partial isomorphism $\bar{c} \mapsto \bar{d}$ from $\langle \mathfrak{A}, \bar{a} \rangle$ to $\langle \mathfrak{A}, \bar{b} \rangle$. Setting

$$\Phi := \operatorname{Th}(\mathfrak{A}_A) \cup \{I_0\} \cup \{\, \beta_k \wedge \psi_k^\varphi \mid k < \omega,\ \varphi \text{ an atomic formula} \,\} \,,$$

we have

$$\mathfrak{B} \vDash \Phi \quad \text{iff} \quad \mathfrak{B} \succeq \mathfrak{A} \quad \text{and} \quad \langle\rangle \mapsto \langle\rangle \in I_\infty(\langle \mathfrak{B}, \bar{a} \rangle, \langle \mathfrak{B}, \bar{b} \rangle) \,.$$

If Φ is satisfiable then we can, therefore, use the Theorem of Löwenheim and Skolem to find the desired structure $\mathfrak{B}$. To prove that Φ is satisfiable let $\Phi_0 \subseteq \Phi$ be finite. There is some $m < \omega$ and a finite set Δ of atomic formulae such that

$$\Phi_0 \subseteq \operatorname{Th}(\mathfrak{A}_A) \cup \{I_0\} \cup \{\, \beta_k \wedge \psi_k^\varphi \mid k < m, \varphi \in \Delta \,\} \,.$$

Let $\bar{a}'$ and $\bar{b}'$ be the subsequences of, respectively, $\bar{a}$ and $\bar{b}$ that appear in Δ. Since $\operatorname{tp}(\bar{a}') = \operatorname{tp}(\bar{b}')$ we can obtain a model $\langle \mathfrak{A}_A, (I_k)_{k<m} \rangle \vDash \Phi_0$ by setting

$$I_k := \left\{\, \bar{c}\bar{d} \in A^{2k} \;\middle|\; \langle \mathfrak{A}, \bar{a}'\bar{c} \rangle \equiv_{m-k} \langle \mathfrak{A}, \bar{b}'\bar{d} \rangle \,\right\} .$$

(b) Let f be a new unary function symbol and set

$$\begin{aligned}\Phi := \mathrm{Th}(\mathfrak{A}_A) &\cup \{ f a_i = b_i \mid i < \alpha \} \\ &\cup \{ \forall x \exists y f y = x \} \\ &\cup \{ \forall \bar{x}(\varphi(\bar{x}) \leftrightarrow \varphi(f\bar{x})) \mid \varphi \text{ an atomic formula} \}.\end{aligned}$$

If $\mathfrak{B} \vDash \Phi$ then $f^{\mathfrak{B}}$ is the desired automorphism. Therefore, it is sufficient to prove that Φ is satisfiable.

Let $\Phi_0 \subseteq \Phi$ be finite. There are finitely many indices $k_0, \dots, k_{n-1} < \alpha$, a finite set $C \subseteq A$, a finite signature $\Sigma_0 \subseteq \Sigma$, and a finite set Δ of atomic formulae over Σ_0 such that

$$\begin{aligned}\Phi_0 \subseteq \mathrm{Th}(\mathfrak{A}_C) &\cup \{ f a_{k_i} = b_{k_i} \mid i < n \} \\ &\cup \{ \forall x \exists y f y = x \} \\ &\cup \{ \forall \bar{x}(\varphi(\bar{x}) \leftrightarrow \varphi(f\bar{x})) \mid \varphi \in \Delta \}.\end{aligned}$$

To simplify notation, set $\bar{a}' = a_{k_0} \dots a_{k_{n-1}}$ and $\bar{b}' = b_{k_0} \dots b_{k_{n-1}}$. By the Theorem of Löwenheim and Skolem, we can find a countable elementary substructure $\mathfrak{A}_0 \preceq \mathfrak{A}|_{\Sigma_0}$ with $C \cup \bar{a}'\bar{b}' \subseteq A_0$.

By (a), there exists a countable elementary extension $\mathfrak{B}_0 \succeq \mathfrak{A}_0$ such that

$$\langle \mathfrak{B}_0, \bar{a}' \rangle \equiv_\infty \langle \mathfrak{B}_0, \bar{b}' \rangle .$$

Hence, by Lemma C4.4.10, it follows that

$$\langle \mathfrak{B}_0, \bar{a}' \rangle \cong \langle \mathfrak{B}_0, \bar{b}' \rangle ,$$

and there is some automorphism $\pi \in \mathrm{Aut}\, \mathfrak{B}_0$ with $\pi(\bar{a}') = \bar{b}'$. Consequently, $\langle \mathfrak{B}_0, \pi \rangle$ is the desired model of Φ_0. □

Exercise 1.2. Let κ be an infinite cardinal. Prove that every structure has a κ-homogeneous elementary extension.

2. Saturated structures

We have shown in the previous section that κ-homogeneous structures can be ordered with respect to the set of types they realise. In this section we consider structures that are maximal in this ordering, i.e., homogeneous structures realising every type.

Definition 2.1. Let $\mathfrak{A}$ be a Σ-structure and κ a cardinal.

(a) $\mathfrak{A}$ is *κ-saturated* if, for all sets $C \subseteq A$ of size $|C| < \kappa$, every type $\mathfrak{p} \in S^{<\omega}(C)$ is realised in $\mathfrak{A}$. A structure $\mathfrak{A}$ is called *saturated* if it is $|A|$-saturated.

(b) $\mathfrak{A}$ is *κ-universal* if there exist elementary embeddings $\mathfrak{B} \to \mathfrak{A}$, for all Σ-structures $\mathfrak{B}$ of size $|B| < \kappa$ such that $\mathfrak{B} \equiv \mathfrak{A}$.

Similarly to homogeneous structures we can characterise κ-saturated structures in terms of the relation $\sqsubseteq^{\kappa}_{\mathrm{FO}}$.

Lemma 2.2. *A structure $\mathfrak{B}$ is κ-saturated if and only if*

$$\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{B}, \bar{b} \rangle \quad \textit{implies} \quad \langle \mathfrak{A}, \bar{a} \rangle \sqsubseteq^{\kappa}_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b} \rangle\,,$$

for all sequences $\bar{a} \in A^{<\kappa}$ and $\bar{b} \in B^{<\kappa}$.

Proof. ($\Rightarrow$) Suppose that $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{B}, \bar{b} \rangle$. We have $\bar{a} \mapsto \bar{b} \in I^{\kappa}_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ and $I^{\kappa}_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ is κ-complete. Therefore, we only need to prove the forth property. Suppose that $\bar{c} \mapsto \bar{d} \in I^{\kappa}_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ and $e \in A$. Set $\mathfrak{p} := \mathrm{tp}(e/\mathfrak{A}_{\bar{c}})$ and let $\mathfrak{q}$ be the type obtained from $\mathfrak{p}$ by replacing the constants $\bar{c}$ by $\bar{d}$. Note that $\mathfrak{q}$ really is a type since $\langle \mathfrak{A}, \bar{c} \rangle \equiv \langle \mathfrak{B}, \bar{d} \rangle$. As $|\bar{d}| < \kappa$ and $\mathfrak{B}$ is κ-saturated we can find some element $f \in B$ realising $\mathfrak{q}$. Therefore,

$$\langle \mathfrak{A}, \bar{c}e \rangle \equiv \langle \mathfrak{B}, \bar{d}f \rangle\,, \quad \text{that is,} \quad \bar{c}e \mapsto \bar{d}f \in I^{\kappa}_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})\,.$$

($\Leftarrow$) Let $C \subseteq B$ be a set of size $|C| < \kappa$ and $\mathfrak{p} \in S^{n}(C)$. There exists an elementary extension $\mathfrak{A} \succeq \mathfrak{B}$ in which $\mathfrak{p}$ is realised by some tuple $\bar{a}$. Let $\bar{c}$ be an enumeration of C. Since $\langle \mathfrak{A}, \bar{c} \rangle \equiv \langle \mathfrak{B}, \bar{c} \rangle$ we have

$$\langle \mathfrak{A}, \bar{c} \rangle \sqsubseteq^{\kappa}_{\mathrm{FO}} \langle \mathfrak{B}, \bar{c} \rangle\,.$$

Hence, by Lemma C4.4.9 we can find a tuple $\bar{b} \in B^n$ such that

$$\langle \mathfrak{A}, \bar{c}\bar{a} \rangle \sqsubseteq^{\kappa}_{\mathrm{FO}} \langle \mathfrak{B}, \bar{c}\bar{b} \rangle .$$

Consequently, $\bar{b}$ is a realisation of $\mathfrak{p}$ in B. □

Corollary 2.3. *For κ-saturated structures $\mathfrak{A}$ and $\mathfrak{B}$, we have*

$$\langle \mathfrak{A}, \bar{a} \rangle \cong^{\kappa}_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b} \rangle \quad \text{iff} \quad \langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{B}, \bar{b} \rangle ,$$

for all $\bar{a} \in A^{<\kappa}$ and $\bar{b} \in B^{<\kappa}$.

We will prove below that every κ-saturated structure is κ-homogeneous. Hence, the next corollary is a special case of Corollary 1.8.

Corollary 2.4. *If $\mathfrak{A}$ and $\mathfrak{B}$ are $\aleph_0$-saturated then*

$$\mathfrak{A} \equiv \mathfrak{B} \quad \textit{implies} \quad \mathfrak{A} \cong_{\infty} \mathfrak{B} .$$

For an example let us take a look at saturated linear orders.

Lemma 2.5. *Every $\aleph_1$-saturated dense linear order is incomplete.*

Proof. Let $a_0 < a_1 < \dots$ be a strictly increasing sequence of length ω and set $A := \{ a_n \mid n < \omega \}$. We claim that $\sup A$ does not exist. For a contradiction, suppose that the supremum c exists. Choose a type $\mathfrak{p}$ over $A \cup \{c\}$ containing the formulae

$$x < c \quad \text{and} \quad a_n < x \quad \text{for } n < \omega .$$

Any realisation b of $\mathfrak{p}$ is an upper bound of A. Hence, $b < c = \sup A$ yields the desired contradiction. □

Lemma 2.6. *A linear order is κ-saturated if, and only if, it is κ-dense.*

Proof. We have already shown in Lemma C4.4.6 that every κ-dense linear order is κ-saturated. For the converse, suppose that $\mathfrak{A} = \langle A, \leq \rangle$

is κ-saturated and let $C, D \subseteq A$ sets of size $|C|, |D| < \kappa$ with $C < D$. Let $\mathfrak{p} \in S^1(C \cup D)$ be any type with

$$\mathfrak{p} \supseteq \{ c < x \mid c \in C \} \cup \{ x < d \mid d \in D \} .$$

Since $\mathfrak{A}$ is κ-saturated there is some element $a \in A$ realising $\mathfrak{p}$. Hence, $C < a < D$ and $\mathfrak{A}$ is κ-dense. □

Lemma 2.7. *Let $(\mathfrak{A}^i)_{i<\lambda}$ be an elementary chain of κ-saturated structures. If $\kappa \leq \operatorname{cf} \lambda$ then the union $\bigcup_i \mathfrak{A}^i$ is also κ-saturated.*

Proof. Let $C \subseteq \bigcup_i A^i$ be a set of size $|C| < \kappa$ and suppose that $\mathfrak{p} \in S^{<\omega}(C)$ is a type over C. Since $|C| < \kappa \leq \operatorname{cf} \lambda$ there is some $\alpha < \lambda$ such that $C \subseteq A^\alpha$. Hence, there is a tuple $\bar{a} \subseteq A^\alpha \subseteq \bigcup_i A^i$ realising $\mathfrak{p}$. □

By definition a structure is κ-saturated if it realises every n-type, for $n < \omega$, with less than κ-parameters. In fact, it is sufficient to realise all 1-types.

Lemma 2.8. *Let $\kappa \geq \aleph_0$. A structure $\mathfrak{A}$ is κ-saturated if, and only if, whenever $C \subseteq A$ is of size $|C| < \kappa$ then every 1-type in $S^1(C)$ is realised in $\mathfrak{A}$.*

Exercise 2.1. Prove the preceding lemma.

Theorem 2.9. *Let $\mathfrak{A}$ be a Σ-structure. The following statements are equivalent:*

(1) *$\mathfrak{A}$ is κ-saturated.*

(2) *$\mathfrak{A}$ is κ-homogeneous and it realises every type in $S^\kappa(\varnothing)$.*

(3) *$\mathfrak{A}$ is κ-homogeneous and it realises every type in $S^{<\kappa}(\varnothing)$.*

If $\kappa \geq |\Sigma| \oplus \aleph_0$ then the following statement is also equivalent to the ones above.

(4) *$\mathfrak{A}$ is κ-homogeneous and κ^+-universal.*

Proof. (1) ⇒ (2) Let $\mathfrak{A}$ be κ-saturated. By Lemma 2.2, $\mathfrak{A} \equiv \mathfrak{A}$ implies $\mathfrak{A} \sqsubseteq^{\kappa}_{\mathrm{FO}} \mathfrak{A}$. Therefore, we have $\mathfrak{A} \cong^{\kappa}_{\mathrm{FO}} \mathfrak{A}$, that is, $\mathfrak{A}$ is κ-homogeneous.

It remains to prove that $\mathfrak{A}$ realises every type $\mathfrak{p} \in S^{\kappa}(\varnothing)$. For $\alpha < \kappa$, let $\mathfrak{p}_\alpha := \mathfrak{p} \cap \mathrm{FO}^{\alpha}[\Sigma]$ be the restriction of $\mathfrak{p}$ to the first α variables. By induction on α, we construct a sequence $(a_\alpha)_{\alpha<\kappa}$ such that the subsequence $(a_i)_{i<\alpha}$ realises $\mathfrak{p}_\alpha$. Suppose we have already defined a_i, for $i < \alpha$. Let

$$\mathfrak{q}_\alpha := \{ \varphi(a_{i_0}, \dots, a_{i_{k-1}}, x_\alpha) \mid \varphi(x_{i_0}, \dots, x_{i_{k-1}}, x_\alpha) \in \mathfrak{p} \text{ for } i_0, \dots, i_{k-1} < \alpha \} .$$

Since $\mathfrak{A}$ is κ-saturated we can find some element a_α such that

$$\mathrm{tp}(a_\alpha/\{ a_i \mid i < \alpha \}) = \mathfrak{q}_\alpha .$$

Hence, $(a_i)_{i\leq\alpha}$ realises $\mathfrak{p}_{\alpha+1}$.

(2) ⇒ (3) is trivial.

(3) ⇒ (1) Let $\mathfrak{p} \in S^n(U)$ where $|U| < \kappa$. Let $(c_i)_{i<\lambda}$ be an enumeration of U and let $\mathfrak{q} \in S^{\lambda+n}(\varnothing)$ be the type

$$\mathfrak{q} := \{ \varphi(x_{i_0}, \dots, x_{i_{k-1}}, x_\lambda, \dots, x_{\lambda+n-1}) \mid \varphi(c_{i_0}, \dots, c_{i_{k-1}}, x_0, \dots, x_{n-1}) \in \mathfrak{p} \} .$$

By assumption we can find sequences $\bar{a} \in A^\lambda$ and $\bar{b} \in A^n$ such that $\mathrm{tp}(\bar{a}\bar{b}) = \mathfrak{q}$. Since

$$\langle \mathfrak{A}, \bar{c} \rangle \equiv \langle \mathfrak{A}, \bar{a} \rangle$$

and $\mathfrak{A}$ is κ-homogeneous it follows that there is some tuple $\bar{d} \in A^n$ such that

$$\langle \mathfrak{A}, \bar{c}\bar{d} \rangle \equiv \langle \mathfrak{A}, \bar{a}\bar{b} \rangle .$$

Consequently $\mathrm{tp}(\bar{d}/\bar{c}) = \mathfrak{p}$.

(2) ⇒ (4) Suppose that $\mathfrak{A}$ realises every type in $S^{\kappa}(\varnothing)$. We claim that $\mathfrak{A}$ is κ^+-universal. Let $\mathfrak{B}$ be a structure of size $|B| \leq \kappa$ with $\mathfrak{B} \equiv \mathfrak{A}$. Choose

an enumeration $\bar{b}$ of B and let $\mathfrak{p} := \mathrm{tp}(\bar{b}/\mathfrak{B})$. Then $\mathfrak{p} \in S^{\leq\kappa}(\varnothing)$. Hence, there exists a sequence $\bar{a} \subseteq A$ realising $\mathfrak{p}$. The function $\bar{b} \mapsto \bar{a}$ is the desired elementary embedding.

(4) ⇒ (1) Suppose that $\mathfrak{A}$ is κ^+-universal. We show that $\mathfrak{A}$ realises every type $\mathfrak{p} \in S^\kappa(\varnothing)$. For each such $\mathfrak{p}$ we can find a structure $\mathfrak{B} \equiv \mathfrak{A}$ and a tuple $\bar{b} \subseteq B$ with $\mathrm{tp}(\bar{b}/\mathfrak{B}) = \mathfrak{p}$. By the Theorem of Löwenheim and Skolem we may assume that $|B| \leq \kappa$. Hence, there exists an elementary embedding $h : \mathfrak{B} \to \mathfrak{A}$. The sequence $h(\bar{b})$ is a realisation of $\mathfrak{p}$ in $\mathfrak{A}$. □

Theorem 2.10. *If $\mathfrak{A} \equiv \mathfrak{B}$ are saturated structures of the same size $|A| = |B|$ then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. Let $\kappa := |A| = |B|$. By Lemma 2.2, we have $\mathfrak{A} \cong^\kappa_{\mathrm{FO}} \mathfrak{B}$. Therefore, the claim follows from Lemma C4.4.10 (a). □

Every structure has a κ-saturated elementary extension. There are two ways to construct such extensions: (i) we can form an ultrapower, or (ii) we can take the union of an infinite elementary chain where each structure realises every type over the universe of the preceding structure. In the following proofs we will employ the first method. Below, where we construct saturated structures and projectively κ-saturated ones, we will choose the second method.

Proposition 2.11. *Let $\mathfrak{u}$ be a regular ultrafilter over an infinite set I and let $(\mathfrak{A}^i)_{i\in I}$ be a family of structures. Every countable partial type $\mathfrak{p}$ over $\prod_i A_i/\mathfrak{u}$ is realised in $\prod_i \mathfrak{A}_i/\mathfrak{u}$.*

Proof. Let $(\varphi_n)_{n<\omega}$ be an enumeration of $\mathfrak{p}$. Since $\mathfrak{u}$ is regular, we can find sets $(s_n)_{n<\omega}$ in $\mathfrak{u}$ such that, for every $i \in I$, the set

$$\{\, n < \omega \mid i \in s_n \,\}$$

is finite. Setting $w_n := s_0 \cap \dots \cap s_n \in \mathfrak{u}$ we obtain a strictly decreasing sequence $w_0 \supset w_1 \supset w_2 \supset \cdots$ of sets $w_n \in \mathfrak{u}$. By choice of $(s_n)_n$ we have

$$\bigcap_{n<\omega} w_n = \bigcap_{n<\omega} s_n = \varnothing\,.$$

Set $\psi_n := \varphi_0 \wedge \dots \wedge \varphi_n$ and let $[\bar{c}^n]_{\mathfrak{u}}$ be the parameters appearing in ψ_n. According to the Theorem of Łoś,

$$\prod_i \mathfrak{A}_i/\mathfrak{u} \vDash \exists \bar{x} \psi_n(\bar{x}; [\bar{c}^n]_{\mathfrak{u}}) \quad \text{implies} \quad [\![\exists \bar{x} \psi_n(\bar{x}; \bar{c}^n)]\!] \in \mathfrak{u} \,.$$

Hence, the sets

$$w_n^{\circ} := \{ i \in w_n \mid \mathfrak{A}_i \vDash \exists \bar{x} \psi_n(\bar{x}; \bar{c}_i^n) \} = w_n \cap [\![\exists \bar{x} \psi_n]\!]$$

are in $\mathfrak{u}$. We define a sequence $(\bar{a}_i)_{i \in I}$ as follows. If $i \notin w_0^{\circ}$, we choose an arbitrary tuple $\bar{a}_i \subseteq A_i$. Otherwise, let n be the maximal number such that $i \in w_n^{\circ}$ and let $\bar{a}_i \subseteq A_i$ be a tuple such that $\mathfrak{A}_i \vDash \psi_n(\bar{a}_i; \bar{c}_i^n)$.

We claim that $[\bar{a}]_{\mathfrak{u}}$ realises $\mathfrak{p}$. Consider $\varphi_n \in \mathfrak{p}$. Then

$$[\![\varphi_n(\bar{a}_i)]\!] \supseteq [\![\psi_n(\bar{a}_i)]\!] \supseteq w_n^{\circ} \in \mathfrak{u} \quad \text{implies} \quad [\![\varphi_n(\bar{a}_i)]\!] \in \mathfrak{u} \,.$$

By the Theorem of Łoś it follows that $\prod_i \mathfrak{A}_i/\mathfrak{u} \vDash \varphi_n([\bar{a}]_{\mathfrak{u}})$. □

Corollary 2.12. *Let $\mathfrak{u}$ be a regular ultrafilter of an infinite set I and let Σ be a countable signature. For every sequence $(\mathfrak{A}_i)_{i \in I}$ of Σ-structures, the ultraproduct $\prod_{i \in I} \mathfrak{A}_i/\mathfrak{u}$ is $\aleph_1$-saturated.*

Proposition 2.13. *Let $\mathfrak{u}$ be an ultrafilter over a set I of size $\kappa := |I|$. The following statements are equivalent:*

(1) *$\mathfrak{u}$ is regular.*

(2) *For each theory T and every family $(\mathfrak{A}_i)_{i \in I}$ of models of T, the ultraproduct $\prod_i \mathfrak{A}_i/\mathfrak{u}$ realises every partial type $\mathfrak{p}$ over $\varnothing$ with $|\mathfrak{p}| \leq \kappa$.*

(3) *For every structure $\mathfrak{M}$, the ultrapower $\mathfrak{M}^{\mathfrak{u}}$ realises every partial type $\mathfrak{p}$ over M with $|\mathfrak{p}| \leq \kappa$.*

Proof. (1) ⇒ (2) Since $|\mathfrak{p}| \leq |I|$ and $\mathfrak{u}$ is regular we can find sets $(s_\varphi)_{\varphi \in \mathfrak{p}}$ in $\mathfrak{u}$ such that the sets

$$\Phi_i := \{ \varphi \in \mathfrak{p} \mid i \in s_\varphi \}$$

are finite. For every $i \in I$, there exists a tuple $\bar{a}^i \subseteq A_i$ realising the finite type Φ_i. We claim that $\bar{a} := (\bar{a}^i)_i$ realises $\mathfrak{p}$. Let $\varphi \in \mathfrak{p}$. For every $k \in s_\varphi$, we have $k \in [\![\varphi(\bar{a}^i)]\!]_i$. Hence, $s_\varphi \subseteq [\![\varphi(\bar{a}^i)]\!]_i \in \mathfrak{u}$ which implies, by the Theorem of Łoś, that $\prod_i \mathfrak{A}_i/\mathfrak{u} \vDash \varphi([\bar{a}]_\mathfrak{u})$.

(2) ⇒ (3) follows by setting $\mathfrak{A}_i := \mathfrak{M}_M$, for each $i \in I$.

(3) ⇒ (1) We consider the structure $\mathfrak{M} := \langle M, \subseteq \rangle$ where

$$M := \{ X \subseteq I \mid |X| < \aleph_0 \},$$

and the type

$$\mathfrak{p} := \{ \{k\} \subseteq x \mid k \in I \},$$

which is finitely satisfiable in $\mathfrak{M}$. By (3), there is an element $[a]_\mathfrak{u}$ of $\mathfrak{M}^\mathfrak{u}$ realising $\mathfrak{p}$. For $k \in I$, we set

$$s_k := \{ i \in I \mid \{k\} \subseteq a_i \} = [\![\{k\} \subseteq a_i]\!].$$

Since $\mathfrak{M}^\mathfrak{u} \vDash \{k\} \subseteq [a]_\mathfrak{u}$ it follows by the Theorem of Łoś that $s_k \in \mathfrak{u}$. Furthermore, each a_i being finite there are only finitely many s_k with $i \in s_k$. Hence, the family $(s_k)_{k \in I}$ witnesses that $\mathfrak{u}$ is regular. □

Proposition 2.14. *Let I be an infinite set, $\mathfrak{u}$ a regular ultrafilter on I, $\kappa := |I|$, and Σ a signature of size $|\Sigma| \leq \kappa$. If $\mathfrak{A}_i$ and $\mathfrak{B}_i$, for $i \in I$, are Σ-structures such that $\mathfrak{A}_i \equiv \mathfrak{B}_i$, for all $i \in I$, then*

$$\prod_{i \in I} \mathfrak{A}_i/\mathfrak{u} \cong^\kappa_{\text{iso}} \prod_{i \in I} \mathfrak{B}_i/\mathfrak{u}.$$

Proof. Below we need our structures to be relational. Therefore, we replace $\mathfrak{A}_i$ and $\mathfrak{B}_i$ by their *relational variants* $\mathfrak{A}_i^*$ and $\mathfrak{B}_i^*$ as follows. Let $\Sigma_{\text{rel}} \subseteq \Sigma$ be the set of relation symbols and $\Sigma_{\text{fun}} \subseteq \Sigma$ the set of function symbols. We replace every function symbol $f \in \Sigma_{\text{fun}}$ of type $\bar{s} \to t$ by a new relation symbol R_f of type $\bar{s}t$. The resulting signature is

$$\Sigma^* := \Sigma_{\text{rel}} \cup \{ R_f \mid f \in \Sigma_{\text{fun}} \}.$$

To every Σ-structure $\mathfrak{M}$, we associate a Σ^*-structure $\mathfrak{M}^*$ by expanding $\mathfrak{M}|_{\Sigma_{\mathrm{rel}}}$ by the graphs

$$R_f^{\mathfrak{M}^*} := \{ \bar{a}b \mid f^{\mathfrak{M}}(\bar{a}) = b \}$$

of the functions $f \in \Sigma_{\mathrm{fun}}$.

Since $\mathfrak{u}$ is regular there exists a sequence $(s_\alpha)_{\alpha<\kappa}$ of sets $s_\alpha \in \mathfrak{u}$ such that, for every $i \in I$, the set $\{ \alpha < \kappa \mid i \in s_\alpha \}$ is finite. Fix an enumeration $\langle \Sigma^*_\alpha, k_\alpha \rangle_{\alpha<\kappa}$ of all pairs $\langle \Sigma^*_\alpha, k_\alpha \rangle$ consisting of finite subsets $\Sigma^*_\alpha \subseteq \Sigma^*$ and $k_\alpha \subseteq \kappa$. For $i \in I$ and $\gamma < \kappa$, set

$$\begin{aligned} \Gamma_i &:= \bigcup \{ \Sigma^*_\alpha \mid i \in s_\alpha \}, \\ K_i &:= \bigcup \{ k_\alpha \mid i \in s_\alpha \}, \\ m_i^\gamma &:= |\{ \alpha \in K_i \mid \alpha \geq \gamma \}|. \end{aligned}$$

We claim that

$$J : \prod_{i \in I} \mathfrak{A}_i/\mathfrak{u} \cong^\kappa_{\mathrm{iso}} \prod_{i \in I} \mathfrak{B}_i/\mathfrak{u},$$

where $J \subseteq \mathrm{pIso}_\kappa(\prod_i \mathfrak{A}_i/\mathfrak{u}, \prod_i \mathfrak{B}_i/\mathfrak{u})$ is the following set of partial isomorphisms $\bar{a} \mapsto \bar{b}$. Let $\bar{a} = (a_\nu)_{\nu<\gamma}$ and $\bar{b} = (b_\nu)_{\nu<\gamma}$ where $\gamma < \kappa$ and $a_\nu = [(a^i_\nu)_{i\in I}]_\mathfrak{u}$ and $b_\nu = [(b^i_\nu)_{i\in I}]_\mathfrak{u}$. Then $\bar{a} \mapsto \bar{b} \in J$ if, and only if,

$$\langle \mathfrak{A}^*_i|_{\Gamma_i}, (a^i_\nu)_{\nu\in K_i} \rangle \cong_{m_i^\gamma} \langle \mathfrak{B}^*_i|_{\Gamma_i}, (b^i_\nu)_{\nu\in K_i} \rangle, \quad \text{for all } i \in I.$$

It is straightforward to check that J is κ-complete and κ-bounded. To show that $\langle\rangle \mapsto \langle\rangle \in J$, note that each Γ_i is finite and relational. Hence, we can use Corollary C4.3.6 to show that

$$\mathfrak{A}^*_i|_{\Gamma_i} \equiv \mathfrak{B}^*_i|_{\Gamma_i} \quad \text{implies} \quad \mathfrak{A}^*_i|_{\Gamma_i} \cong_\omega \mathfrak{B}^*_i|_{\Gamma_i}.$$

It remains to prove that J has the back-and-forth property with respect to itself. By symmetry, it is sufficient to prove the forth property. Let $\bar{a} \mapsto \bar{b} \in J$ and $c = [(c^i)_{i\in I}]_\mathfrak{u} \in \prod_i A_i/\mathfrak{u}$. To find a matching element

$d = [(d^i)_{i \in I}]_{\mathfrak{u}} \in \prod_i B_i / \mathfrak{u}$ we consider each component d_i separately. Let $\bar{a} = (a_\nu)_{\nu<\gamma}$ and $\bar{b} = (b_\nu)_{\nu<\gamma}$ as above. By definition, $\bar{a} \mapsto \bar{b} \in J$ implies that

$$\langle \mathfrak{A}_i^*|_{\Gamma_i}, (a_\nu^i)_{\nu \in K_i} \rangle \cong_{m_i^\gamma} \langle \mathfrak{B}_i^*|_{\Gamma_i}, (b_\nu^i)_{\nu \in K_i} \rangle .$$

If $\gamma \notin K_i$, we take an arbitrary element $d_i \in B_i$. Otherwise, there exists some $d_i \in B_i$ such that

$$\langle \mathfrak{A}_i^*|_{\Gamma_i}, (a_\nu^i)_{\nu \in K_i}, c^i \rangle \cong_{m_i^\gamma - 1} \langle \mathfrak{B}_i^*|_{\Gamma_i}, (b_\nu^i)_{\nu \in K_i}, d^i \rangle .$$

Since $\gamma \in K_i$ implies $m_i^{\gamma+1} = m_i^\gamma - 1$, it follows in both cases that

$$\langle \mathfrak{A}_i^*|_{\Gamma_i}, (a_\nu^i)_{\nu \in K_i}, c^i \rangle \cong_{m_i^{\gamma+1}} \langle \mathfrak{B}_i^*|_{\Gamma_i}, (b_\nu^i)_{\nu \in K_i}, d^i \rangle . \qquad \square$$

We have seen that we can find κ-saturated elementary extensions, for all cardinals κ. For saturated elementary extensions the situation is different. The next results give conditions on when such extensions exist.

Proposition 2.15. *Let T be a countable complete first-order theory with infinite models. The following statements are equivalent:*

(1) *T has a countable saturated model.*

(2) *T has a countable $\aleph_1$-universal model.*

(3) $|S^{\bar{s}}(T)| \leq \aleph_0$, *for all finite tuples $\bar{s}$.*

Proof. (1) ⇒ (2) follows from Theorem 2.9.

(2) ⇒ (3) Let $\mathfrak{M}$ be a countable $\aleph_1$-universal model of T. Each type $\mathfrak{p} \in S^{\bar{s}}(T)$ is realised in some countable model. Hence, it is also realised in $\mathfrak{M}$. Since $\mathfrak{M}$ is countable it follows that $|S^{\bar{s}}(T)| \leq \aleph_0$.

(3) ⇒ (1) First, let us show that $|S^{<\omega}(A)| \leq \aleph_0$, for every finite set A. Let $\bar{a}$ be an enumeration of A and $\bar{t}$ the sorts of $\bar{a}$. For every finite tuple of sorts $\bar{s}$ there exists an injective function $f : S^{\bar{s}}(A) \to S^{\bar{s}\bar{t}}(T)$ sending a type $\mathfrak{p} \in S^{\bar{s}}(A)$ to the type

$$f(\mathfrak{p}) := \{ \varphi(\bar{x}, \bar{y}) \mid \varphi(\bar{x}, \bar{a}) \in \mathfrak{p} \} .$$

Consequently, $|S^{\bar{s}}(A)| \leq |S^{\bar{s}\bar{t}}(T)| \leq \aleph_0$. Since T is countable there are only countably many sorts. Therefore it follows that $S^{<\omega}(A)$ is countable as well.

To find the desired saturated model of T we construct an elementary chain $(\mathfrak{M}_n)_{n<\omega}$ of countable models of T such that each $\mathfrak{M}_{n+1}$ realises every type over a finite subset $A \subseteq \mathfrak{M}_n$. Then the union $\mathfrak{M}_\omega := \bigcup_{n<\omega} \mathfrak{M}_n$ will be the desired countable $\aleph_0$-saturated model of T.

We start with an arbitrary countable model $\mathfrak{M}_0$ of T. Given $\mathfrak{M}_n$ we construct $\mathfrak{M}_{n+1}$ as follows. Let F be the class of all finite subsets of M_n and set $P := \bigcup_{A \in F} S^{<\omega}(A)$. By the above remarks it follows that P is countable. Fix an enumeration $(\mathfrak{p}_k)_{k<\omega}$ of P. Using Lemma C3.5.2 we construct an elementary chain $(\mathfrak{A}_n^k)_{k<\omega}$ of countable structures with $\mathfrak{A}_n^0 := \mathfrak{M}_n$ such that $\mathfrak{p}_k$ is realised in $\mathfrak{A}_n^{k+1}$. Their union $\bigcup_k \mathfrak{A}_n^k$ is the desired structure $\mathfrak{M}_{n+1}$. □

For the existence of uncountable saturated structures we can only give a sufficient condition at the moment. A more precise characterisation will be presented in Theorem **??** below.

Theorem 2.16. *Let T be a complete theory with infinite models. If T is κ-stable, for a regular cardinal $\kappa \geq |T|$, then T has a saturated model of size κ.*

Proof. We construct an elementary chain $(\mathfrak{A}_i)_{i\leq\kappa}$ of models $\mathfrak{A}_i \vDash T$ with $|A_i| = \kappa$. We start with an arbitrary model $\mathfrak{A}_0$ of size κ. For limit ordinals δ, we set $\mathfrak{A}_\delta := \bigcup_{i<\delta} \mathfrak{A}_i$. For the successor step, suppose that we have already defined $\mathfrak{A}_i$. Since T is κ-stable we have $|S^s(A_i)| \leq \kappa$, for all sorts s. Furthermore, there are at most $|T| \leq \kappa$ sorts. Hence, we can use Corollary C3.5.3 to find an elementary extension $\mathfrak{A}_{i+1} \succeq \mathfrak{A}_i$ of size κ that realises every type in $\bigcup_s S^s(A_i)$.

We claim that the limit $\mathfrak{A}_\kappa$ is saturated. It is sufficient to prove that every 1-type over a set $U \subseteq A_\kappa$ of size $|U| < \kappa$ is realised in $\mathfrak{A}_\kappa$. Since κ is regular there exists an index $\alpha < \kappa$ with $U \subseteq A_\alpha$. Consequently, every 1-type over U is realised in $\mathfrak{A}_{\alpha+1} \preceq \mathfrak{A}_\kappa$. □

We conclude this section with a closer look at definable relations in κ-saturated structures. We have already proved in Lemma C5.6.17 that the closure ordinal of a least fixed point on an $\aleph_0$-saturated structure is at most ω.

Lemma 2.17. *Suppose that $\mathfrak{A}$ is κ-saturated and let $\varphi(\bar{x})$ be a first-order formula with $|\bar{x}| < \omega$. Either $|\varphi^{\mathfrak{A}}| < \aleph_0$ or $|\varphi^{\mathfrak{A}}| \geq \kappa$.*

Proof. Suppose that $\varphi^{\mathfrak{A}}$ is infinite. We construct a sequence $(\bar{a}^i)_{i<\kappa}$ of distinct tuples satisfying φ. Suppose that we have already defined $\bar{a}^i$, for $i < \alpha$. The set

$$\Gamma_\alpha(\bar{x}) := \{\varphi(\bar{x})\} \cup \{\, \bar{x} \neq \bar{a}^i \mid i < \alpha \,\}$$

is a partial type since $\varphi^{\mathfrak{A}}$ is infinite. Since $\mathfrak{A}$ is κ-saturated we can therefore find a tuple $\bar{a}^\alpha$ realising $\Gamma_\alpha(\bar{x})$. □

Proposition 2.18. *A first-order theory T admits quantifier elimination if and only if we have*

$$\mathfrak{A} \equiv_0 \mathfrak{B} \quad \textit{implies} \quad \mathfrak{A} \cong_0^{\aleph_0} \mathfrak{B}\,,$$

for all $\aleph_0$-saturated models $\mathfrak{A}, \mathfrak{B}$ of T.

Proof. ($\Leftarrow$) follows from Corollary D1.2.12. For ($\Rightarrow$), note that, according to Theorem D1.2.6, if $\mathfrak{A}$ and $\mathfrak{B}$ are models of T then we have

$$I_0^{\aleph_0}(\mathfrak{A}, \mathfrak{B}) = I_{\mathrm{FO}}^{\aleph_0}(\mathfrak{A}, \mathfrak{B})\,.$$

Furthermore, if $\mathfrak{A}$ and $\mathfrak{B}$ are $\aleph_0$-saturated then we have

$$I_{\mathrm{FO}}^{\aleph_0}(\mathfrak{A}, \mathfrak{B}) = I_\infty^{\aleph_0}(\mathfrak{A}, \mathfrak{B})\,,$$

by Corollary 2.3. Since $\mathfrak{A} \equiv_0 \mathfrak{B}$ implies $\langle\rangle \mapsto \langle\rangle \in I_\infty^{\aleph_0}(\mathfrak{A}, \mathfrak{B})$, it follows that $\mathfrak{A} \cong_0^{\aleph_0} \mathfrak{B}$. □

Proposition 2.19. *If $\mathfrak{A}$ is κ-saturated then so is $\mathcal{I}(\mathfrak{A})$, for every first-order interpretation $\mathcal{I}$.*

Proof. Recall that interpretations are comorphisms, that is, for every formula $\varphi(\bar{x})$, there is a formula $\varphi^{\mathcal{I}}(\bar{x})$ such that

$$\mathcal{I}(\mathfrak{A}) \vDash \varphi(\mathcal{I}(\bar{a})) \quad \text{iff} \quad \mathfrak{A} \vDash \varphi^{\mathcal{I}}(\bar{a}) .$$

Suppose that $\mathfrak{p} \in S^n(U)$ where $U \subseteq \mathcal{I}[A]$ is of size $|U| < \kappa$. Then there is some set $V \subseteq A$ of size $|V| = |U|$ with $U = \mathcal{I}[V]$. Since $\mathfrak{A}$ is κ-saturated we can find a tuple $\bar{a} \in A^n$ realising the partial type

$$\mathfrak{p}^{\mathcal{I}} := \{ \varphi^{\mathcal{I}}(\bar{x}, \bar{c}) \mid \varphi(\bar{x}, \mathcal{I}(\bar{c})) \in \mathfrak{p},\ \bar{c} \subseteq V \}$$

over V. It follows that $\mathcal{I}(\bar{a})$ realises $\mathfrak{p}$. □

3. *Projectively saturated structures*

In a saturated structure every type over sets of a certain size is realised. We can extend this requirement by also including types with *second-order* variables. Structures that realise also all types of this form are called *projectively saturated.*

Definition 3.1. Let Σ and Ξ be disjoint signatures and $T \subseteq \mathrm{FO}^{\circ}[\Sigma]$ a first-order theory.

(a) A *Ξ-type* is a subset $\mathfrak{p} \subseteq \mathrm{FO}^{\circ}[\Sigma \cup \Xi]$ such that $T \cup \mathfrak{p}$ is consistent. $\mathfrak{p}$ is *complete* if $\mathfrak{p} = \mathrm{Th}(\mathfrak{A})$ for some $(\Sigma \cup \Xi)$-structure $\mathfrak{A}$ satisfying T. The set of all complete Ξ-types is denoted by $S^{\Xi}(T)$.

(b) A Σ-structure $\mathfrak{A}$ *realises* a Ξ-type $\mathfrak{p}$ if it has a $(\Sigma \cup \Xi)$-expansion $\mathfrak{A}_+$ with $\mathfrak{A}_+ \vDash \mathfrak{p}$.

(c) We call a structure $\mathfrak{A}$ *projectively κ-saturated* if it realises every $\{\xi\}$-type over a set of less than κ parameters, for all relation symbols and function symbols ξ.

Lemma 3.2. *Every projectively κ-saturated structure is κ-saturated and strongly κ-homogeneous.*

Proof. Let $\mathfrak{M}$ be a projectively κ-saturated Σ-structure.

First, we show that $\mathfrak{M}$ is κ-saturated. Let $A \subseteq M$ be a subset of size $|A| < \kappa$ and let $\mathfrak{p} \in S^n(A)$. We have to find some $\bar{c} \in M^n$ with $\operatorname{tp}(\bar{c}/A) = \mathfrak{p}$. Let $\mathfrak{N}$ be some elementary extension of $\mathfrak{M}$ that realises $\mathfrak{p}$ and fix a tuple $\bar{d} \in N^n$ of type $\mathfrak{p}$. Let $R \notin \Sigma$ be a new n-ary relation symbol and set $R^{\mathfrak{N}} = \{\bar{d}\}$. Since $\mathfrak{M}$ is projectively κ-saturated there exists a relation $R^{\mathfrak{M}}$ such that

$$\langle \mathfrak{M}, R^{\mathfrak{M}}, \bar{a} \rangle \equiv \langle \mathfrak{N}, R^{\mathfrak{N}}, \bar{a} \rangle ,$$

where $\bar{a}$ is some enumeration of A. It follows that $R^{\mathfrak{M}}$ contains exactly one tuple $\bar{c}$ and we have $\operatorname{tp}(\bar{c}/A) = \operatorname{tp}(\bar{d}/A) = \mathfrak{p}$.

It remains to show that $\mathfrak{M}$ is strongly κ-homogeneous. Let $\bar{a}, \bar{b} \in M^\alpha$, for $\alpha < \kappa$, be sequences such that $\langle \mathfrak{M}, \bar{a} \rangle \equiv \langle \mathfrak{M}, \bar{b} \rangle$. Set

$$\begin{aligned} \Phi(f) := {} & \operatorname{Th}(\mathfrak{M}, \bar{a}, \bar{b}) \\ & \cup \{ f a_i = b_i \mid i < \alpha \} \\ & \cup \{ \forall x \exists y f y = x \} \\ & \cup \{ \forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \varphi(f \bar{x})) \mid \varphi \in \mathrm{FO} \} , \end{aligned}$$

where $f \notin \Sigma$ is a new unary function symbol. By Lemma 1.11, we know that $\Phi(f)$ is satisfiable. Hence, $\Phi(f)$ is an $\{f\}$-type over $\bar{a}\bar{b}$ and there exist a function $\pi : M \to M$ such that $\langle \mathfrak{M}, \bar{a}\bar{b} \rangle \vDash \Phi(\pi)$. In particular, π is an automorphism of $\mathfrak{M}$ with $\pi(\bar{a}) = \bar{b}$. □

Theorem 3.3. *Let $\mathfrak{A}$ be a Σ-structure and $\kappa > |\Sigma| \oplus \aleph_0$ a regular cardinal. There exists a projectively κ-saturated elementary extension $\mathfrak{B} \succeq \mathfrak{A}$ of size $|B| \leq |A|^{<\kappa}$.*

Proof. If $\mathfrak{A}$ is finite then it is already projectively κ-saturated, for all κ. Therefore, we may assume that $\mathfrak{A}$ is infinite. Let us write $\mathfrak{C} \sqsubseteq \mathfrak{D}$ if $\mathfrak{D}$ is an expansion of some elementary extension of $\mathfrak{C}$. If $(\mathfrak{C}_i)_{i<\alpha}$ is a $\sqsubseteq$-chain then we can form its union $\bigcup_{i<\alpha} \mathfrak{C}_i$ and, by the same proof as for elementary chains, it follows that $\mathfrak{C}_k \sqsubseteq \bigcup_{i<\alpha} \mathfrak{C}_i$.

Set $\mu := |\Sigma| \oplus \aleph_0$ and $\lambda := (|A| \oplus \mu^+)^{<\kappa}$. Then $\lambda^{<\kappa} = \lambda \geq \kappa$. We will construct a $\sqsubseteq$-chain $(\mathfrak{C}_\alpha)_{\alpha<\lambda\kappa}$ of length $\lambda\kappa$ where the structure $\mathfrak{C}_\alpha$ is of size $|C_\alpha| = \lambda \otimes (\alpha \oplus 1)$. For simplicity, we assume that C_α is the set of ordinals less than $\lambda(\alpha + 1)$. The Σ-reduct of the union $\bigcup_{\alpha<\lambda\kappa} \mathfrak{C}_\alpha$ will be the desired structure $\mathfrak{B} \succeq \mathfrak{A}$. Note that $B = \lambda\kappa$ has size $|B| = \lambda \otimes \kappa = \lambda$.

For every finite tuple $\bar{s}$ of sorts and each sort t fix a new relation symbol $R_{\bar{s}}$ of type $\bar{s}$ and a new function symbol $f_{\bar{s}t}$ of type $\bar{s} \to t$. Let Ξ be the set of these symbols. For $U \subseteq B$ we can consider $T := \mathrm{Th}(\mathfrak{A})$ as an incomplete theory over the signature Σ_U. Hence, we have the type space $S^\Xi(U) := S(\mathrm{FO}[\Sigma_U \cup \Xi]/T)$. Fix an enumeration $(\mathfrak{p}_i)_{i<\lambda\kappa}$ of all $\{\xi\}$-types $\mathfrak{p}_i \in S^{\{\xi\}}(U_i)$, for all possible $\xi \in \Xi$ and all subsets $U_i \subseteq B$ of size $|U_i| < \kappa$. For every $\nu < \kappa$, there are $|B|^\nu = \lambda^\nu \leq \lambda^{<\kappa} = \lambda$ subsets of size ν and $2^{\mu\oplus\nu} \leq \lambda^{<\kappa} = \lambda$ different $\{\xi\}$-types with ν parameters. Therefore, the above enumeration contains $\lambda \otimes \lambda = \lambda$ different types. Consequently, we can choose the sequence $\langle \mathfrak{p}_i \rangle_{i<\lambda\kappa}$ such that, for every $\alpha < \kappa$, each $\{\xi\}$-type $\mathfrak{p}$ appears at least once with some index $\lambda\alpha \leq i < \lambda(\alpha + 1)$. In particular, we assume that every type appears cofinally often in our enumeration.

We start the construction of $(\mathfrak{C}_i)_i$ with an arbitrary elementary extension $\mathfrak{C}_0 \succeq \mathfrak{A}$ of size $|C_0| = \lambda$. For limit ordinals δ, we set $\mathfrak{C}_\delta := \bigcup_{\alpha<\delta} \mathfrak{C}_\alpha$. For the successor step, suppose that $\mathfrak{C}_\alpha$ has already been defined.

If $U_\alpha \nsubseteq C_\alpha = \lambda(\alpha + 1)$ or if $\mathfrak{p}_\alpha$ is inconsistent with $\mathrm{Th}((\mathfrak{C}_\alpha)_{C_\alpha})$ then we choose an arbitrary elementary extension $\mathfrak{C}_{\alpha+1} \succeq \mathfrak{C}_\alpha$ with universe $\lambda(\alpha+2)$. Otherwise, let $\mathfrak{D}$ be a model of $\mathfrak{p}_\alpha \cup \mathrm{Th}((\mathfrak{C}_\alpha)_{C_\alpha})$. By the Theorem of Löwenheim and Skolem we can choose $\mathfrak{D}$ of size $|D| = \lambda$. Hence, we may assume that $D = \lambda(\alpha + 2)$. By construction, we have $\mathfrak{C}_\alpha \sqsubseteq \mathfrak{D}$ and we can set $\mathfrak{C}_{\alpha+1} := \mathfrak{D}$.

This concludes the construction of $(\mathfrak{C}_\alpha)_\alpha$. Let $\mathfrak{D} := \bigcup_{\alpha<\lambda\kappa} \mathfrak{C}_\alpha$. We claim that $\mathfrak{B} := \mathfrak{D}|_\Sigma$ is a projectively κ-saturated elementary extension of $\mathfrak{A}$. Since $\mathfrak{A} \preceq \mathfrak{C}_0 \sqsubseteq \mathfrak{D}$ we have $\mathfrak{A} \preceq \mathfrak{B}$. Let $V \subseteq B$ be a set of size $|V| < \kappa$ and let $\mathfrak{p}$ be a $\{\xi\}$-type over V. We have to find a relation or function $\xi^{\mathfrak{B}}$ such that $\langle \mathfrak{B}_V, \xi^{\mathfrak{B}} \rangle \vDash \mathfrak{p}$. Since $V \subseteq \lambda\kappa$, $|V| < \kappa$, and κ is regular there is some ordinal α such that $V \subseteq \lambda\alpha$. By construction, there is some index i in the range $\lambda\alpha \leq i < \lambda(\alpha+1)$ such that $\mathfrak{p} = \mathfrak{p}_i$ and $V = U_i$. Consequently, $(\mathfrak{C}_{i+1})_{U_i} \vDash \mathfrak{p}_i$ implies $\langle \mathfrak{B}_V, \xi^{\mathfrak{C}_{i+1}} \rangle \vDash \mathfrak{p}$. □

Corollary 3.4. *Let $\kappa \geq |\Sigma| \oplus \aleph_0$. Every Σ-structure $\mathfrak{A}$ has a projectively κ^+-saturated elementary extension of size at most $|A|^\kappa$.*

In the definition of a projectively saturated structure we only require that every type with one free second-order variable is realised. In fact, we can add several relations at the same time.

Proposition 3.5. *Let $\mathfrak{A}$ be a projectively κ-saturated Σ-structure. Then $\mathfrak{A}$ realises every Ξ-type over less than κ parameters with $|\Xi| < \kappa$.*

Proof. Let $\mathfrak{p}$ be a Ξ-type and $\mathfrak{B} \vDash \mathfrak{p}$ a structure of size κ realising $\mathfrak{p}$. Fix an arbitrary bijection $f : B \times B \to B$ and let $(\xi_i)_{i<\alpha}$ be an enumeration of Ξ. We choose α different elements $c_i \in B$, $i < \alpha$. Using the pairing function f we can replace each relation or function ξ_i by a unary relation P_i. Finally, we define a 4-ary relation R by

$$R := \{\, \langle a, a, b, f(a,b) \rangle \mid a, b \in B \,\} \\ \cup \{\, \langle c_i, a, a, b \rangle \mid b \in P_i, a \in B, a \neq c_i \,\} .$$

Note that $\mathfrak{B}$ is definable in the structure $\mathfrak{B}' := \langle \mathfrak{B}|_\Sigma, R, (P_i)_i, (c_i)_i \rangle$. Since $\mathfrak{A}$ is projectively κ-saturated it has an expansion $\mathfrak{A}' \equiv \mathfrak{B}'$. We can apply the definition of $\mathfrak{B}$ in $\mathfrak{B}'$ to the structure $\mathfrak{A}'$ to obtain the desired $(\Sigma \cup \Xi)$-expansion $\mathfrak{A}_+$ of $\mathfrak{A}$ with $\mathfrak{A}_+ \equiv \mathfrak{B}$. □

4. Pseudo-saturated structures

Depending on the model of set theory there can be first-order theories without saturated models. But if we slightly weaken the definition of saturation then we can prove that such models always exist.

Definition 4.1. A structure $\mathfrak{A}$ is *pseudo-saturated,* or *special,* if there exists an elementary chain $(\mathfrak{A}_\kappa)_{\kappa<|A|}$, indexed by cardinals κ, such that $\mathfrak{A} = \bigcup_\kappa \mathfrak{A}_\kappa$ and every $\mathfrak{A}_\kappa$ is κ^+-saturated.

Lemma 4.2. *Every saturated structure is pseudo-saturated.*

Proof. If $\mathfrak{A}$ is saturated then it is κ^+-saturated, for all $\kappa < |A|$. Therefore, we can obtain the desired chain $(\mathfrak{A}_\kappa)_\kappa$ be setting $\mathfrak{A}_\kappa := \mathfrak{A}$ for all κ. □

By a *strong limit cardinal* we mean a cardinal of the form $\beth_\delta$ where δ is either 0 or a limit ordinal.

Theorem 4.3. *Let $\mathfrak{A}$ be an infinite Σ-structure and κ a strong limit cardinal with $\kappa > |A| \oplus |\Sigma|$. Then $\mathfrak{A}$ has a pseudo-saturated elementary extensions of size κ.*

Proof. Suppose that $\kappa = \beth_\delta$. Fix a strictly increasing sequence $(\lambda_i)_{i<\mathrm{cf}\,\delta}$ of cardinals $\lambda_i < \beth_\delta$ such that

$$\beth_\delta = \sup\{\, \lambda_i \mid i < \mathrm{cf}\,\delta \,\} = \sup\{\, 2^{\lambda_i} \mid i < \mathrm{cf}\,\delta \,\}.$$

By removing some elements of this sequence, we may assume that $\lambda_0 > |A| \oplus |\Sigma|$. We construct an elementary chain $(\mathfrak{B}_i)_{i<\mathrm{cf}\,\delta}$ such that

- $\mathfrak{B}_0 = \mathfrak{A}$,
- each $\mathfrak{B}_{i+1}$ is a λ_i^+-saturated structure of size $|B_{i+1}| = 2^{\lambda_i}$, and
- $|B_\gamma| \leq 2^{\lambda_\gamma}$, for limit ordinals γ.

The first structure $\mathfrak{B}_0$ is already defined. If $i = j + 1$ is a successor then $|B_j| \leq 2^{\lambda_j}$ implies that we can apply Corollary 3.4 to find a λ_i^+-saturated elementary extension $\mathfrak{B}_{j+1} \succeq \mathfrak{B}_j$ of size $|B_i| = |B_j|^{\lambda_i} = 2^{\lambda_i}$. Finally, for limit ordinals γ, we can set $\mathfrak{B}_\gamma := \bigcup_{i<\gamma} \mathfrak{B}_i$ since

$$|B_\gamma| = \sup\{\, 2^{\lambda_i} \mid i < \gamma \,\} \leq 2^{\lambda_\gamma}.$$

The structure $\mathfrak{B} := \bigcup_i \mathfrak{B}_i$ is an elementary extension of $\mathfrak{B}_0 = \mathfrak{A}$ of size $|B| = \sup\{\, 2^{\lambda_i} \mid i < \mathrm{cf}\,\delta \,\} = \kappa$. We claim that $\mathfrak{B} := \bigcup_i \mathfrak{B}_i$ is pseudo-saturated. Let g be an increasing function from the set of all cardinals less than κ to the ordinal $\mathrm{cf}\,\delta$ such that $\lambda_{g(\mu)} \geq \mu$, for all $\mu < \kappa$. Then $\mathfrak{B}_{g(\mu)+1}$ is $\lambda_{g(\mu)}^+$-saturated and the chain $(\mathfrak{B}_{g(\mu)+1})_{\mu<\kappa}$ witnesses that $\mathfrak{B}$ is pseudo-saturated. □

Corollary 4.4. *Let $T \subseteq \mathrm{FO}[\Sigma]$ be a consistent first-order theory.*

(a) *T has a pseudo-saturated model.*

(b) *If T has infinite models and $\kappa > |\mathrm{FO}[\Sigma]|$ is a strong limit cardinal then T has a pseudo-saturated model of size κ.*

Proof. (b) By the Theorem of Löwenheim and Skolem T has a model $\mathfrak{A}$ of size $|A| = |\mathrm{FO}[\Sigma]|$. Therefore, we can apply the preceding theorem to obtain a pseudo-saturated elementary extension $\mathfrak{B} \succeq \mathfrak{A}$ of size κ.

(a) If T has infinite models then the claim follows from (b). Otherwise, T has a finite model and every finite structure is saturated. □

Theorem 4.5. *If $\mathfrak{A} \equiv \mathfrak{B}$ are pseudo-saturated structures of the same size $|A| = |B|$ then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. Suppose that $\mathfrak{A} = \bigcup_\kappa \mathfrak{A}_\kappa$ and $\mathfrak{B} = \bigcup_\kappa \mathfrak{B}_\kappa$. Choose subsets $C_\kappa \subseteq A_\kappa$ and $D_\kappa \subseteq B_\kappa$ of size $|C_\kappa| = |D_\kappa| = \kappa$ such that

$$\bigcup_\kappa C_\kappa = A \quad \text{and} \quad \bigcup_\kappa D_\kappa = B\,.$$

By induction on κ, we construct an increasing chain of partial isomorphisms $(p_\kappa)_\kappa$ with $p_\kappa \in I^\kappa_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ such that

$$C_\kappa \subseteq \operatorname{dom} p_\kappa \subseteq A_\kappa \quad \text{and} \quad D_\kappa \subseteq \operatorname{rng} p_\kappa \subseteq B_\kappa\,.$$

The union $p := \bigcup_\kappa p_\kappa$ is the desired isomorphism.

Let $p_0 := \langle\rangle \mapsto \langle\rangle$. If κ is a limit cardinal then we set $p_\kappa := \bigcup_{\lambda<\kappa} p_\lambda$. Since $I^\kappa_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ is κ-complete, we have $p_\kappa \in I^\kappa_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$. Finally, suppose that $\kappa = \lambda^+$ and $p_\lambda = \bar{a} \mapsto \bar{b} \in I^\lambda_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ has already been defined. Let $\bar{c}$ be an enumeration of C_κ and $\bar{d}$ one of D_κ. Since $\mathfrak{A}_\kappa$ and $\mathfrak{B}_\kappa$ are κ^+-saturated, we have

$$\langle \mathfrak{A}, \bar{a} \rangle \cong^{\kappa^+}_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b} \rangle\,.$$

As $|\bar{c}| = |\bar{d}| = \kappa < \kappa^+$ we can apply Lemma c4.4.9 to find sequences $\bar{e} \in (A_\kappa)^\kappa$ and $\bar{f} \in (B_\kappa)^\kappa$ such that

$$\langle \mathfrak{A}, \bar{a}\bar{c}\bar{e} \rangle \cong^{\kappa^+}_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b}\bar{f}\bar{d} \rangle\,.$$

In particular, $p_\kappa := \bar{a}\bar{c}\bar{e} \mapsto \bar{b}\bar{f}\bar{d} \in I^\kappa_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$. □

Lemma 4.6. *Let $\mathfrak{A}$ be a pseudo-saturated Σ-structure of size $|A| = \kappa$.*

(a) *The expansion $\langle \mathfrak{A}, \bar{a} \rangle$ is pseudo-saturated, for every sequence $\bar{a} \in A^\alpha$ of length $\alpha < \operatorname{cf} \kappa$.*

(b) *The reduct $\mathfrak{A}|_\Gamma$ is pseudo-saturated, for every $\Gamma \subseteq \Sigma$.*

Proof. (b) follows immediately from the definition.

(a) Let $\mathfrak{A} = \bigcup_{\lambda<\kappa} \mathfrak{A}_\lambda$ where $\mathfrak{A}_\lambda$ is λ^+-saturated. Since $\alpha < \operatorname{cf} \kappa$ there is some index $\mu < \kappa$ with $\bar{a} \subseteq A_\mu$. It follows that $\langle \mathfrak{A}_\lambda, \bar{a} \rangle$ is λ^+-saturated, for every $\lambda \geq \mu$. Consequently, $\langle \mathfrak{A}, \bar{a} \rangle = \bigcup_{\lambda<\kappa} \langle \mathfrak{A}_{\lambda \oplus \mu}, \bar{a} \rangle$ is pseudo-saturated. □

As an easy corollary of Theorem 4.5 we see that every pseudo-saturated structure $\mathfrak{A}$ is $\operatorname{cf}(|A|)$-homogeneous. In fact, we will show below that it is even projectively $\operatorname{cf}(|A|)$-saturated.

Proposition 4.7. *Every pseudo-saturated structure $\mathfrak{A}$ of size $|A| = \kappa$ is strongly $\operatorname{cf}(\kappa)$-homogeneous.*

Proof. Suppose that $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle$, for $\bar{a}, \bar{b} \in A^\alpha$ with $\alpha < \operatorname{cf} \kappa$. The expansions $\langle \mathfrak{A}, \bar{a} \rangle$ and $\langle \mathfrak{A}, \bar{b} \rangle$ are pseudo-saturated, by Lemma 4.6 (a). Consequently, it follows by Theorem 4.5 that they are isomorphic. □

Every pseudo-saturated structure of size κ is projectively $\operatorname{cf}(\kappa)$-saturated and κ^+-universal. To prove this fact we need some technical lemmas.

Lemma 4.8. *Let $\mathfrak{A}$ be a Σ-structure and $\mathfrak{B}$ a Σ_+-structure with $\Sigma \subseteq \Sigma_+$. If $\mathfrak{A}$ and $\mathfrak{B}$ are pseudo-saturated, $\mathfrak{A} \equiv \mathfrak{B}|_\Sigma$, and $|\Sigma_+| \leq |A| \leq |B|$ then there exists an elementary embedding $h : \mathfrak{A} \to \mathfrak{B}|_\Sigma$ such that the set $\operatorname{rng} h$ induces a substructure of $\mathfrak{B}$.*

Proof. Suppose that $\mathfrak{A} = \bigcup_\lambda \mathfrak{A}_\lambda$ and $\mathfrak{B} = \bigcup_\lambda \mathfrak{B}_\lambda$. Let $(a_\alpha)_{\alpha<\kappa}$ be an enumeration of A such that $a_\alpha \in A_{|\alpha|}$, for all α. We choose a bijection $\tau : \kappa \to T[\Sigma_+, A]$ such that

$$\tau(\alpha) = t(a_{i_0}, \dots, a_{i_{n-1}}) \quad \text{implies} \quad i_0, \dots, i_{n-1} < \alpha\,.$$

To define h we construct an increasing sequence $(p_\alpha)_{\alpha<\kappa}$ of partial elementary maps $p_\alpha \in I_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ such that, for all $\alpha < \kappa$,

- $\operatorname{dom} p_\alpha \subseteq A_{|\alpha|}$ and $\operatorname{rng} p_\alpha \subseteq B_{|\alpha|}$,
- $|p_\alpha| \leq |2\alpha|$,
- $a_\alpha \in \operatorname{dom} p_{\alpha+1}$,
- if $\tau(\alpha) = t(\bar{a})$ then $t^{\mathfrak{B}}[p_\alpha(\bar{a})] \in \operatorname{rng} p_{\alpha+1}$.

The limit $h := \bigcup_\alpha p_\alpha$ will be the desired elementary embedding.

We start the construction with $p_0 := \varnothing$. For limit ordinals δ, we set $p_\delta := \bigcup_{\alpha<\delta} p_\alpha$. For the successor step, suppose that $p_\alpha = \bar{c} \mapsto \bar{d}$ has already been defined. Suppose that $\tau(\alpha) = t(\bar{a})$ and let $y := t^{\mathfrak{B}}[p_\alpha(\bar{a})]$. As $\mathfrak{A}_{|\alpha|}$ is $|\alpha|^+$-saturated there is some element $x \in A_{|\alpha|}$ such that

$$\langle \mathfrak{A}, \bar{c}x \rangle \equiv \langle \mathfrak{B}, \bar{d}y \rangle .$$

Similarly, since $\mathfrak{B}_{|\alpha|}$ is $|\alpha|^+$-saturated we can find an element $z \in B_{|\alpha|}$ with

$$\langle \mathfrak{A}, \bar{c}xa_\alpha \rangle \equiv \langle \mathfrak{B}, \bar{d}yz \rangle .$$

We set $p_{\alpha+1} := \bar{c}xa_\alpha \mapsto \bar{d}yz$. □

Theorem 4.9. *Let $\mathfrak{A}$ be a pseudo-saturated Σ-structure and Ξ a signature disjoint from Σ. If $|A| \geq |\Sigma| \oplus |\Xi|$ then $\mathfrak{A}$ realises every Ξ-type $\mathfrak{p} \in S^{\Xi}(\varnothing)$.*

Proof. Let $\mathfrak{p}^* \subseteq \mathrm{FO}^0[\Gamma]$ be a Skolemisation of $\mathfrak{p}$ and fix a pseudo-saturated model $\mathfrak{B}$ realising $\mathfrak{p}^*$ such that $\mathfrak{B}|_\Sigma \equiv \mathfrak{A}$ and $|B| \geq |A|$. We can use Lemma 4.8 to find exists an elementary embedding $h : \mathfrak{A} \to \mathfrak{B}|_\Sigma$ whose range $B_0 := \operatorname{rng} h$ induces a substructure $\mathfrak{B}_0$ of $\mathfrak{B}$. We define a Γ-expansion $\mathfrak{A}_*$ of $\mathfrak{A}$ by setting

$$\xi^{\mathfrak{A}_*} := h^{-1}[\xi^{\mathfrak{B}_0}] , \quad \text{for } \xi \in \Gamma \setminus \Sigma .$$

It follows that $h : \mathfrak{A}_* \cong \mathfrak{B}_0$. Since $\mathfrak{p}^*$ is a Skolem theory we have $\mathfrak{B}_0 \preceq \mathfrak{B}$. This implies that $\mathfrak{A}_* \cong \mathfrak{B}_0 \vDash \mathfrak{p}^*$. Consequently, $\mathfrak{A}_+ := \mathfrak{A}_*|_{\Sigma \cup \Xi}$ is the desired model of $\mathfrak{p}$. □

Corollary 4.10. *Let $\mathfrak{A}$ be a pseudo-saturated structure of size $|A| = \kappa$ and let Δ be a set of first-order formulae that is closed under conjunctions. If $\mathfrak{B}$ is any structure of size $|B| \leq \kappa$ with $\mathfrak{B} \leq_{\exists\Delta} \mathfrak{A}$ then there exists a Δ-embedding $\mathfrak{B} \to \mathfrak{A}$.*

Proof. Let $\Phi := \mathrm{Th}_\Delta(\mathfrak{B}_B)$. If we can show that $\Phi \cup \mathrm{Th}(\mathfrak{A})$ is consistent then we can use Theorem 4.9 to find an expansion $\mathfrak{A}_C$ of $\mathfrak{A}$ satisfying Φ. Hence, the Diagram Lemma implies that there exists a Δ-embedding $\mathfrak{B} \to \mathfrak{A}$.

It remains to prove that $\Phi \cup \mathrm{Th}(\mathfrak{A})$ is consistent. Suppose otherwise. Then there are finitely many formulae $\varphi_0(\bar{b}_0), \dots, \varphi_{n-1}(\bar{b}_{n-1}) \in \Phi$ with parameters $\bar{b}_i \subseteq B$ such that

$$\mathrm{Th}(\mathfrak{A}) \vDash \neg\varphi_0(\bar{b}_0) \vee \dots \vee \neg\varphi_{n-1}(\bar{b}_{n-1}) .$$

Since Φ is closed under conjunction we may assume w.l.o.g. that $n = 1$. Consequently,

$$\mathfrak{A} \vDash \neg\exists\bar{x}\varphi_0(\bar{x}) .$$

But $\mathfrak{B} \vDash \exists\bar{x}\varphi_0(\bar{x})$ and $\mathfrak{B} \leq_{\exists\Delta} \mathfrak{A}$ implies that $\mathfrak{A} \vDash \exists\bar{x}\varphi_0(\bar{x})$. Contradiction. □

Theorem 4.11. *A pseudo-saturated structure of size κ is κ^+-universal and projectively $\mathrm{cf}(\kappa)$-saturated.*

Proof. Let $\mathfrak{A}$ be pseudo-saturated. If $\mathfrak{B} \equiv \mathfrak{A}$ is a structure of size $|B| \leq \kappa$ then we can use Corollary 4.10 to find an elementary embedding $\mathfrak{B} \to \mathfrak{A}$. Consequently, $\mathfrak{A}$ is κ^+-universal.

For the second claim suppose that $\bar{a} \in A^\alpha$ is a sequence of $\alpha < \mathrm{cf}\,\kappa$ elements. Then $\langle\mathfrak{A}, \bar{a}\rangle$ is pseudo-saturated by Lemma 4.6 (a). It follows by Theorem 4.9 that $\langle\mathfrak{A}, \bar{a}\rangle$ is projectively 1-saturated. Consequently, $\mathfrak{A}$ is projectively $\mathrm{cf}(\kappa)$-saturated. □

Corollary 4.12. *If $\mathfrak{A}$ is pseudo-saturated and $|A|$ is regular then $\mathfrak{A}$ is saturated.*

Corollary 4.13. *Every saturated structure of size κ is projectively κ-saturated.*

Proof. Suppose that $\mathfrak{A}$ is saturated. Then so is $\langle \mathfrak{A}, \bar{a} \rangle$, for every $\bar{a} \in A^{<\kappa}$. Since saturated structures are pseudo-saturated it follows that every expansion $\langle \mathfrak{A}, \bar{a} \rangle$ by less than κ constants is projectively 1-saturated. Consequently, $\mathfrak{A}$ is projectively κ-saturated. □

We conclude this section with a few results about definable relations in pseudo-saturated and projectively saturated structures. We start with an analogue of Lemma 2.17.

Lemma 4.14. *Suppose that $\mathfrak{A}$ is pseudo-saturated and let $\varphi(\bar{x}, \bar{c})$ be a first-order formula with parameters $\bar{c} \subseteq A$ where $|\bar{x}| < \omega$. Then $\varphi(\bar{x}, \bar{c})^{\mathfrak{A}}$ is either finite or $|\varphi(\bar{x}, \bar{c})^{\mathfrak{A}}| = |A|$.*

Proof. Suppose that $\mathfrak{A} = \bigcup_\lambda \mathfrak{A}_\lambda$. If $\varphi^{\mathfrak{A}}$ is infinite then, by Lemma 2.17, we have $|\varphi^{\mathfrak{A}_\lambda}| \geq \lambda^+$. Consequently,

$$|\varphi^{\mathfrak{A}}| \geq |\varphi^{\mathfrak{A}_\lambda}| \geq \lambda^+ , \quad \text{for all } \lambda < |A| ,$$

implies that $|\varphi^{\mathfrak{A}}| = |A|$. □

Lemma 4.15. *If $\mathfrak{A}$ is pseudo-saturated then so is $\mathcal{I}(\mathfrak{A})$, for every first-order interpretation $\mathcal{I}$.*

Proof. Suppose that $\mathfrak{A} = \bigcup_\kappa \mathfrak{A}_\kappa$ where each $\mathfrak{A}_\kappa$ is κ^+-saturated. Note that

$$\mathfrak{A}_\kappa \preceq \mathfrak{A}_\lambda \quad \text{implies} \quad \mathcal{I}(\mathfrak{A}_\kappa) \preceq \mathcal{I}(\mathfrak{A}_\lambda) , \quad \text{for } \kappa \leq \lambda .$$

Hence, the structures $\mathcal{I}(\mathfrak{A}_\kappa)$ form an elementary chain with limit

$$\bigcup_{\kappa < |A|} \mathcal{I}(\mathfrak{A}_\kappa) = \mathcal{I}(\mathfrak{A}) .$$

Furthermore, according to Proposition 2.19, each structure $\mathcal{I}(\mathfrak{A}_\kappa)$ is κ^+-saturated. Hence, $\mathcal{I}(\mathfrak{A})$ is pseudo-saturated. □

Lemma 4.16. *Let $\mathcal{I}$ be a first-order interpretation from Σ to Γ and let $\kappa > |\Sigma| \oplus |\Gamma|$ be a cardinal. If $\mathfrak{A}$ is projectively κ-saturated then so is $\mathcal{I}(\mathfrak{A})$.*

Proof. Let $\bar{a} \subseteq \mathcal{I}(A)$ be a sequence of less than κ-parameters and suppose that $\mathfrak{p}$ is a $\{\xi\}$-type over $\bar{a}$. We can find parameters $\bar{c} \subseteq A$ and an interpretation $\mathcal{J}$ with $\mathcal{J}(\mathfrak{A}, \bar{c}) = \langle \mathcal{I}(\mathfrak{A}), \bar{a} \rangle$. Replacing $\mathfrak{A}$ by $\langle \mathfrak{A}, \bar{c} \rangle$ and $\mathcal{I}$ by $\mathcal{J}$ we can therefore simplify notation by omitting the parameters.

To show that $\mathfrak{p}$ is realised in $\mathcal{I}(\mathfrak{A})$ fix a $(\Gamma \cup \{\xi\})$-structure $\mathfrak{B} \vDash \mathfrak{p}$ realising $\mathfrak{p}$. Let λ be a strong limit cardinal with $\lambda > |\Sigma| \oplus |\Gamma|$ and choose pseudo-saturated structures $\mathfrak{A}_+$ and $\mathfrak{B}_+$ of size λ such that $\mathfrak{A}_+ \equiv \mathfrak{A}$ and $\mathfrak{B}_+ \equiv \mathfrak{B}$. Then $\mathcal{I}(\mathfrak{A}_+) \equiv \mathfrak{B}_+|_\Gamma$ implies, by Theorem 4.5, that $\mathcal{I}(\mathfrak{A}_+) \cong \mathfrak{B}_+|_\Gamma$. Let $\xi^{\mathcal{I}(\mathfrak{A}_+)}$ be the relation on $\mathcal{I}(\mathfrak{A}_+)$ induced by this isomorphism and let $\xi^{\mathfrak{A}_+}$ be its preimage under $\mathcal{I}$. Similarly, for every $\zeta \in \Gamma$, let $\zeta^{\mathfrak{A}_+}$ be the preimage of $\zeta^{\mathcal{I}(\mathfrak{A}_+)}$ under $\mathcal{I}$. W.l.o.g. assume that Σ and Γ are disjoint. Let $\mathfrak{A}_*$ be the $(\Sigma \cup \Gamma \cup \{\xi\})$-expansion of $\langle \mathfrak{A}_+, \xi^{\mathfrak{A}_+} \rangle$ by all these relations and functions $\zeta^{\mathfrak{A}_+}$. We can extend $\mathcal{I}$ to an interpretation $\mathcal{J}$ with

$$\mathcal{J}(\mathfrak{A}_*) = \left\langle \mathcal{I}(\mathfrak{A}_+), \xi^{\mathcal{I}(\mathfrak{A}_+)} \right\rangle .$$

Since $\kappa > |\Sigma| \oplus |\Gamma|$ we can use Proposition 3.5 to find a $(\Sigma \cup \Gamma \cup \{\xi\})$-expansion $\mathfrak{A}'$ of $\mathfrak{A}$ with $\mathfrak{A}' \equiv \mathfrak{A}_*$. It follows that $\mathcal{J}(\mathfrak{A}')$ is an $(\Gamma \cup \{\xi\})$-expansion of $\mathcal{I}(\mathfrak{A})$ with $\mathcal{J}(\mathfrak{A}') \equiv \mathcal{J}(\mathfrak{A}_*) \equiv \mathfrak{B}_+ \equiv \mathfrak{B}$. □

E2. *Definability and automorphisms*

1. *Definability in projectively saturated models*

As an application of the notions introduced in the previous chapter we study the relationship between definable relations and automorphisms.

Definition 1.1. Let L be an algebraic logic, $\mathfrak{M}$ a structure, and $U \subseteq M$ a set of parameters.

(a) A tuple $\bar{a} \subseteq M$ is *L-definable over U* if there is an L-formula $\varphi(\bar{x}; \bar{c})$ with parameters $\bar{c} \subseteq U$ such that $\varphi(\bar{x}; \bar{c})^{\mathfrak{M}} = \{\bar{a}\}$.

(b) The *L-definitional closure* of U is the set

$$\mathrm{dcl}_L(U) := \{ a \in M \mid a \text{ is } L\text{-definable over } U \} .$$

The set U is *L-definitional closed* if it is a fixed point of dcl_L.

(c) We say that an L-formula $\varphi(\bar{x}; \bar{c})$ with parameters $\bar{c} \subseteq M$ is *algebraic* if $\varphi(\bar{x}; \bar{c})^{\mathfrak{M}}$ is finite. An L-type $\mathfrak{p}$ is *algebraic* if it implies an algebraic formula.

We call a tuple $\bar{a} \subseteq M$ *L-algebraic over U* if there is an algebraic L-formula $\varphi(\bar{x}; \bar{c})$ with parameters $\bar{c} \subseteq U$ such that $\mathfrak{M} \vDash \varphi(\bar{a}; \bar{c})$.

(d) The *L-algebraic closure* of U is the set

$$\mathrm{acl}_L(U) := \{ a \in M \mid a \text{ is } L\text{-algebraic over } U \} .$$

The set U is *L-algebraically closed* if it is a fixed point of acl_L.

(e) For $L = \mathrm{FO}$ we simply say that $\bar{a}$ is *definable* or *algebraic* over U and we write $\mathrm{dcl}(U)$ and $\mathrm{acl}(U)$ without the index L.

Lemma 1.2. *Let $\mathfrak{M}$ be a structure. The operators $\mathrm{dcl}_{\mathrm{FO}}$ and $\mathrm{acl}_{\mathrm{FO}}$ are closure operators on M with finite character.*

Proof. Every element $a \in U$ is definable over U by the formula $x = a$. Consequently, $U \subseteq \mathrm{dcl}_{\mathrm{FO}}(U) \subseteq \mathrm{acl}_{\mathrm{FO}}(U)$.

If a is definable or algebraic over U by the formula $\varphi(x; \bar{c})$, the same formula can be used to show that a is definable or algebraic over any set $V \supseteq \bar{c}$. Consequently, $U \subseteq V$ implies $\mathrm{dcl}(U) \subseteq \mathrm{dcl}(V)$ and $\mathrm{acl}(U) \subseteq \mathrm{acl}(V)$. Furthermore, it follows that $a \in \mathrm{dcl}(\bar{c})$ or $a \in \mathrm{acl}(\bar{c})$, respectively. Hence, these operators have finite character.

Finally, suppose that a is definable over $\mathrm{dcl}(U)$. Let $\varphi(x; \bar{c}, \bar{d})$ be the corresponding formula where $\bar{d} \subseteq U$ and $\bar{c} \subseteq \mathrm{dcl}(U) \setminus U$. For every element c_i, there is a formula ψ_i over U with $\psi_i^{\mathbb{M}} = \{c_i\}$. We can define a over U by the formula

$$\varphi'(x; \bar{d}) := \exists \bar{y} \Big[\bigwedge_i \psi_i(y_i) \wedge \varphi(x; \bar{y}, \bar{d}) \Big] .$$

The proof for acl is analogous. Suppose that a is algebraic over $\mathrm{acl}(U)$ and let $\varphi(x; \bar{c}, \bar{d})$ be the formula witnessing this fact where $\bar{d} \subseteq U$ and $\bar{c} \subseteq \mathrm{acl}(U) \setminus U$. For every element c_i, fix a formula ψ_i over U such that $\psi_i^{\mathfrak{M}}$ is a finite set containing c_i. Set $m := |\varphi(x, \bar{c}, \bar{d})^{\mathfrak{M}}|$. The following formula shows that a is algebraic over U.

$$\varphi'(x; \bar{d}) := \exists \bar{y} \Big[\bigwedge_i \psi_i(y_i) \wedge \vartheta(\bar{y}) \wedge \varphi(x; \bar{y}, \bar{d}) \Big] ,$$

where

$$\vartheta(\bar{y}) := \forall z_0 \cdots \forall z_m \Big[\bigwedge_i \varphi(z_i; \bar{y}, \bar{d}) \to \bigvee_{i<k} z_i = z_k \Big]$$

states that there are at most m elements z satisfying $\varphi(z; \bar{y}, \bar{d})$. □

For strongly κ-homogeneous structures there is a tight relationship between types and automorphisms.

Lemma 1.3. *Let $\mathfrak{M}$ be strongly κ-homogeneous and $U \subseteq M$ a set of size $|U| < \kappa$. For $\bar{a}, \bar{b} \in M^{<\kappa}$, the following statements are equivalent:*

(1) $\operatorname{tp}(\bar{a}/U) = \operatorname{tp}(\bar{b}/U)$

(2) *There is some automorphism* $\pi \in \operatorname{Aut}\mathfrak{M}$ *with*

$$\pi \upharpoonright U = \mathrm{id}_U \quad \textit{and} \quad \pi(\bar{a}) = \bar{b}\,.$$

Proof. (1) ⇒ (2) follows from the definition of a strongly κ-homogeneous structure, while (2) ⇒ (1) follows from the fact that isomorphisms preserve first-order formulae. □

As a consequence we can express the definitional closure and the algebraic closure in terms of automorphisms.

Definition 1.4. Let $\mathfrak{M}$ be a structure and $U \subseteq M$.

(a) Let ξ and ζ be two tuples or two relations in M. We say that ζ is a *conjugate* of ξ *over* U if ξ is mapped to ζ by an automorphism of $\mathfrak{M}$ that fixes U pointwise.

For a sets of formulae Φ and Ψ we similarly say that Ψ is a *conjugate* of Φ over U if there exists an automorphism π fixing U pointwise such that

$$\Psi = \big\{\, \varphi(\bar{x}; \pi(\bar{c})) \;\big|\; \varphi(\bar{x}; \bar{c}) \in \Phi \,\big\}\,.$$

(b) We define the following two closure operators on M:

$$\begin{aligned} \mathrm{dcl}_{\mathrm{Aut}}(U) &:= \{\, a \in M \mid a \text{ has exactly one conjugate over } U \,\}\,, \\ \mathrm{acl}_{\mathrm{Aut}}(U) &:= \{\, a \in M \mid a \text{ has only finitely many conjugates} \\ &\qquad\qquad \text{over } U \,\}\,. \end{aligned}$$

Exercise 1.1. Let $\mathfrak{M}$ be a structure. Prove that $\mathrm{dcl}_{\mathrm{Aut}}$ and $\mathrm{acl}_{\mathrm{Aut}}$ are closure operators on M.

Example. Let $\mathfrak{V}$ be a vector space and let $U \subseteq V$. Then

$$\mathrm{dcl}_{\mathrm{Aut}}(U) = \langle\!\langle U \rangle\!\rangle_{\mathfrak{V}}\,.$$

Remark. Let $\mathfrak{M}$ be a structure and $U \subseteq M$. We can write the pointwise stabiliser of U in $\operatorname{Aut}\mathfrak{M}$ and its setwise stabiliser as

$$(\operatorname{Aut}\mathfrak{M})_{(U)} = \operatorname{Aut}\mathfrak{M}_U \quad \text{and} \quad (\operatorname{Aut}\mathfrak{M})_{\{U\}} = \operatorname{Aut}\langle\mathfrak{M}, U\rangle\,.$$

In arbitrary structures the relationship between dcl_L and $\mathrm{dcl}_{\mathrm{Aut}}$ and between acl_L and $\mathrm{acl}_{\mathrm{Aut}}$ is as follows.

Lemma 1.5. *Let L be an algebraic logic, $\mathfrak{M}$ a structure, and $U \subseteq M$.*

(a) $\mathrm{dcl}_L(U) \subseteq \mathrm{dcl}_{\mathrm{Aut}}(U)$

(b) $\mathrm{acl}_L(U) \subseteq \mathrm{acl}_{\mathrm{Aut}}(U)$

Proof. (a) If there is an automorphism π with $\pi \upharpoonright U = \mathrm{id}_U$ and $\pi(a) = b$, for $a \neq b$, then

$$\mathfrak{M} \vDash \varphi(a;\bar{c}) \leftrightarrow \varphi(b;\bar{c})\,,$$

for all L-formulae φ and all parameters $\bar{c} \subseteq U$. Consequently, a is not L-definable over U.

(b) Similarly, if the orbit of a under $\operatorname{Aut}\mathfrak{M}_U$ is infinite then every formula satisfied by a is also satisfied by infinitely many other elements. Hence, a is not L-algebraic over U. □

For sufficiently saturated structures the two closure operators coincide.

Theorem 1.6. *Let $\mathfrak{M}$ be κ-saturated and strongly κ-homogeneous, $a \in M$ an element, and let $U \subseteq M$ be a set of size $|U| < \kappa$.*

(a) *The following statements are equivalent:*

  (1) $a \in \mathrm{dcl}_{\mathrm{FO}}(U)$

  (2) $a \in \mathrm{dcl}_{\mathrm{Aut}}(U)$

  (3) $\mathrm{tp}(a/U)$ *has exactly one realisation in* $\mathfrak{M}$.

(b) *The following statements are equivalent:*

  (1) $a \in \mathrm{acl}_{\mathrm{FO}}(U)$

  (2) $a \in \mathrm{acl}_{\mathrm{Aut}}(U)$

(3) $\operatorname{tp}(a/U)$ *has only finitely many realisations in* $\mathfrak{M}$.

Proof. (a) (2) ⇔ (3) follows by Lemma 1.3.

(1) ⇒ (3) Fix a formula $\varphi(x)$ over U that defines a. Since $\varphi \in \operatorname{tp}(a/U)$, it follows that a is the only realisation of $\operatorname{tp}(a/U)$.

(3) ⇒ (1) Suppose that $a \notin \operatorname{dcl}_{\mathrm{FO}}(U)$. It follows that, for every finite set Φ of first-order formulae over U, there is some element $b \neq a$ such that

$$\mathfrak{M} \vDash \bigwedge \Phi(a) \leftrightarrow \bigwedge \Phi(b)\,.$$

By the Compactness Theorem and the fact that $\mathfrak{M}$ is κ-saturated, it follows that we can find some element $b \neq a$ with

$$\operatorname{tp}(a/U) = \operatorname{tp}(b/U)\,.$$

(b) (2) ⇔ (3) follows by Lemma 1.3.

(1) ⇒ (3) Fix a formula $\varphi(x)$ over U such that $\varphi^{\mathfrak{M}}$ is a finite set containing a. Since $\varphi \in \operatorname{tp}(a/U)$ it follows that there are at most $|\varphi^{\mathfrak{M}}|$ realisations of $\operatorname{tp}(a/U)$.

(3) ⇒ (1) We can use an analogous argument as in (a) to show that $a \notin \operatorname{acl}_{\mathrm{FO}}(U)$ implies that there are infinitely many realisations of $\operatorname{tp}(a/U)$. □

Corollary 1.7. *Let* $\mathfrak{M}$ *be a structure and* $U \subseteq M$. *Then*

$$\pi[\operatorname{acl}(U)] = \operatorname{acl}(U)\,, \quad \textit{for all } \pi \in \operatorname{Aut} \mathfrak{M}_U\,.$$

Proof. Let $a \in \operatorname{acl}(U)$. To show that $\pi(a) \in \operatorname{acl}(U)$ we consider the set $A \subseteq M$ of all realisations of $\operatorname{tp}(a/U)$. By Theorem 1.6, A is a finite set with $A \subseteq \operatorname{acl}(U)$. Consequently, $\pi(a) \in A \subseteq \operatorname{acl}(U)$. □

Corollary 1.8. *Let* $\mathfrak{M}$ *be* κ*-saturated and strongly* κ*-homogeneous, and let* $A, B \subseteq M$ *be sets of size* $|A|, |B| < \kappa$.

(a) *The following statements are equivalent:*

(1) $A \subseteq \operatorname{dcl}(B)$

(2) $\mathrm{dcl}(A) \subseteq \mathrm{dcl}(B)$
(3) $\mathrm{Aut}\,\mathfrak{M}_A \supseteq \mathrm{Aut}\,\mathfrak{M}_B$.

(b) *The following statements are equivalent:*

(1) $A \subseteq \mathrm{dcl}(B)$ *and* $B \subseteq \mathrm{dcl}(A)$
(2) $\mathrm{dcl}(A) = \mathrm{dcl}(B)$
(3) $\mathrm{Aut}\,\mathfrak{M}_A = \mathrm{Aut}\,\mathfrak{M}_B$.

Proof. (b) follows from (a).

(a) (1) ⇔ (2) Clearly, $\mathrm{dcl}(A) \subseteq \mathrm{dcl}(B)$ implies $A \subseteq \mathrm{dcl}(A) \subseteq \mathrm{dcl}(B)$. Conversely, $A \subseteq \mathrm{dcl}(B)$ implies $\mathrm{dcl}(A) \subseteq \mathrm{dcl}(\mathrm{dcl}(B)) = \mathrm{dcl}(B)$.

(1) ⇒ (3) Suppose that $A \subseteq \mathrm{dcl}(B)$ and let $\pi \in \mathrm{Aut}\,\mathfrak{M}_B$. Then it follows by Theorem 1.6 and definition of $\mathrm{dcl}_{\mathrm{Aut}}(B)$ that

$$\pi(a) = a\,, \quad \text{for all } a \in \mathrm{dcl}_{\mathrm{Aut}}(B) = \mathrm{dcl}(B) \supseteq A\,.$$

Hence, $\pi \in \mathrm{Aut}\,\mathfrak{M}_A$.

(3) ⇒ (1) Suppose that $\mathrm{Aut}\,\mathfrak{M}_A \supseteq \mathrm{Aut}\,\mathfrak{M}_B$ and let $a \in A$. Then $a \in \mathrm{dcl}_{\mathrm{Aut}}(A)$ implies that

$$\pi(a) = a\,, \quad \text{for all } \pi \in \mathrm{Aut}\,\mathfrak{M}_A\,.$$

In particular, we have

$$\pi(a) = a\,, \quad \text{for all } \pi \in \mathrm{Aut}\,\mathfrak{M}_B\,.$$

By Theorem 1.6 and definition of $\mathrm{dcl}_{\mathrm{Aut}}(B)$, it follows that

$$a \in \mathrm{dcl}_{\mathrm{Aut}}(B) = \mathrm{dcl}(B)\,. \qquad \square$$

As an application of Theorem 1.6, we present the following characterisation of the algebraic closure.

Lemma 1.9. *Let $\mathfrak{M}$ be a Σ-structure that is κ-saturated and strongly κ-homogeneous, for some cardinal $\kappa > |\Sigma|$, and let $U \subseteq M$ be a set of size $|U| < \kappa$. Then*

$$\mathrm{acl}(U) = \bigcap \{\, A \mid \mathfrak{A} \preceq \mathfrak{M} \text{ with } U \subseteq A \,\}\,.$$

Proof. (⊆) Let $\mathfrak{A} \preceq \mathfrak{M}$ be an elementary substructure containing U. To show that $\operatorname{acl}(U) \subseteq A$, consider an element $a \in \operatorname{acl}(U)$. There exists an algebraic formula $\varphi(x)$ over U with $a \in \varphi^{\mathfrak{M}}$. Let $m := |\varphi^{\mathfrak{M}}|$. Then

$$\mathfrak{M} \vDash \exists^m x \varphi(x) \quad \text{implies} \quad \mathfrak{A} \vDash \exists^m x \varphi(x) .$$

Since $\varphi^{\mathfrak{A}} \subseteq \varphi^{\mathfrak{M}}$ it follows that $\varphi^{\mathfrak{A}} = \varphi^{\mathfrak{M}}$. Hence, $a \in \varphi^{\mathfrak{A}} \subseteq A$.

(⊇) Suppose that $a \notin \operatorname{acl}(U)$. We have to find an elementary substructure $\mathfrak{A} \preceq \mathfrak{M}$ containing U such that $a \notin A$. By Theorem 1.6 and the fact that $\mathfrak{M}$ is κ-saturated, there exists a sequence $(b_\alpha)_{\alpha<\kappa}$ of distinct elements such that

$$\operatorname{tp}(b_\alpha/U) = \operatorname{tp}(a/U) , \quad \text{for all } \alpha < \kappa .$$

Using the Theorem of Löwenheim and Skolem, we can find an elementary substructure $\mathfrak{A}_0 \preceq \mathfrak{M}$ containing U with

$$|A_0| \leq |U| \oplus |\Sigma| < \kappa .$$

There exists an index $\alpha < \kappa$ with $b_\alpha \notin A_0$. Since $\mathfrak{M}$ is strongly κ-homogeneous, we can find an automorphism π with $\pi \restriction U = \mathrm{id}_U$ and $\pi(b_\alpha) = a$. Set $\mathfrak{A} := \pi[\mathfrak{A}_0]$. Then $\mathfrak{A} \preceq \mathfrak{M}$ contains U but not a. □

After considering the definability of single elements we now study the relationship between automorphisms and definable relations. Our first result gives a characterisation of those relations that are definable over a set U of parameters.

Lemma 1.10. *Suppose that $\mathfrak{M}$ is κ-saturated and strongly κ-homogeneous and let $U \subseteq M$ be a set of size $|U| < \kappa$. An M-definable relation $R \subseteq M^n$ is U-definable if, and only if, $\pi[R] = R$, for all $\pi \in \operatorname{Aut} \mathfrak{M}_U$.*

Proof. Clearly, a U-definable relation is invariant under all automorphisms of $\mathfrak{M}$ that fix U pointwise. For the converse, suppose that R is defined by the formula $\varphi(\bar{x}; \bar{c})$ with $\bar{c} \subseteq M$. Consider the set

$$\begin{aligned} \Phi := {} & \{\varphi(\bar{x}; \bar{c}) \wedge \neg\varphi(\bar{x}'; \bar{c})\} \\ & \cup \{\, \psi(\bar{x}) \leftrightarrow \psi(\bar{x}') \mid \psi \text{ a formula over } U \,\} . \end{aligned}$$

If $\Phi(\bar{x}, \bar{x}') \cup \mathrm{Th}(\mathfrak{M}_M)$ is satisfiable then Φ is a partial type and, since $\mathfrak{M}$ is κ-saturated, there are elements $\bar{a}, \bar{b} \in M^n$ satisfying Φ. Let $\pi_0 : U \cup \bar{a} \to U \cup \bar{b}$ be the function with $\pi_0 \restriction U = \mathrm{id}_U$ and $\pi_0(\bar{a}) = \bar{b}$. By choice of $\bar{a}$ and $\bar{b}$ this is an elementary partial function. Since $\mathfrak{M}$ is strongly κ-homogeneous, we can extend it to an automorphism $\pi : M \to M$. But we have $\bar{a} \in \varphi^{\mathfrak{M}} = R$ and $\pi(\bar{a}) = \bar{b} \notin \varphi^{\mathfrak{M}} = R$. Hence, R is not invariant under automorphisms of $\mathrm{Aut}\,\mathfrak{M}_U$. A contradiction.

Consequently, $\Phi \cup \mathrm{Th}(\mathfrak{M}_M)$ is not satisfiable. Hence, there are finitely many formulae $\psi_0, \dots, \psi_{m-1}$ over U such that

$$\mathfrak{M} \vDash \forall \bar{x} \forall \bar{x}' \Big[\bigwedge_i [\psi_i(\bar{x}) \leftrightarrow \psi_i(\bar{x}')] \to [\varphi(\bar{x}; \bar{c}) \leftrightarrow \varphi(\bar{x}'; \bar{c})] \Big].$$

For $I \subseteq [m]$, define

$$\chi_I(\bar{x}) := \bigwedge_{i \in I} \psi_i(\bar{x}) \wedge \bigwedge_{i \notin I} \neg\psi_i(\bar{x}),$$

and let

$$S := \big\{ I \subseteq [m] \bigm| \mathfrak{M} \vDash \chi_I(\bar{a}) \text{ for some } \bar{a} \in R \big\}.$$

It follows that

$$\bar{a} \in R \quad \text{iff} \quad \mathfrak{M} \vDash \bigvee_{I \in S} \chi_I(\bar{a}).$$

Consequently, the formula $\bigvee_{I \in S} \chi_I(\bar{x})$ defines R over U. □

An analogous result for relations with finitely many conjugates will be given in Lemma 3.11 below.

If the structure $\mathfrak{M}$ is even projectively saturated, we can drop the assumption that the relation R is M-definable. In particular, the following result implies that FO has the Beth property.

Theorem 1.11. *Let Σ, Ξ be disjoint signatures, $\kappa > |\Xi|$, and $T \subseteq \mathrm{FO}^0[\Sigma]$ a first-order theory. For a complete Ξ-type $\mathfrak{p} \in S^{\Xi}(T)$ and a relation symbol $R \in \Xi$, the following statements are equivalent:*

(1) *There is an* $\mathrm{FO}^{<\omega}[\Sigma]$*-formula* $\varphi(\bar{x})$ *such that*

$$\mathfrak{p} \vDash \forall \bar{x}[R\bar{x} \leftrightarrow \varphi(\bar{x})] .$$

(2) *If* $\mathfrak{M}$ *is a model of* T *and* $\mathfrak{N}_0, \mathfrak{N}_1$ *are realisations of* $\mathfrak{p}$ *in* $\mathfrak{M}$ *then* $R^{\mathfrak{N}_0} = R^{\mathfrak{N}_1}$.

(3) *There is a model* $\mathfrak{M}$ *of* T *which is either projectively* κ*-saturated, or saturated and of cardinality at least* $|\Sigma \cup \Xi|$*, such that*

$$R^{\mathfrak{N}_0} = R^{\mathfrak{N}_1} , \quad \text{for every pair } \mathfrak{N}_0, \mathfrak{N}_1 \text{ of realisations of } \mathfrak{p} \text{ in } \mathfrak{M}.$$

(4) *There is a model* $\mathfrak{M}$ *of* T *which is either projectively* κ*-saturated, or saturated and of cardinality at least* $|\Sigma \cup \Xi|$*, such that*

$$\pi[R^{\mathfrak{M}_+}] = R^{\mathfrak{M}_+} , \quad \text{for every realisation } \mathfrak{M}_+ \text{ of } \mathfrak{p} \text{ in } \mathfrak{M} \text{ and each automorphism } \pi \in \operatorname{Aut} \mathfrak{M}.$$

Proof. The implications (1) ⇒ (2) and (3) ⇒ (4) are trivial. (2) ⇒ (3) is also trivial, except for the existence of $\mathfrak{M}$ which follows by Corollary E1.3.4.

(4) ⇒ (1) The proof is similar to that of the preceding lemma. Let $\bar{s}$ be the type of R. We choose new constant symbols $\bar{c}$ and $\bar{d}$ and we set

$$\Phi := \mathfrak{p} \cup \{R\bar{c}, \neg R\bar{d}\} \cup \{ \psi(\bar{c}) \leftrightarrow \psi(\bar{d}) \mid \psi \in \mathrm{FO}^{\bar{s}}[\Sigma] \} .$$

If Φ is inconsistent, there are finitely many formulae $\psi_0, \dots, \psi_{m-1} \in \mathrm{FO}^{\bar{s}}[\Sigma]$ such that

$$\mathfrak{p} \vDash \forall \bar{x}\bar{y}\Big[\bigwedge_{i<m} [\psi_i(\bar{x}) \leftrightarrow \psi_i(\bar{y})] \to (R\bar{x} \leftrightarrow R\bar{y})\Big] .$$

As above we define

$$\chi_I(\bar{x}) := \bigwedge_{i \in I} \psi_i(\bar{x}) \wedge \bigwedge_{i \notin I} \neg\psi_i(\bar{x}) , \quad \text{for } I \subseteq [m] .$$

For every $I \subseteq [m]$, it follows that we either have

$$\mathfrak{p} \vDash \chi_I(\bar{x}) \to R\bar{x} \quad \text{or} \quad \mathfrak{p} \vDash \chi_I(\bar{x}) \to \neg R\bar{x} \,.$$

Consequently, we can define R by the formula

$$\varphi(\bar{x}) := \bigvee_{I \in S} \chi_I(\bar{x}) \quad \text{where} \quad S := \{ I \subseteq [m] \mid \mathfrak{p} \vDash \chi_I(\bar{x}) \to R\bar{x} \} \,.$$

It remains to consider the case where Φ has a model $\mathfrak{A}$. We claim that this is impossible. Since $\mathfrak{p}$ is complete it follows that $\mathfrak{A}|_\Sigma \equiv \mathfrak{M}|_\Sigma$. Consequently, we can use Proposition E1.3.5 to expand $\mathfrak{M}|_\Sigma$ to a model $\mathfrak{M}^+$ of Φ. Let $\bar{a}$ and $\bar{b}$ be the values of the constants $\bar{c}$ and $\bar{d}$ in $\mathfrak{M}^+$, respectively. Then

$$\langle \mathfrak{M}|_\Sigma, \bar{a} \rangle \equiv \langle \mathfrak{M}|_\Sigma, \bar{b} \rangle \,.$$

Since $\mathfrak{M}|_\Sigma$ is strongly $\aleph_0$-homogeneous it follows that there is some automorphism $\pi \in \operatorname{Aut} \mathfrak{M}|_\Sigma$ with $\pi(\bar{a}) = \bar{b}$. But $\bar{a} \in R^{\mathfrak{M}^+}$ and $\pi(\bar{a}) = \bar{b} \notin R^{\mathfrak{M}^+}$ contradicts our choice of $\mathfrak{M}$. □

Corollary 1.12. *Let Σ, Ξ be disjoint signatures, $R \in \Xi$ a relation symbol, and $T \subseteq \mathrm{FO}^0[\Sigma]$ a complete first-order theory. If $\mathfrak{p} \in S^\Xi(T)$ is a complete Ξ-type such that, for every realisation $\mathfrak{M}$ of $\mathfrak{p}$ and all automorphisms $\pi \in \operatorname{Aut} \mathfrak{M}|_\Sigma$, we have*

$$\pi[R^\mathfrak{M}] = R^\mathfrak{M},$$

then there is an $\mathrm{FO}^{<\omega}[\Sigma]$-formula $\varphi(\bar{x})$ such that

$$\mathfrak{p} \vDash \forall \bar{x}[R\bar{x} \leftrightarrow \varphi(\bar{x})] \,.$$

Proof. Since T has a projectively $|\Xi|^+$-saturated model, the claim follows from Theorem 1.11. □

Corollary 1.13. *Let Σ, Ξ be disjoint signatures, $R \in \Xi$ a relation symbol, and $T \subseteq \mathrm{FO}^0[\Sigma]$ a first-order theory. If $\mathfrak{p}$ is a Ξ-type such that, for every realisation $\mathfrak{M}$ of $\mathfrak{p}$ and all automorphisms $\pi \in \operatorname{Aut} \mathfrak{M}|_\Sigma$, we have*

$$\pi[R^{\mathfrak{M}}] = R^{\mathfrak{M}},$$

then there are finitely many formulae $\varphi_0(\bar{x}), \dots, \varphi_{n-1}(\bar{x}) \in \mathrm{FO}^{<\omega}[\Sigma]$ such that

$$\mathfrak{p} \vDash \bigvee_{i<n} \forall \bar{x}[R\bar{x} \leftrightarrow \varphi_i(\bar{x})] .$$

Proof. If $\mathfrak{q} \supseteq \mathfrak{p}$ is a complete Ξ-type, we can use the preceding corollary to find a formula $\varphi_{\mathfrak{q}}(\bar{x})$ defining R modulo $\mathfrak{q}$. Consequently,

$$\mathfrak{p} \vDash \bigvee \{ R\bar{x} \leftrightarrow \varphi_{\mathfrak{q}}(\bar{x}) \mid \mathfrak{q} \supseteq \mathfrak{p} \text{ complete} \} .$$

By compactness, it follows that there are finitely many complete types $\mathfrak{q}_0, \dots, \mathfrak{q}_{n-1} \supseteq \mathfrak{p}$ with

$$\mathfrak{p} \vDash \bigvee_{i<n} [R\bar{x} \leftrightarrow \varphi_{\mathfrak{q}_i}(\bar{x})] .$$

□

Below we will frequently work in projectively saturated elementary extensions of a given model. In order to simplify the presentation and to avoid having to include phrases like 'there exists an elementary extension such that', it turned out to be a good idea to fix such an extension once and for all. If this structure is sufficiently saturated, we can use the Amalgamation Theorem and Theorem E1.2.9 to embed all other models we consider into it.

Thus, let us fix a projectively κ-saturated model $\mathbb{M}$ of T where κ is some very large cardinal. We call $\mathbb{M}$ the *monster model* of T. All models $\mathfrak{M}$ of T we will consider are tacitly assumed to be elementary substructures of $\mathbb{M}$ of size $|M| < \kappa$.

We call a relation $R \subseteq \mathbb{M}^n$ *small* if $|R| < \kappa$. Otherwise, it is *large*. To distinguish small and large relations we denote the latter by blackboard bold symbols $\mathbb{A}, \mathbb{B}, \mathbb{C}, \dots$. Note that, by Lemma E1.2.17, definable relations are either finite or large. Mostly, we will only consider types $\mathfrak{p} \in S^{\bar{s}}(U)$ over small sets U of parameters. Note that every such type is realised in $\mathbb{M}$. Similarly, we will tacitly assume that all parameter-definable relations are defined over a small set of parameters.

To simplify notation, we will drop the model $\mathbb{M}$ and write just $\bar{a} \equiv_U \bar{b}$ instead of $\langle \mathbb{M}_U, \bar{a} \rangle \equiv \langle \mathbb{M}_U, \bar{b} \rangle$. By Lemma 1.3, it follows that $\bar{a} \equiv_U \bar{b}$ if, and only if, there exists a U-automorphism π of $\mathbb{M}$ mapping $\bar{a}$ to $\bar{b}$. We extend this notation to sequences of sets $A_0, \dots, A_n, B_0, \dots, B_n \subseteq \mathbb{M}$ by defining

$$A_0 \dots A_n \equiv_U B_0 \dots B_n$$

if there are enumerations $\bar{a}_i$ of A_i and $\bar{b}_i$ of B_i such that

$$\operatorname{tp}(\bar{a}_0 \dots \bar{a}_n / U) = \operatorname{tp}(\bar{b}_0 \dots \bar{b}_n / U) .$$

2. *Imaginary elements and canonical parameters*

In this section we present a construction adding to a given structure new elements representing all definable relations. More generally, we add elements for every class of a definable equivalence relation.

Definition 2.1. Let $\mathfrak{M}$ be an S-sorted structure. An *equivalence formula* is a formula $\chi(\bar{x}, \bar{y})$ without parameters defining an equivalence relation on $M^{\bar{s}}$, for some $\bar{s} \in S^{<\omega}$. The tuple $\bar{s}$ is called the *type* of χ. We denote the equivalence class of a tuple $\bar{a} \in M^{\bar{s}}$ by $[\bar{a}]_\chi$. The elements of the quotient $M^{\bar{s}}/\chi^{\mathfrak{M}}$ are called *imaginary elements.*

Given $\mathfrak{M}$ we construct a new structure $\mathfrak{M}^{\mathrm{eq}}$ by adding all imaginary elements.

Definition 2.2. Let $\mathfrak{M}$ be an S-sorted Σ-structure.

(a) Set

$$S^{\mathrm{eq}} := \{ \chi \mid \chi \text{ an equivalence formula} \} ,$$
$$\Sigma^{\mathrm{eq}} := \Sigma \cup \{ p_\chi \mid \chi \in S^{\mathrm{eq}} \} .$$

We regard S as a subset of S^{eq} via the identification of $s \in S$ with the formula $(x = y) \in S^{eq}$, where x and y are variables of sort s.

We construct an S^{eq}-sorted Σ^{eq}-structure $\mathfrak{M}^{eq}$ as follows. For every equivalence formula χ of type $\bar{s}$, the domain of sort χ is

$$M^{eq}_{\chi} := M^{\bar{s}} / \chi^{\mathfrak{M}}.$$

By the identification of $s \in S$ with $(x = y) \in S^{eq}$, we obtain an embedding of M into M^{eq}. We interpret the symbols of $\Sigma \subseteq \Sigma^{eq}$ in $\mathfrak{M}^{eq}$ according to this embedding. The new function symbols p_{χ} are interpreted as the canonical projections $M^{\bar{s}} \to M^{\bar{s}} / \chi^{\mathfrak{M}}$.

(b) To avoid ambiguities we denote the definable closure and the algebraic closure of a subset $U \subseteq M^{eq}$ by $\mathrm{dcl}^{eq}(U)$ and $\mathrm{acl}^{eq}(U)$, respectively, while $\mathrm{dcl}(U)$ and $\mathrm{acl}(U)$ are the closures of U in the original structure $\mathfrak{M}$.

Remark. (a) Every finite tuple $\bar{a} \in M^{\bar{s}}$ is encoded in $\mathfrak{M}^{eq}$ as a single element $[\bar{a}]_{\chi} \in M^{eq}$ of sort

$$\chi(\bar{x}, \bar{y}) := x_0 = y_0 \wedge \dots \wedge x_{n-1} = y_{n-1} ,$$

where the variables x_i and y_i have sort s_i.

(b) For each formula $\varphi(\bar{x})$, we can define the equivalence formula

$$\chi(\bar{x}, \bar{y}) := \varphi(\bar{x}) \leftrightarrow \varphi(\bar{y}) .$$

There are two imaginary elements of sort χ: one representing $\varphi^{\mathfrak{M}}$, the other one representing $\neg\varphi^{\mathfrak{M}}$. Consequently, $\mathfrak{M}^{eq}$ contains imaginary elements for all relations definable without parameters.

The next proposition shows that, when considering the logical properties of a structure, the transition from $\mathfrak{M}$ to $\mathfrak{M}^{eq}$ does not change much. But we will see below that, when studying automorphisms, this construction allows us in certain cases to replace setwise stabilisers by pointwise ones.

Proposition 2.3. *Let $\mathfrak{M}$ be a structure.*

(a) *$\mathfrak{M}$ is a relativised reduct of $\mathfrak{M}^{\mathrm{eq}}$.*

(b) *There exists a first-order interpretation mapping $\mathfrak{M}$ to $\mathfrak{M}^{\mathrm{eq}}$.*

(c) *For every formula $\varphi(\bar{x}) \in \mathrm{FO}^{\bar{s}}[\Sigma^{\mathrm{eq}}]$, we can construct a formula $\varphi'(\bar{x}) \in \mathrm{FO}^{\bar{s}}[\Sigma]$ such that*

$$\mathfrak{M}^{\mathrm{eq}} \vDash \varphi(\bar{a}) \quad \text{iff} \quad \mathfrak{M} \vDash \varphi'(\bar{a}) , \quad \textit{for all } \bar{a} \in M^{\bar{s}} .$$

(d) *$\mathfrak{A} \equiv \mathfrak{B}$ implies $\mathfrak{A}^{\mathrm{eq}} \equiv \mathfrak{B}^{\mathrm{eq}}$.*

(e) *$M^{\mathrm{eq}} = \langle\!\langle M \rangle\!\rangle_{\mathfrak{M}^{\mathrm{eq}}}$.*

(f) *Every element of M^{eq} is definable over M.*

(g) *Every elementary embedding $g : \mathfrak{A} \to \mathfrak{B}$ can be extended to an elementary embedding $\mathfrak{A}^{\mathrm{eq}} \to \mathfrak{B}^{\mathrm{eq}}$ in a unique way.*

(h) *The restriction map*

$$\rho : \mathfrak{Aut}\,\mathfrak{M}^{\mathrm{eq}} \to \mathfrak{Aut}\,\mathfrak{M} : \pi \mapsto \pi \restriction M$$

is a group isomorphism.

(i) *For every $U \subseteq M$, we have*

$$\mathrm{dcl}(U) = \mathrm{dcl}^{\mathrm{eq}}(U) \cap M \quad \textit{and} \quad \mathrm{acl}(U) = \mathrm{acl}^{\mathrm{eq}}(U) \cap M .$$

Proof. (a) and (b) follow immediately from the definition of $\mathfrak{M}^{\mathrm{eq}}$.

(c) and (d) follow from (b) via Lemma C1.5.9 and Corollary C1.5.13, respectively.

(e) Every imaginary element $[\bar{a}]_\chi \in M^{\mathrm{eq}}$ is denoted by a term $p_\chi \bar{a}$ with parameters $\bar{a} \subseteq M$.

(f) follows immediately from (e).

(g) Let $g : \mathfrak{A} \to \mathfrak{B}$ be an elementary embedding. It follows by (b) and Lemma C2.2.10 that the map $[\bar{a}]_\chi \mapsto [g(\bar{a})]_\chi$ is an elementary embedding $\mathfrak{A}^{\mathrm{eq}} \to \mathfrak{B}^{\mathrm{eq}}$ extending g. For uniqueness, suppose that there are elementary embeddings $h_0, h_1 : \mathfrak{A}^{\mathrm{eq}} \to \mathfrak{B}^{\mathrm{eq}}$ with $h_0 \restriction A = h_1 \restriction A$. By Theorem B3.1.9, it follows that $h_0 \restriction \langle\!\langle A \rangle\!\rangle_{\mathfrak{A}^{\mathrm{eq}}} = h_1 \restriction \langle\!\langle A \rangle\!\rangle_{\mathfrak{A}^{\mathrm{eq}}}$. Hence, (e) implies that $h_0 = h_1$.

(h) First, note that ρ is well-defined since it follows by Lemma C2.2.10 and (a) that, for all $\pi \in \operatorname{Aut} \mathfrak{M}^{\mathrm{eq}}$, the restriction $\pi \restriction M$ is indeed an automorphism of $\mathfrak{M}$. Furthermore, ρ is obviously a group homomorphism. Hence, it remains to show that it is bijective. For surjectivity, note that, by (b), every automorphism of $\mathfrak{M}$ can be extended to one of $\mathfrak{M}^{\mathrm{eq}}$. For injectivity, note that, by (g), every automorphism of $\mathfrak{M}$ can be extended to at most one of $\mathfrak{M}^{\mathrm{eq}}$.

(i) To see that $\operatorname{acl}(U) \subseteq \operatorname{acl}^{\mathrm{eq}}(U)$ note that, if there is a formula φ over U defining a finite set X in $\mathfrak{M}$ then the same formula can be used to define X in $\mathfrak{M}^{\mathrm{eq}}$. For the converse, suppose that φ is a formula over U defining a finite set $X \subseteq M$ in $\mathfrak{M}^{\mathrm{eq}}$. By (c), we can find a formula φ' over U defining the same set in $\mathfrak{M}$. The claim for the definable closure is proved analogously. □

According to the preceding proposition, the first-order theory of $\mathfrak{M}^{\mathrm{eq}}$ only depends on the theory of $\mathfrak{M}$. Consequently, we can extend the operation $^{\mathrm{eq}}$ to theories.

Definition 2.4. For a complete first-order theory T, we denote the theory $\operatorname{Th}(\mathbb{M}^{\mathrm{eq}})$ by T^{eq}.

It also follows that adding imaginary elements does not change the structure of the type spaces.

Corollary 2.5. *Let $U \subseteq \mathbb{M}^{\mathrm{eq}}$ and $U_0 \subseteq \mathbb{M}$ be sets.*

$$\operatorname{dcl}^{\mathrm{eq}}(U) = \operatorname{dcl}^{\mathrm{eq}}(U_0) \quad \textit{implies} \quad \mathfrak{S}^{\bar{s}}(T^{\mathrm{eq}}(U)) \cong \mathfrak{S}^{\bar{s}}(T(U_0))\,.$$

Proof. Since $\operatorname{dcl}^{\mathrm{eq}}(U) = \operatorname{dcl}^{\mathrm{eq}}(U_0)$, it follows by Proposition 2.3 and Lemma C3.3.4 that $\mathrm{FO}^{\bar{s}}[\Sigma_{U_0}]/T(U_0)$ is a retract of $\mathrm{FO}^{\bar{s}}[\Sigma^{\mathrm{eq}}_U]/T^{\mathrm{eq}}(U)$. Consequently, the claim follows by Corollary C3.3.3. □

As a consequence, many logical properties of $\mathfrak{M}$ and T transfer to $\mathfrak{M}^{\mathrm{eq}}$ and T^{eq}. We give two examples.

Lemma 2.6. *Let T be a complete first-order theory, $\mathfrak{M}$ a structure, and κ an infinite cardinal.*

(a) $\mathfrak{M}$ *is* κ*-saturated if, and only if,* $\mathfrak{M}^{\mathrm{eq}}$ *is* κ*-saturated.*

(b) T *is* κ*-stable if, and only if,* T^{eq} *is* κ*-stable.*

Proof. (a) We have seen in Proposition E1.2.19 that κ-saturation is preserved under interpretations.

(b) ($\Leftarrow$) Suppose that T^{eq} is κ-stable. To show that T is κ-stable, consider a set $U \subseteq \mathbb{M}$ of size $|U| \leq \kappa$. By Corollary 2.5, we have

$$\mathfrak{S}^{\bar{s}}(T(U)) \cong \mathfrak{S}^{\bar{s}}(T^{\mathrm{eq}}(U)) .$$

Consequently, $|S^{\bar{s}}(T(U))| = |S^{\bar{s}}(T^{\mathrm{eq}}(U))| \leq \kappa$.

($\Rightarrow$) Suppose that T is κ-stable and let $U \subseteq \mathbb{M}^{\mathrm{eq}}$ be a set of size $|U| \leq \kappa$. There exists a set $C \subseteq \mathbb{M}$ of size $|C| \leq |U| \oplus \aleph_0 \leq \kappa$ with $U \subseteq \mathrm{dcl}^{\mathrm{eq}}(C)$. By Corollary 2.5, we have

$$\mathfrak{S}^{\bar{s}}(T(C)) \cong \mathfrak{S}^{\bar{s}}(T^{\mathrm{eq}}(U \cup C)) .$$

Consequently, $|S^{\bar{s}}(T^{\mathrm{eq}}(U))| \leq |S^{\bar{s}}(T^{\mathrm{eq}}(U \cup C))| = |S^{\bar{s}}(T(C))| \leq \kappa$. □

We have seen that the operation of adding imaginary elements is well-behaved. But what do we gain by it? As an example, consider the following problem. Suppose that a relation $\mathbb{R}$ is defined by a formula $\varphi(\bar{x};\bar{c})$ with parameters $\bar{c}$. There might be many other parameters $\bar{d}$ such that $\varphi(\bar{x};\bar{d})$ defines the same relation $\mathbb{R}$. Sometimes, we would like the parameter $\bar{c}$ to be unique. Using imaginary elements, this can be done. We start by defining the equivalence formula

$$\chi(\bar{y},\bar{y}') := \forall\bar{x}\big[\varphi(\bar{x};\bar{y}) \leftrightarrow \varphi(\bar{x};\bar{y}')\big] .$$

Then two tuples $\bar{a}$ and $\bar{b}$ are equivalent if $\varphi(\bar{x};\bar{a})$ and $\varphi(\bar{x};\bar{b})$ define the same relation. Consequently, the tuples in $[\bar{c}]_\chi$ are precisely those defining $\mathbb{R}$. The imaginary element $e := [\bar{c}]_\chi$ is a unique representative of this set. We obtain a formula

$$\psi(\bar{x};z) := \exists y\big[\varphi(\bar{x};\bar{y}) \wedge p_\chi\bar{y} = z\big]$$

such that e is the unique element such that $\psi(\bar{x};e)$ defines $\mathbb{R}$. Let us formalise this construction.

Definition 2.7. Let $\varphi(\bar{x};\bar{y})$ be a formula.

(a) The *parameter equivalence* for φ is the formula

$$\chi(\bar{y},\bar{y}') := \forall\bar{x}\big[\varphi(\bar{x};\bar{y}) \leftrightarrow \varphi(\bar{x};\bar{y}')\big].$$

(b) A tuple $\bar{c}$ is a *canonical parameter* of a relation $\mathbb{R}$ if there exists a formula $\psi(\bar{x};\bar{y})$ such that $\bar{c}$ is the unique tuple satisfying

$$\psi(\bar{x};\bar{c})^{\mathbb{M}} = \mathbb{R}.$$

In this case, we call the formula $\psi(\bar{x};\bar{c})$ a *canonical definition* of $\mathbb{R}$.

In this terminology we can state the above remark as follows.

Lemma 2.8. *Let χ be the parameter equivalence of a formula $\varphi(\bar{x};\bar{y})$. For every tuple $\bar{c}$, the imaginary element $[\bar{c}]_\chi \in \mathbb{M}^{\mathrm{eq}}_\chi$ is a canonical parameter of $\varphi(\bar{x};\bar{c})^{\mathbb{M}}$.*

Proof. The formula

$$\psi(\bar{x};[\bar{c}]_\chi) := \exists\bar{y}\big[\varphi(\bar{x};\bar{y}) \wedge p_\chi\bar{y} = [\bar{c}]_\chi\big]$$

is a canonical definition of $\varphi(\bar{x};\bar{c})^{\mathbb{M}}$. □

Corollary 2.9. *Every relation $\mathbb{R} \subseteq \mathbb{M}^{\bar{s}}$ that is definable over a set $U \subseteq \mathbb{M}$ has a canonical parameter $e \in \mathrm{dcl}^{\mathrm{eq}}(U)$.*

Thus, all parameter-definable relations $\mathbb{R} \subseteq \mathbb{M}^{\bar{s}}$ have canonical parameters in $\mathbb{M}^{\mathrm{eq}}$. We will see in Corollary 2.12 below that the same is true for parameter-definable relations in $\mathbb{M}^{\mathrm{eq}}$. The reason for this is that performing the operation $^{\mathrm{eq}}$ twice does not offer any additional benefit: according to the following proposition there exist, for every sort $\chi \in (S^{\mathrm{eq}})^{\mathrm{eq}}$, a sort $\eta \in S^{\mathrm{eq}}$ and a definable bijection $(M^{\mathrm{eq}})^{\mathrm{eq}}_\chi \to M^{\mathrm{eq}}_\eta$. Hence, every doubly imaginary element is already present as a singly imaginary one.

Proposition 2.10. *For every equivalence formula $\chi(\bar{x}, \bar{y})$ with type $\bar{\zeta} \in (S^{\mathrm{eq}})^n$, there exist a sort $\eta \in S^{\mathrm{eq}}$ and a definable, surjective function*

$$f : (M^{\mathrm{eq}})^{\bar{\zeta}} \to (M^{\mathrm{eq}})^{\eta}$$

such that $\ker f = \chi^{\mathfrak{M}^{\mathrm{eq}}}$.

Proof. Each sort $\zeta_i \in S^{\mathrm{eq}}$ is itself an equivalence formula of some type $\bar{s}_i \in S^{<\omega}$. We set

$$\eta(\bar{x}_0 \dots \bar{x}_{n-1}, \bar{y}_0 \dots \bar{y}_{n-1}) := \chi\big(p_{\zeta_0}\bar{x}_0, \dots, p_{\zeta_{n-1}}\bar{x}_{n-1},\ p_{\zeta_0}\bar{y}_0, \dots, p_{\zeta_{n-1}}\bar{y}_{n-1}\big)\,.$$

Then $\eta \in S^{\mathrm{eq}}$ is an equivalence formula of type $\bar{s}_0 \dots \bar{s}_{n-1}$. We claim that the desired function $f : (M^{\mathrm{eq}})^{\bar{\zeta}} \to (M^{\mathrm{eq}})^{\eta}$ is defined by the formula

$$\varphi(\bar{x}, y) := \exists \bar{z}_0 \cdots \exists \bar{z}_{n-1} \Big[\bigwedge_{i<n} x_i = p_{\zeta_i}\bar{z}_i \wedge p_{\eta}\bar{z}_0 \dots \bar{z}_{n-1} = y \Big]\,.$$

Note that

$$\mathfrak{M}^{\mathrm{eq}} \vDash \varphi(\bar{\alpha}, b)$$

if, and only if, there are tuples $\bar{a}_0, \dots, \bar{a}_{n-1}$ such that

$$\bar{\alpha} = \langle [\bar{a}_0]_{\zeta_0}, \dots, [\bar{a}_{n-1}]_{\zeta_{n-1}} \rangle \quad \text{and} \quad b = [\bar{a}_0 \dots \bar{a}_{n-1}]_{\eta}\,.$$

Since the equivalence class $[\bar{a}_0 \dots \bar{a}_{n-1}]_{\eta}$ does not depend on the particular choice of representatives $\bar{a}_i \in [\bar{a}_i]_{\zeta_i}$, the element b is uniquely determined by $\bar{\alpha}$. Thus, φ defines a function $f : (M^{\mathrm{eq}})^{\bar{\zeta}} \to (M^{\mathrm{eq}})^{\eta}$.

To see that f is surjective, note that, for every element $[\bar{a}_0 \dots \bar{a}_{n-1}]_{\eta} \in (M^{\mathrm{eq}})^{\eta}$, we have

$$[\bar{a}_0 \dots \bar{a}_{n-1}]_{\eta} = f([\bar{a}_0]_{\zeta_0}, \dots, [\bar{a}_{n-1}]_{\zeta_{n-1}}) \in \operatorname{rng} f\,.$$

Hence, it remains to compute the kernel. Let $\bar\alpha, \bar\alpha' \in (M^{eq})^{\bar\zeta}$ and suppose that $\bar\alpha = \langle [\bar a_0]_{\zeta_0}, \dots, [\bar a_{n-1}]_{\zeta_{n-1}} \rangle$ and $\bar\alpha' = \langle [\bar a'_0]_{\zeta_0}, \dots, [\bar a'_{n-1}]_{\zeta_{n-1}} \rangle$. Then

$$\begin{aligned} f(\bar\alpha) = f(\bar\alpha') \quad &\text{iff} \quad \mathfrak{M}^{eq} \vDash \exists y[\varphi(\bar\alpha, y) \wedge \varphi(\bar\alpha', y)] \\ &\text{iff} \quad [\bar a_0 \dots \bar a_{n-1}]_\eta = [\bar a'_0 \dots \bar a'_{n-1}]_\eta \\ &\text{iff} \quad \mathfrak{M}^{eq} \vDash \eta(\bar a_0 \dots \bar a_{n-1}, \bar a'_0 \dots \bar a'_{n-1}) \\ &\text{iff} \quad \mathfrak{M}^{eq} \vDash \chi(\bar\alpha, \bar\alpha'). \end{aligned}$$ □

We obtain the following generalisation of Lemma 2.8.

Corollary 2.11. *Let $\mathfrak{M}$ be a structure. For every formula $\varphi(\bar x; \bar y)$, there exists a formula $\psi(\bar x; \bar z)$ such that, for every tuple $\bar b \subseteq M^{eq}$, there is a unique tuple $\bar c \subseteq M^{eq}$ with*

$$\varphi(\bar x; \bar b)^{\mathfrak{M}^{eq}} = \psi(\bar x; \bar c)^{\mathfrak{M}^{eq}}.$$

Proof. Let $\varphi(\bar x; \bar y)$ be a formula with parameter equivalence $\chi(\bar y, \bar y')$. According to Proposition 2.10 there exists a definable and surjective function $f : (M^{eq})^{\bar\zeta} \to (M^{eq})^\eta$ such that $\ker f = \chi^{\mathfrak{M}}$. We claim that the formula

$$\psi(\bar x; \bar z) := \exists \bar y[\varphi(\bar x; \bar y) \wedge f(\bar y) = \bar z]$$

has the desired properties.

We start by proving that $\varphi(\bar x; \bar b)^{\mathfrak{M}^{eq}} = \psi(\bar x; \bar c)^{\mathfrak{M}^{eq}}$ where $\bar c := f(\bar b)$. Clearly, every tuple satisfying $\varphi(\bar x; \bar b)$ also satisfies $\psi(\bar x; \bar c)$. Conversely, suppose that $\bar a$ satisfies $\psi(\bar x; \bar c)$. Then there is some tuple $\bar b' \in f^{-1}(\bar c)$ such that $\bar a \in \varphi(\bar x; \bar b')^{\mathfrak{M}^{eq}}$. By definition of f, it follows that $\bar b' \in [\bar b]_\chi$. Hence, $\varphi(\bar x; \bar b')^{\mathfrak{M}^{eq}} = \varphi(\bar x; \bar b)^{\mathfrak{M}^{eq}}$. Consequently, $\bar a$ satisfies $\varphi(\bar x; \bar b)$.

It remains to show that $\bar c$ is unique. Hence, suppose that $\bar c'$ is some tuple with $\varphi(\bar x; \bar b)^{\mathfrak{M}^{eq}} = \psi(\bar x; \bar c')^{\mathfrak{M}^{eq}}$. As f is surjective, there exists an element $\bar b' \in f^{-1}(\bar c')$. Since

$$\varphi(\bar x; \bar b')^{\mathfrak{M}^{eq}} = \psi(\bar x; \bar c')^{\mathfrak{M}^{eq}} = \varphi(\bar x; \bar b)^{\mathfrak{M}^{eq}},$$

it follows that $\mathfrak{M} \vDash \chi(\bar b, \bar b')$. Consequently, $\bar c' = f(\bar b') = f(\bar b) = \bar c$. □

Corollary 2.12. *Every parameter-definable relation in* $\mathbb{M}^{eq}$ *has a canonical parameter.*

3. *Galois bases*

We can characterise canonical parameters also in a more algebraic way via automorphisms.

Definition 3.1. A *Galois base,* or *canonical base,* of a relation $\mathbb{R} \subseteq \mathbb{M}^{\bar{s}}$ is a set $B \subseteq \mathbb{M}$ such that

$$\pi[\mathbb{R}] = \mathbb{R} \quad \text{iff} \quad \pi \upharpoonright B = \mathrm{id}_B\,, \quad \text{for all } \pi \in \mathrm{Aut}\,\mathbb{M}\,.$$

Remark. According to the definition, B is a Galois base or $\mathbb{R}$ if, and only if, in $\mathrm{Aut}\,\mathbb{M}$ the setwise stabiliser of $\mathbb{R}$ coincides with the pointwise stabiliser of B, i.e., if $\mathrm{Aut}\langle\mathbb{M}, \mathbb{R}\rangle = \mathrm{Aut}\,\mathbb{M}_B$.

From the results of Section 1 it follows that, for parameter-definable relations, Galois bases are the same as canonical parameters. But note that the notion of a Galois base also applies to relations that are not definable. Before giving the proof, let us present some technical lemmas. The first one is an immediate consequence of Lemma 1.10.

Lemma 3.2. *If B is a Galois base of a parameter-definable relation* $\mathbb{R}$*, then* $\mathbb{R}$ *is definable over B.*

Lemma 3.3. *Let* $\mathbb{R} \subseteq \mathbb{M}^{\bar{s}}$ *be a relation and* $B \subseteq \mathbb{M}$ *a set. The following statements are equivalent:*

(1) *B is a Galois base of* $\mathbb{R}$ *in the structure* $\mathbb{M}$.

(2) *B is a Galois base of* $\mathbb{R}$ *in the structure* $\mathbb{M}^{eq}$.

Proof. As the restriction map $\pi \mapsto \pi \upharpoonright M$ is an isomorphism between $\mathfrak{Aut}\,\mathbb{M}^{eq}$ and $\mathfrak{Aut}\,\mathbb{M}$, the following two statements are equivalent:

- $\pi[\mathbb{R}] = \mathbb{R} \quad \Leftrightarrow \quad \pi \upharpoonright B = \mathrm{id}_B\,, \quad \text{for all } \pi \in \mathrm{Aut}\,\mathbb{M}\,.$
- $\pi[\mathbb{R}] = \mathbb{R} \quad \Leftrightarrow \quad \pi \upharpoonright B = \mathrm{id}_B\,, \quad \text{for all } \pi \in \mathrm{Aut}\,\mathbb{M}^{eq}.$

□

Lemma 3.4. *Let $\mathbb{R}$ be a relation and A, B sets.*

(a) *If* $\operatorname{dcl}(A) = \operatorname{dcl}(B)$, *then A is a Galois base of $\mathbb{R}$ if, and only if, B is a Galois base of $\mathbb{R}$.*

(b) *If A and B are both Galois bases of $\mathbb{R}$, then* $\operatorname{dcl}(A) = \operatorname{dcl}(B)$.

Proof. (a) Suppose that A is a Galois base of $\mathbb{R}$. By Corollary 1.8, it follows that

$$\operatorname{Aut} \mathbb{M}_B = \operatorname{Aut} \mathbb{M}_A = \operatorname{Aut}\langle \mathbb{M}, \mathbb{R} \rangle .$$

Hence, B is a Galois base of $\mathbb{R}$.

(b) Since both A and B are Galois bases, we have

$$\operatorname{Aut} \mathbb{M}_B = \operatorname{Aut}\langle \mathbb{M}, \mathbb{R} \rangle = \operatorname{Aut} \mathbb{M}_A .$$

Therefore it follows by Corollary 1.8 that $\operatorname{dcl}(A) = \operatorname{dcl}(B)$. □

With these preparations we can prove that, for parameter-definable relations, Galois bases and canonical parameters are the same.

Proposition 3.5. *Let $\mathbb{R}$ be a parameter-definable relation and $\bar{b}$ a tuple. The following statements are equivalent:*

(1) *$\bar{b}$ is a Galois base of $\mathbb{R}$.*

(2) *$\bar{b}$ is a canonical parameter of $\mathbb{R}$.*

(3) $\operatorname{dcl}^{\mathrm{eq}}(\bar{b})$ *is the least* $\operatorname{dcl}^{\mathrm{eq}}$*-closed set over which $\mathbb{R}$ is definable.*

Proof. (2) ⇒ (1) Suppose that $\psi(\bar{x}; \bar{b})$ is a canonical definition of $\mathbb{R}$. To show that $\bar{b}$ is a Galois base of $\mathbb{R}$, consider an automorphism π of $\mathbb{M}$. Then

$$\pi(\bar{b}) = \bar{b} \quad \text{implies} \quad \pi[\mathbb{R}] = \psi(\bar{x}; \pi(\bar{b}))^{\mathbb{M}} = \psi(\bar{x}; \bar{b})^{\mathbb{M}} = \mathbb{R} .$$

Conversely,

$$\pi[\mathbb{R}] = \mathbb{R} \quad \text{implies} \quad \psi(\bar{x}; \pi(\bar{b}))^{\mathbb{M}} = \psi(\bar{x}; \bar{b})^{\mathbb{M}}.$$

By uniqueness of $\bar{b}$, it follows that $\pi(\bar{b}) = \bar{b}$.

(1) ⇒ (2) Suppose that $\bar{b}$ is a Galois base of $\mathbb{R}$. By Lemma 3.2, there exists a formula $\varphi(\bar{x}; \bar{z})$ such that

$$\mathbb{R} = \varphi(\bar{x}; \bar{b})^{\mathbb{M}}.$$

First, let us show that there is no tuple $\bar{b}' \neq \bar{b}$ with

$$\bar{b}' \equiv_{\varnothing} \bar{b} \quad \text{and} \quad \varphi(\bar{x}; \bar{b}')^{\mathbb{M}} = \varphi(\bar{x}; \bar{b})^{\mathbb{M}}.$$

For a contradiction, suppose otherwise. Since $\bar{b}$ and $\bar{b}'$ have the same type, there exists an automorphism π with $\pi(\bar{b}) = \bar{b}'$. It follows that

$$\pi[\mathbb{R}] = \pi[\varphi(\bar{x}; \bar{b})^{\mathbb{M}}] = \varphi(\bar{x}; \pi(\bar{b}))^{\mathbb{M}} = \varphi(\bar{x}; \bar{b}')^{\mathbb{M}} = \mathbb{R}.$$

Since $\bar{b}$ is a Galois base of $\mathbb{R}$, this implies that $\pi(\bar{b}) = \bar{b}$. Hence, $\bar{b}' = \bar{b}$. Contradiction.

Set $\Phi(\bar{x}) := \operatorname{tp}(\bar{b})$. We have shown that

$$\Phi(\bar{y}) \cup \Phi(\bar{y}') \cup \{\forall \bar{x}[\varphi(\bar{x}; \bar{y}) \leftrightarrow \varphi(\bar{x}; \bar{y}')]\} \vDash \bar{y} = \bar{y}'.$$

By compactness, there exists a finite subset $\Phi_0 \subseteq \Phi$ such that

$$\Phi_0(\bar{y}) \cup \Phi_0(\bar{y}') \cup \{\forall \bar{x}[\varphi(\bar{x}; \bar{y}) \leftrightarrow \varphi(\bar{x}; \bar{y}')]\} \vDash \bar{y} = \bar{y}'.$$

Consequently, we obtain a canonical definition of $\mathbb{R}$ by setting

$$\psi(\bar{x}; \bar{b}) := \varphi(\bar{x}; \bar{b}) \wedge \bigwedge \Phi_0(\bar{b}).$$

(2) ⇒ (3) Let $\bar{b}$ be a Galois base of $\mathbb{R}$. We have seen in Lemma 3.2 that $\mathbb{R}$ is definable over $\bar{b}$. Suppose that $\mathbb{R}$ is definable over a $\operatorname{dcl}^{\text{eq}}$-closed set $A \subseteq \mathbb{M}^{\text{eq}}$. For $\pi \in \operatorname{Aut} \mathbb{M}^{\text{eq}}$, it follows that

$$\pi \upharpoonright A = \operatorname{id}_A \quad \text{implies} \quad \pi[\mathbb{R}] = \mathbb{R} \quad \text{implies} \quad \pi(\bar{b}) = \bar{b}.$$

Consequently, $\operatorname{Aut}\mathbb{M}^{\mathrm{eq}}_A \subseteq \operatorname{Aut}\mathbb{M}^{\mathrm{eq}}_{\bar b}$ and it follows by Corollary 1.8 that $\bar b \subseteq \mathrm{dcl}^{\mathrm{eq}}(A)$.

(3) ⇒ (1) We have seen in Corollary 2.9 that $\mathbb{R}$ has a canonical parameter $e \in \mathbb{M}^{\mathrm{eq}}$. By (3), this implies that $\mathrm{dcl}^{\mathrm{eq}}(\bar b) \subseteq \mathrm{dcl}^{\mathrm{eq}}(e)$. Conversely, since $\mathbb{R}$ is definable over $\bar b$, it follows by the already proved implication (2) ⇒ (3) that $\mathrm{dcl}^{\mathrm{eq}}(e) \subseteq \mathrm{dcl}^{\mathrm{eq}}(\bar b)$. Consequently, $\mathrm{dcl}^{\mathrm{eq}}(e) = \mathrm{dcl}^{\mathrm{eq}}(\bar b)$. Note that, by the already established implication (1) ⇒ (2), e is a Galois base of $\mathbb{R}$. Therefore, we can use Lemma 3.4 (a) to show that $\bar b$ is also a Galois base of $\mathbb{R}$. □

Relations that are not definable still might have a Galois base. Of particular interest are relations that are definable by types.

Definition 3.6. A *Galois base* of a type $\mathfrak{p} \in S^{\bar s}(\mathbb{M})$ is a Galois base of the relation $\mathfrak{p}^{\mathbb{M}}$ defined by it.

For types, Galois bases do not necessarily exists. But if they do, they are unique up to definable equivalence.

Definition 3.7. For a type $\mathfrak{p}$ with Galois base B, we set

$$\mathrm{Gb}(\mathfrak{p}) := \mathrm{dcl}^{\mathrm{eq}}(B)\,.$$

Remark. By the Lemma 3.4, it follows that $\mathrm{Gb}(\mathfrak{p})$ is the maximal Galois base of $\mathfrak{p}$ and that it does not depend on the choice of B.

Lemma 3.8. *Let T be a complete first-order theory and $\mathfrak{p} \in S^{\bar s}(\mathbb{M})$ a type. If $\mathfrak{p}$ is definable over $U \subseteq \mathbb{M}$, it has a Galois base $B \subseteq \mathrm{dcl}^{\mathrm{eq}}(U)$ of size $|B| \leq |T|$.*

Proof. Let $\varphi(\bar x; \bar y)$ be a formula without parameters and let $\delta_\varphi(\bar y)$ be a φ-definition of $\mathfrak{p}$ over U. By Corollary 2.9, the relation $\mathbb{R}_\varphi := (\delta_\varphi)^{\mathbb{M}}$ has a Galois base $b_\varphi \in \mathrm{dcl}^{\mathrm{eq}}(U)$. Set $B := \{\, b_\varphi \mid \varphi \text{ a formula} \,\}$. Then $|B| \leq |T|$ and $B \subseteq \mathrm{dcl}^{\mathrm{eq}}(U)$. To show that B is a Galois base of $\mathfrak{p}$, consider an automorphism $\pi \in \operatorname{Aut} \mathbb{M}^{\mathrm{eq}}$. Then

$$\begin{aligned} \pi(\mathfrak{p}) = \mathfrak{p} \quad & \text{iff} \quad \pi[\mathbb{R}_\varphi] = \mathbb{R}_\varphi\,, \quad \text{for all } \varphi \\ & \text{iff} \quad \pi(b_\varphi) = b_\varphi\,, \quad \text{for all } \varphi \\ & \text{iff} \quad \pi \restriction B = \mathrm{id}_B\,, \end{aligned}$$

as desired. □

Corollary 3.9. *In a stable first-order theory T, every complete type over a set U has a Galois base in* $\mathrm{dcl}^{\mathrm{eq}}(U)$.

Proof. Let $\mathfrak{p}$ be a complete type over U. According to Theorem C3.5.17, $\mathfrak{p}$ is definable over U. Hence, the claim follows by Lemma 3.8. □

Lemma 3.10. *Let $\mathfrak{p} \in S^{\bar{s}}(\mathbb{M})$ be a definable type and $U \subseteq \mathbb{M}$ a set of parameters. Then $\mathfrak{p}$ is definable over U if, and only if,* $\mathrm{Gb}(\mathfrak{p}) \subseteq \mathrm{dcl}^{\mathrm{eq}}(U)$.

Proof. ($\Rightarrow$) follows by Lemma 3.8.

($\Leftarrow$) According to Lemma 3.8, $\mathfrak{p}$ has a Galois base B. Since $\mathfrak{p}$ is definable we can find, for every formula $\varphi(\bar{x};\bar{y})$, a definable relation $\mathbb{R}_\varphi$ such that

$$\varphi(\bar{x};\bar{c}) \in \mathfrak{p} \quad \text{iff} \quad \bar{c} \in \mathbb{R}_\varphi\,.$$

Since $B \subseteq \mathrm{Gb}(\mathfrak{p}) \subseteq \mathrm{dcl}^{\mathrm{eq}}(U)$, it is sufficient to show that $\mathbb{R}_\varphi$ is definable over B. For each automorphism $\pi \in \operatorname{Aut} \mathbb{M}^{\mathrm{eq}}_B$, we have $\pi[\mathfrak{p}] = \mathfrak{p}$. Consequently, $\pi[\mathbb{R}_\varphi] = \mathbb{R}_\varphi$. Therefore, Lemma 3.2 implies that $\mathbb{R}_\varphi$ is definable over B. □

We conclude this section with a characterisation of the algebraic closure in $\mathbb{M}^{\mathrm{eq}}$. We start with an analogue of Lemma 1.10 for the algebraic closure.

Lemma 3.11. *A parameter-definable relation $\mathbb{R}$ has finitely many conjugates over a set $U \subseteq \mathbb{M}$ if, and only if, $\mathbb{R}$ is definable over* $\mathrm{acl}^{\mathrm{eq}}(U)$.

Proof. ($\Leftarrow$) Suppose that $\mathbb{R}$ is definable over $\bar{c} \subseteq \mathrm{acl}^{\mathrm{eq}}(U)$. Then

$$|\{ \pi[\mathbb{R}] \mid \pi \in \mathrm{Aut}\, \mathbb{M}^{\mathrm{eq}}_U \}| \leq |\{ \pi(\bar{c}) \mid \pi \in \mathrm{Aut}\, \mathbb{M}^{\mathrm{eq}}_U \}| < \aleph_0 \,.$$

Hence, $\mathbb{R}$ has only finitely many conjugates over U.

($\Rightarrow$) Suppose that $\mathbb{R}$ has only finitely many conjugates over U and let $\bar{b}$ be a Galois base of $\mathbb{R}$. Then

$$|\{ \pi(\bar{b}) \mid \pi \in \mathrm{Aut}\, \mathbb{M}^{\mathrm{eq}}_U \}| \leq |\{ \pi[\mathbb{R}] \mid \pi \in \mathrm{Aut}\, \mathbb{M}^{\mathrm{eq}}_U \}| < \aleph_0 \,.$$

By Theorem 1.6, it follows that $\bar{b} \subseteq \mathrm{acl}^{\mathrm{eq}}(U)$. Furthermore, we have seen in Lemma 3.2 that $\mathbb{R}$ is definable over $\bar{b}$. □

The algebraic closure of a set U in $\mathbb{M}^{\mathrm{eq}}$ can be characterised as follows.

Definition 3.12. Let $U \subseteq \mathbb{M}$ be a set of parameters and $\bar{s}$ a finite tuple of sorts. We denote by $\mathrm{FE}^{\bar{s}}(U)$ the set of all formulae $\chi(\bar{x}, \bar{y})$ over U where $\bar{x}$ and $\bar{y}$ have sort $\bar{s}$ such that $\chi^{\mathbb{M}}$ is an equivalence relation on $\mathbb{M}^{\bar{s}}$ with finitely many classes.

Lemma 3.13. *Let $\bar{a}, \bar{b} \in \mathbb{M}^{\bar{s}}$ be finite tuples and $U \subseteq \mathbb{M}$ a set of parameters. Then*

$$\bar{a} \equiv_{\mathrm{acl}^{\mathrm{eq}}(U)} \bar{b} \quad \text{iff} \quad \mathbb{M} \vDash \chi(\bar{a}, \bar{b}) \quad \textit{for all } \chi \in \mathrm{FE}^{\bar{s}}(U) \,.$$

Proof. ($\Rightarrow$) Let $\chi \in \mathrm{FE}^{\bar{s}}(U)$ and let $\mathbb{B} := [\bar{b}]_{\chi^{\mathbb{M}}} \subseteq \mathbb{M}^{\bar{s}}$ be the $\chi^{\mathbb{M}}$-class of $\bar{b}$. The conjugates of $\mathbb{B}$ over U are $\chi^{\mathbb{M}}$-classes. Since there are only finitely many such classes, it follows by Lemma 3.11 (b) that $\mathbb{B}$ is definable over $\mathrm{acl}^{\mathrm{eq}}(U)$. Therefore, we can use Proposition 3.5 and Corollary 2.9 to find a canonical definition $\psi(\bar{x}; e)$ of $\mathbb{B}$ where $e \in \mathrm{dcl}^{\mathrm{eq}}(\mathrm{acl}^{\mathrm{eq}}(U)) = \mathrm{acl}^{\mathrm{eq}}(U)$. Since

$$\bar{a} \equiv_{\mathrm{acl}^{\mathrm{eq}}(U)} \bar{b} \,,$$

it follows that

$$\mathbb{M} \vDash \psi(\bar{b}; e) \quad \text{implies} \quad \mathbb{M} \vDash \psi(\bar{a}; e) \,.$$

Hence, $\bar{a} \in \mathbb{B}$ implies $\mathbb{M} \vDash \chi(\bar{a}, \bar{b})$.

($\Leftarrow$) Suppose that $\mathbb{M} \vDash \varphi(\bar{a}; \bar{c})$, for $\bar{c} \subseteq \mathrm{acl}^{\mathrm{eq}}(U)$. We have to show that $\mathbb{M} \vDash \varphi(\bar{b}; \bar{c})$. There exists a formula $\psi(\bar{x})$ over U such that $\psi^{\mathbb{M}}$ is a finite set containing $\bar{c}$. The formula

$$\chi(\bar{x}, \bar{y}) := (\forall \bar{z}.\psi(\bar{z}))[\varphi(\bar{x}; \bar{z}) \leftrightarrow \varphi(\bar{y}; \bar{z})]$$

defines an equivalence relation with finitely many classes. Therefore, $\chi \in \mathrm{FE}^{\bar{s}}(U)$ and $\mathbb{M} \vDash \chi(\bar{a}, \bar{b})$. Since $\bar{c} \in \psi^{\mathbb{M}}$, it follows that

$$\mathbb{M} \vDash \varphi(\bar{a}; \bar{c}) \quad \text{implies} \quad \mathbb{M} \vDash \varphi(\bar{b}; \bar{c}) . \qquad \square$$

4. *Elimination of imaginaries*

In the abstract we can capture the property of $\mathbb{M}^{\mathrm{eq}}$ exhibited in Proposition 2.10 by the following definition.

Definition 4.1. A structure $\mathfrak{M}$ has *uniform elimination of imaginaries* if, for every equivalence formula $\chi(\bar{x}, \bar{y})$ of type $\bar{s}$, there exist sorts $\bar{t}$ and a definable function $f : M^{\bar{s}} \to M^{\bar{t}}$ such that $\ker f = \chi^{\mathfrak{M}}$.

We say that a theory T has *uniform elimination of imaginaries* if every model of T does.

We have shown in Proposition 2.10 that structures of the form $\mathfrak{M}^{\mathrm{eq}}$ have uniform elimination of imaginaries.

Proposition 4.2. *Every structure of the form* $\mathfrak{M}^{\mathrm{eq}}$ *has uniform elimination of imaginaries.*

Exercise 4.1. Show that the structure $\langle \mathbb{N}, +, \cdot \rangle$ has uniform elimination of imaginaries.

Frequently, the following weaker condition is equivalent to having uniform elimination of imaginaries.

Definition 4.3. A structure $\mathfrak{M}$ has *elimination of imaginaries* if, for each equivalence formula $\chi(\bar{x}, \bar{y})$ of type $\bar{s}$ and all tuples $\bar{a} \in M^{\bar{s}}$, the equivalence class $[\bar{a}]_\chi$ has a canonical parameter.

We say that a theory T has *elimination of imaginaries* if every model of T does.

For structures where $\mathrm{dcl}(\varnothing)$ is non-trivial, elimination of imaginaries already implies uniform elimination of imaginaries.

Lemma 4.4. *Let $\mathfrak{M}$ be a structure. The following statements are equivalent:*

(1) *$\mathfrak{M}$ has uniform elimination of imaginaries.*

(2) *$\mathfrak{M}$ has elimination of imaginaries and at least one of the following conditions holds:*

- *There is some sort u with $|\mathrm{dcl}(\varnothing) \cap M^u| > 1$.*
- *$|M^s| \leq 1$, for all sorts s.*

Proof. (1) ⇒ (2) To show that $\mathfrak{M}$ has elimination of imaginaries, consider an equivalence formula $\chi(\bar{x}, \bar{y})$ and a tuple $\bar{a}$ in M. By (1), there exists a definable function f with $\ker f = \chi^{\mathfrak{M}}$. Then $[\bar{a}]_\chi$ has the canonical definition

$$\psi(\bar{x}; \bar{b}) := (f(\bar{x}) = \bar{b}) \quad \text{where} \quad \bar{b} := f(\bar{a}) .$$

To conclude the proof, suppose that there is some sort s with $|M^s| > 1$. We have to find a sort u with $|\mathrm{dcl}(\varnothing) \cap M^u| > 1$. Consider the equivalence formula

$$\chi(xx', yy') := (x = x') \leftrightarrow (y = y')$$

of type ss. By (1), there exists a definable function f with $\ker f = \chi^{\mathfrak{M}}$. Fix distinct elements $c, d \in M^s$. It follows that the tuples $\bar{a} := f(c, c)$ and $\bar{b} := f(c, d)$ are definable and distinct. Fixing an index i with $a_i \neq b_i$, we obtain distinct elements a_i and b_i in $\mathrm{dcl}(\varnothing)$ of the same sort.

(2) ⇒ (1) If $|M^s| \leq 1$, for all sorts s, every equivalence formula χ defines the equality relation. Hence, the identity function has kernel $\chi^{\mathfrak{M}}$ and we are done.

It therefore remains to consider the case where $|\mathrm{dcl}(\varnothing) \cap M^u| > 1$, for some sort u. Let $\chi(\bar{x}, \bar{y})$ be an equivalence formula of type $\bar{s}$. For every tuple $\bar{a} \in M^{\bar{s}}$, fix a canonical definition $\delta_{\bar{a}}(\bar{x}; \bar{b}_{\bar{a}})$ of $[\bar{a}]_\chi$. Let $\bar{t}_{\bar{a}}$ be the sorts of $\bar{b}_{\bar{a}}$. We obtain a formula

$$\psi_{\bar{a}}(\bar{x}; \bar{y}) := \delta_{\bar{a}}(\bar{x}, \bar{y}) \wedge \forall \bar{z}[\delta_{\bar{a}}(\bar{z}; \bar{y}) \leftrightarrow \chi(\bar{x}, \bar{z})]$$

that defines a partial function $f_{\bar{a}} : U_{\bar{a}} \to M^{\bar{t}_{\bar{a}}}$ with kernel $\chi^{\mathfrak{M}}|_{U_{\bar{a}}}$. Note that the domain $U_{\bar{a}}$ of $f_{\bar{a}}$ is a union of χ-classes and that it is definable by the formula

$$\vartheta_{\bar{a}}(\bar{x}) := \exists \bar{y} \psi_{\bar{a}}(\bar{x}, \bar{y}) .$$

Hence,

$$M^{\bar{s}} = \bigcup_{\bar{a} \in M^{\bar{s}}} U_{\bar{a}} \quad \text{implies} \quad \mathrm{Th}(\mathfrak{M}) \vDash \bigvee_{\bar{a} \in M^{\bar{s}}} \vartheta_{\bar{a}} .$$

By compactness, there are finitely many tuples $\bar{a}_0, \dots, \bar{a}_n \in M^{\bar{s}}$ such that $M^{\bar{s}} = U_{\bar{a}_0} \cup \dots \cup U_{\bar{a}_n}$. Fix distinct elements $c, d \in \mathrm{dcl}(\varnothing) \cap M^u$. The formula

$$\begin{aligned} &\varphi(\bar{x}; \bar{y}_0, \dots, \bar{y}_n, \bar{z}) := \\ &\quad \bigvee_{i \le n} \Big[\psi_{\bar{a}_i}(\bar{x}; \bar{y}_i) \wedge \bar{x} \in U_{\bar{a}_i} \setminus (U_{\bar{a}_0} \cup \dots \cup U_{\bar{a}_{i-1}}) \\ &\qquad\qquad \wedge \bigwedge_{j \neq i} \bar{y}_j = \langle c, \dots, c \rangle \\ &\qquad\qquad \wedge \bar{z} = \langle \underbrace{c, \dots, c}_{i \text{ times}}, d \dots, d \rangle \Big] \end{aligned}$$

defines a function $f : M^{\bar{s}} \to M^{\bar{t}_{\bar{a}_0} \dots \bar{t}_{\bar{a}_n} u \dots u}$ with $\ker f = \chi^{\mathfrak{M}}$. □

As an example, we consider o-minimal structures and, in particular, real closed fields. We say that a theory T has *definable Skolem functions* if, for every formula $\varphi(\bar{x}, y)$, there exists a definable function f such that

$$T \vDash \forall \bar{x}[\exists y \varphi(\bar{x}, y) \to \varphi(\bar{x}, f(\bar{x}))] .$$

Proposition 4.5. *Every o-minimal structure $\mathfrak{M}$ with definable Skolem functions has elimination of imaginaries.*

Proof. We start by proving that every parameter-definable set $P \subseteq M$ has a canonical definition. Suppose that $P \subseteq M$ is parameter-definable. By o-minimality, P is of the form

$$P = (a_0, b_0) \cup \dots \cup (a_{m-1}, b_{m-1}) \cup \{c_0, \dots, c_{n-1}\},$$

for elements $a_i, b_i, c_i \in M$ satisfying

$$a_0 < b_0 < a_1 < b_1 < \dots < a_{m-1} < b_{m-1} \quad \text{and} \quad c_0 < \dots < c_{n-1}.$$

Fix such a decomposition of P where m and n are minimal. Then

$$\begin{aligned}\psi(x; \bar{a}, \bar{b}, \bar{c}) := \Big[\bigvee_{i<m} (a_i < x \wedge x < b_i) \vee \bigvee_{i<n} x = c_i \Big] \\ \wedge \Big[\bigwedge_{i<m} a_i < b_i \wedge \bigwedge_{i<m-1} b_i < a_{i+1} \wedge \bigwedge_{i<n-1} c_i < c_{i+1} \Big]\end{aligned}$$

is a canonical definition of P.

To show that $\mathfrak{M}$ has elimination of imaginaries, let $\chi(\bar{x}, \bar{y})$ be an equivalence formula of type $\bar{s}$ and let $\bar{a} \in M^{\bar{s}}$. To find a canonical definition of $[\bar{a}]_\chi$, we define, by induction on $i < n := |\bar{s}|$, a formula $\psi_i(y_i; \bar{z}_i)$, parameters $\bar{b}_i$, and a definable function s_i such that

- $\psi_i(y_i; \bar{b}_i)$ is a canonical definition of the relation defined by

$$\begin{aligned}&\vartheta_i(y_i; \bar{a}, \bar{b}_0, \dots, \bar{b}_{i-1}) := \\ &\quad \exists y_{i+1} \cdots \exists y_{n-1} \chi(\bar{a}, s_0(\bar{b}_0), \dots, s_{i-1}(\bar{b}_{i-1}), \\ &\qquad y_i, y_{i+1}, \dots, y_{n-1}),\end{aligned}$$

- $\mathfrak{M} \vDash \psi_i(s_i(\bar{b}_i); \bar{b}_i)$.

Suppose that we have already defined the formulae $\psi_0(y_0; \bar{b}_0), \dots, \psi_{i-1}(y_{i-1}; \bar{b}_{i-1})$ and the functions $s_0, \dots, s_{i-1}$. Since ϑ_i defines a set, we can use the statement we have proved above to find a canonical definition $\psi_i(y_i;\bar b_i)$ of $\vartheta_i^{\mathfrak{M}}(y_i;\bar a,\bar b_0,\dots,\bar b_{i-1})$. Let s_i be a definable Skolem function for the formula $\psi_i(y_i;\bar z_i)$. This concludes the inductive step.

We claim that the formula

$$\begin{aligned}\psi(\bar x;\bar b_0,\dots,\bar b_{n-1}) :=&\\ &\chi(\bar x,\ s_0(\bar b_0),\dots,s_{n-1}(\bar b_{n-1}))\\ &\wedge \bigwedge_{i<n} \forall y_i\big[\psi_i(y_i;\bar b_i) \leftrightarrow \vartheta_i(y_i;\bar x,\bar b_0,\dots,\bar b_{i-1})\big]\end{aligned}$$

is a canonical definition of $[\bar a]_\chi$. By construction, we have

$$\psi(\bar x;\bar b_0,\dots,\bar b_{n-1})^{\mathfrak{M}} = [\bar a]_\chi\,.$$

Suppose that $\bar b'_0,\dots,\bar b'_{n-1}$ are tuples such that

$$\psi(\bar x;\bar b'_0,\dots,\bar b'_{n-1})^{\mathfrak{M}} = [\bar a]_\chi\,.$$

Then

$$\psi_i(y_i;\bar b'_i)^{\mathfrak{M}} = \vartheta_i(\bar y_i;\bar a,\bar b'_0,\dots,b'_{i-1})^{\mathfrak{M}}\,.$$

By choice of ψ_i we can use induction on i to show that $\bar b'_i = \bar b_i$. □

Corollary 4.6. *The theory* RCF *of real closed fields has uniform elimination of imaginaries.*

Proof. After we have shown that RCF has definable Skolem functions, we can use Proposition 4.5 to show that RCF has elimination of imaginaries. Since $0,1 \in \mathrm{dcl}(\emptyset)$, it therefore follows by Lemma 4.4 that it even has uniform elimination of imaginaries.

Hence, it remains to show that RCF has definable Skolem functions. Let $\varphi(\bar x,y)$ be a formula. By o-minimality, for every choice of values $\bar c$ for the variables $\bar x$, the relation $\varphi(\bar a,y)^{\mathbb{M}}$ is of the form

$$\varphi(\bar c,y)^{\mathbb{M}} = (a_0,b_0)\cup\dots\cup(a_{m-1},b_{m-1})\cup\{d_0,\dots,d_{n-1}\}\,,$$

for elements $a_i, b_i, c_i \in M$ satisfying

$$a_0 < b_0 < a_1 < b_1 < \cdots < a_{m-1} < b_{m-1} \quad \text{and} \quad d_0 < \cdots < d_{n-1} .$$

Furthermore, it follows by Theorem D3.3.11 that there exists a bound $k < \omega$ such that, for every tuple $\bar{c}$, we can choose a decomposition as above where the numbers m and n are less than k.

Let $\psi(\bar{x}; y)$ be a formula stating that, for the given value of $\bar{x}$, there are numbers $m, n < k$ and tuples $\bar{a}, \bar{b}, \bar{d}$ such that

- $\varphi(\bar{x}, y')^{\mathbb{M}} = (a_0, b_0) \cup \cdots \cup (a_m, b_m) \cup \{d_0, \dots, d_n\}$,
- m and n are the minimal numbers such that $\varphi(\bar{x}, y')^{\mathbb{M}}$ can be written in this form,
- $a_0 < b_0 < a_1 < b_1 < \cdots < a_{m-1} < b_{m-1}$ and $d_0 < \cdots < d_{n-1}$,
- $y = \begin{cases} d_0 & \text{if } n > 0 , \\ (a_0 + b_0)/2 & \text{if } n = 0 ,\ m > 0 ,\ \text{and } -\infty < a_0 < b_0 < \infty , \\ b_0 - 1 & \text{if } n = 0 ,\ m > 0 ,\ \text{and } -\infty = a_0 < b_0 < \infty , \\ a_0 + 1 & \text{if } n = 0 ,\ m > 0 ,\ \text{and } -\infty < a_0 < b_0 = \infty , \\ 0 & \text{otherwise} . \end{cases}$

Then $\psi(\bar{x}, y)$ defines a Skolem function for $\varphi(\bar{x}, y)$. □

We can use Galois bases to characterise theories with elimination of imaginaries.

Proposition 4.7. *Let T be a complete first-order theory. The following statements are equivalent:*

(1) *T has elimination of imaginaries.*

(2) *Every parameter-definable relation has a canonical parameter.*

(3) *Every parameter-definable relation has a finite Galois base.*

(4) *For every parameter-definable relation $\mathbb{R}$, there exists a least* $\mathrm{dcl}^{\mathrm{eq}}$-*closed set $B \subseteq \mathbb{M}$ over which $\mathbb{R}$ is definable.*

(5) *For every imaginary element* $e \in \mathbb{M}^{\mathrm{eq}}$*, there is a finite set* $B \subseteq \mathbb{M}$ *with* $\mathrm{dcl}^{\mathrm{eq}}(e) = \mathrm{dcl}^{\mathrm{eq}}(B)$.

Proof. (3) ⇒ (4) ⇔ (2) follows by Proposition 3.5.

(2) ⇒ (1) Let $\chi(\bar{x}, \bar{y})$ be an equivalence formula. If every parameter-definable relation has a canonical parameter then, in particular, this is true for every relation of the form $[\bar{a}]_\chi$.

(1) ⇒ (5) Let $e \in \mathbb{M}^{\mathrm{eq}}_\chi$ be an imaginary element and $\mathbb{E} := p_\chi^{-1}(e)$ the corresponding equivalence class. Since T has elimination of imaginaries, there exists a canonical definition $\psi(\bar{x}; \bar{b})$ of $\mathbb{E}$. Obviously, we can choose the tuple $\bar{b}$ to be finite. According to Proposition 3.5, $\bar{b}$ is a Galois base of $\mathbb{E}$. Note that, in the structure $\mathbb{M}^{\mathrm{eq}}$, $\{e\}$ is a Galois base of $\mathbb{E}$. Consequently, it follows by Lemmas 3.3 and 3.4 that

$$\mathrm{dcl}^{\mathrm{eq}}(e) = \mathrm{dcl}^{\mathrm{eq}}(\bar{b}) \,.$$

(5) ⇒ (3) Let $\mathbb{R}$ be a parameter-definable relation. We fix a formula $\varphi(\bar{x}; \bar{c})$ with parameters $\bar{c}$ defining $\mathbb{R}$. Let $\chi(\bar{y}, \bar{y}')$ be the parameter equivalence for $\varphi(\bar{x}; \bar{y})$ and set $e := [\bar{c}]_\chi$. By assumption, there exists a finite set $B \subseteq \mathbb{M}$ such that $\mathrm{dcl}^{\mathrm{eq}}(e) = \mathrm{dcl}^{\mathrm{eq}}(B)$. We claim that B is a Galois base of $\mathbb{R}$. Note that, by Lemma 3.3, it is sufficient to prove that B is a Galois base of $\mathbb{R}$ in the structure $\mathbb{M}^{\mathrm{eq}}$. Furthermore, it follows by Lemma 2.8 and Proposition 3.5 that e is a Galois base of $\mathbb{R}$. Therefore, Lemma 3.4 (a) implies that B is also a Galois base of $\mathbb{R}$. □

5. *Weak elimination of imaginaries*

In this section we take a look at a weaker condition than elimination of imaginaries.

Definition 5.1. (a) A tuple $\bar{c}$ is a *weak canonical parameter* of a relation $\mathbb{R}$ if there exist a formula $\psi(\bar{x}; \bar{y})$ such that $\bar{c}$ is one of only finitely many tuples satisfying

$$\psi(\bar{x}; \bar{c})^{\mathbb{M}} = \mathbb{R} \,.$$

In this case, we call the formula $\psi(\bar{x};\bar{c})$ a *weak canonical definition* of $\mathbb{R}$.

(b) A complete first-order theory T has *weak elimination of imaginaries* if, for each equivalence formula $\chi(\bar{x},\bar{y})$ of type $\bar{s}$ and all tuples $\bar{a} \in \mathbb{M}^{\bar{s}}$, the equivalence class $[\bar{a}]_\chi$ has a weak canonical parameter.

We start with an analogue of Proposition 3.5.

Lemma 5.2. *Let $\mathbb{R}$ be a parameter-definable relation and U a set. The following statements are equivalent:*

(1) $\mathbb{R}$ *has a weak canonical parameter $\bar{c}$ with* $\mathrm{acl}(\bar{c}) = \mathrm{acl}(U)$.

(2) $\mathrm{acl}(U)$ *is the least algebraically closed set over which $\mathbb{R}$ is definable.*

Proof. (1) ⇒ (2) Let $\psi(\bar{x};\bar{c})$ be a weak canonical definition of $\mathbb{R}$. We claim that $\mathrm{acl}(\bar{c})$ is the least algebraically closed set over which $\mathbb{R}$ is definable. Obviously, $\mathbb{R}$ is definable over $\mathrm{acl}(\bar{c})$. To show that $\mathrm{acl}(\bar{c})$ is the least such set, let $\varphi(\bar{x};\bar{b})$ be an arbitrary formula defining $\mathbb{R}$. We have to prove that $\mathrm{acl}(\bar{c}) \subseteq \mathrm{acl}(\bar{b})$. The formula

$$\vartheta(\bar{y};\bar{b}) := \forall\bar{x}[\psi(\bar{x};\bar{y}) \leftrightarrow \varphi(\bar{x};\bar{b})]$$

defines the finite set $\{\, \bar{c}' \mid \psi(\bar{x};\bar{c}')^{\mathbb{M}} = \mathbb{R} \,\}$. This implies that $\bar{c} \subseteq \mathrm{acl}(\bar{b})$, as desired.

(2) ⇒ (1) Suppose that $\mathrm{acl}(U)$ is the least algebraically closed set over which $\mathbb{R}$ is definable. Fix a formula $\psi(\bar{x};\bar{c})$ with parameters $\bar{c} \subseteq \mathrm{acl}(U)$ defining $\mathbb{R}$. Note that, by assumption on U, it follows that $\mathrm{acl}(\bar{c}) = \mathrm{acl}(U)$.

We start by proving that there are only finitely many tuples $\bar{c}'$ such that

$$\bar{c}' \equiv_\varnothing \bar{c} \quad \text{and} \quad \psi(\bar{x};\bar{c}')^{\mathbb{M}} = \mathbb{R}\,.$$

For a contradiction, suppose otherwise. By compactness, we can then find a tuple $\bar{c}'$ such that

$$\bar{c}' \nsubseteq \mathrm{acl}(\bar{c})\,, \quad \bar{c}' \equiv_\varnothing \bar{c}\,, \quad \text{and} \quad \psi(\bar{x};\bar{c}')^{\mathbb{M}} = \mathbb{R}\,.$$

Since $\mathbb{R}$ is definable over $\bar{c}'$ it follows by assumption on U that

$$\bar{c} \subseteq \mathrm{acl}(U) \subseteq \mathrm{acl}(\bar{c}')\,.$$

As $\bar{c}' \equiv_\varnothing \bar{c}$, there exists an automorphism π with $\pi(\bar{c}') = \bar{c}$. Setting $\bar{c}'' := \pi(\bar{c})$ it follows that

$$\bar{c} \nsubseteq \mathrm{acl}(\bar{c}'') \quad \text{and} \quad \bar{c}'' \subseteq \mathrm{acl}(\bar{c})\,,$$

Since, for every tuple $\bar{a}$,

$$\begin{aligned} \mathbb{M} \vDash \varphi(\bar{a};\bar{c}'') \quad &\text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a};\pi(\bar{c})) \\ &\text{iff} \quad \mathbb{M} \vDash \varphi(\pi^{-1}(\bar{a});\bar{c}) \\ &\text{iff} \quad \mathbb{M} \vDash \varphi(\pi^{-1}(\bar{a});\bar{c}') \\ &\text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a};\pi(\bar{c}')) \quad \text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a};\bar{c}')\,, \end{aligned}$$

it furthermore follows that $\psi(\bar{x};\bar{c}'')^{\mathbb{M}} = \mathbb{R}$. But, by assumption on U, this implies that $\bar{c} \subseteq \mathrm{acl}(U) \subseteq \mathrm{acl}(\bar{c}'')$. A contradiction.

Set $\Phi(\bar{y}) := \mathrm{tp}(\bar{c})$. We have shown that there exists a number $n < \omega$ such that

$$\Phi(\bar{y}_0) \cup \dots \cup \Phi(\bar{y}_n) \cup \big\{ \forall \bar{x}[\psi(\bar{x};\bar{y}_i) \leftrightarrow \psi(\bar{x};\bar{y}_k)] \bigm| i,k \leq n \big\}$$

is inconsistent. By compactness, we can find a finite subset $\Phi_0 \subseteq \Phi$ such that

$$\Phi_0(\bar{y}_0) \cup \dots \cup \Phi_0(\bar{y}_n) \cup \big\{ \forall \bar{x}[\psi(\bar{x};\bar{y}_i) \leftrightarrow \psi(\bar{x};\bar{y}_k)] \bigm| i,k \leq n \big\}$$

is already inconsistent. Consequently, the formula

$$\psi(\bar{x};\bar{c}) \wedge \bigwedge \Phi_0(\bar{c})$$

is a weak canonical definition of $\mathbb{R}$ with $\mathrm{acl}(\bar{c}) = \mathrm{acl}(U)$. □

Corollary 5.3. *If $\bar{a}$ and $\bar{b}$ are weak canonical parameters of a relation $\mathbb{R}$, then* $\mathrm{acl}(\bar{a}) = \mathrm{acl}(\bar{b})$.

For relations that do have a Galois base, we can be more precise.

Lemma 5.4. *Let $\mathbb{R}$ be a parameter-definable relation with Galois base $\bar{b}$. A tuple $\bar{c}$ is a weak canonical parameter of $\mathbb{R}$ if, and only if,*

$$\bar{b} \subseteq \operatorname{dcl}(\bar{c}) \quad \textit{and} \quad \bar{c} \subseteq \operatorname{acl}(\bar{b})\,.$$

Proof. By Proposition 3.5, we can fix a canonical definition $\hat{\psi}(\bar{x};\bar{b})$ of $\mathbb{R}$.

($\Rightarrow$) Suppose that $\psi(\bar{x};\bar{c})$ is a weak canonical definition of $\mathbb{R}$. Then $\bar{b} \subseteq \operatorname{dcl}(\bar{c})$ since $\bar{b}$ is the unique tuple satisfying

$$\vartheta(\bar{z};\bar{c}) := \forall\bar{x}[\psi(\bar{x};\bar{c}) \leftrightarrow \hat{\psi}(\bar{x};\bar{z})]\,.$$

Furthermore, $\bar{c} \subseteq \operatorname{acl}(\bar{b})$ since the formula

$$\varphi(\bar{y};\bar{b}) := \forall\bar{x}[\psi(\bar{x};\bar{y}) \leftrightarrow \hat{\psi}(\bar{x};\bar{b})]$$

defines a finite set containing $\bar{c}$.

($\Leftarrow$) Let us first consider the special case where $\mathbb{R} = \varnothing$. Then $\varnothing$ is a Galois base of $\mathbb{R}$ and it follows by Lemma 3.4 that $\bar{b} \subseteq \operatorname{dcl}(\varnothing)$. Hence, $\bar{c} \subseteq \operatorname{acl}(\varnothing)$ and there exists a formula $\vartheta(\bar{y})$ that defines a finite relation containing the tuple $\bar{c}$. It follows that the formula

$$\psi(\bar{x};\bar{c}) := \neg\vartheta(\bar{c})$$

is a weak canonical definition of $\mathbb{R} = \varnothing$.

It remains to consider the case where $\mathbb{R} \neq \varnothing$. Fix formulae $\vartheta(\bar{z};\bar{y})$ and $\varphi(\bar{y};\bar{z})$ such that $\vartheta(\bar{z};\bar{c})^{\mathbb{M}} = \{\bar{b}\}$ and $\varphi(\bar{y};\bar{b})^{\mathbb{M}}$ is a finite set containing $\bar{c}$. We claim that the formula

$$\psi(\bar{x};\bar{c}) := \exists\bar{z}\big[\vartheta(\bar{z};\bar{c}) \wedge \hat{\psi}(\bar{x};\bar{z}) \wedge \varphi(\bar{c};\bar{z})\big]$$

is a weak canonical definition of $\mathbb{R}$. Clearly, $\psi(\bar{x};\bar{c})^{\mathbb{M}} = \mathbb{R}$. Furthermore, suppose that $\bar{c}'$ is a tuple such that $\psi(\bar{x};\bar{c}')^{\mathbb{M}} = \mathbb{R}$. Fix a tuple $\bar{a} \in \mathbb{R}$ and let $\bar{b}'$ be a tuple such that

$$\mathbb{M} \vDash \vartheta(\bar{b}';\bar{c}') \wedge \hat{\psi}(\bar{a};\bar{b}') \wedge \varphi(\bar{c}';\bar{b}')\,.$$

Then $\mathbb{R} = \psi(\bar{x}; \bar{c}')^{\mathbb{M}} = \hat{\psi}(\bar{x}; \bar{b}')^{\mathbb{M}}$ implies that $\bar{b}' = \bar{b}$. Hence, we have $\mathbb{M} \vDash \varphi(\bar{c}'; \bar{b})$. Since there are only finitely many such tuples $\bar{c}'$, it follows that $\psi(\bar{x}; \bar{c})^{\mathbb{M}}$ is a weak canonical definition of $\mathbb{R}$. □

We obtain a characterisation of theories with weak elimination of imaginaries along the same lines as Proposition 4.7.

Proposition 5.5. *Let T be a complete first-order theory. The following statements are equivalent:*

(1) *T has weak elimination of imaginaries.*

(2) *All parameter-definable relations have weak canonical parameters.*

(3) *For every parameter-definable relation $\mathbb{R}$, there is a least algebraically closed set over which $\mathbb{R}$ is definable.*

(4) *For every element $e \in \mathbb{M}^{\mathrm{eq}}$, there is a finite set $B \subseteq \mathbb{M}$ such that*

$$e \in \mathrm{dcl}^{\mathrm{eq}}(B) \quad \textit{and} \quad B \subseteq \mathrm{acl}^{\mathrm{eq}}(e)\,.$$

(5) *For every imaginary element $e \in \mathbb{M}^{\mathrm{eq}}$, there exists a finite tuple $\bar{s}$ of sorts and a finite relation $C \subseteq \mathbb{M}^{\bar{s}}$ such that*

$$\mathrm{dcl}^{\mathrm{eq}}(e) = \mathrm{dcl}^{\mathrm{eq}}(B)\,, \quad \textit{for every Galois base } B \textit{ of } C\,.$$

Proof. (4) ⇒ (1) Let $e \in \mathbb{M}^{\mathrm{eq}}_{\chi}$ be an imaginary element and $\mathbb{E} := p_{\chi}^{-1}(e)$ its equivalence class. By assumption, there exists a finite tuple $\bar{c} \subseteq \mathbb{M}$ such that $e \in \mathrm{dcl}^{\mathrm{eq}}(\bar{c})$ and $\bar{c} \subseteq \mathrm{acl}^{\mathrm{eq}}(e)$. Since e is a Galois base of $\mathbb{E}$ it follows by Lemma 5.4 that $\bar{c}$ is a weak canonical parameter of $\mathbb{E}$.

(1) ⇒ (3) Let $\mathbb{R}$ be a relation defined by the formula $\varphi(\bar{x}; \bar{b})$ and let χ be the parameter equivalence of φ. By assumption, there exists a finite relation C and a formula $\psi(\bar{z}; \bar{y})$ such that

$$\psi(\bar{z}; \bar{c})^{\mathbb{M}} = [\bar{b}]_{\chi} \quad \text{iff} \quad \bar{c} \in C\,.$$

We claim that $\mathrm{acl}(\bigcup C)$ is the desired algebraically closed set.

First, note that $\mathbb{R}$ is defined over $\bar{c} \subseteq \operatorname{acl}(\bigcup C)$ by the formula

$$\vartheta(\bar{x};\bar{c}) :\equiv \exists \bar{z}\big[\psi(\bar{z};\bar{c}) \wedge \varphi(\bar{x};\bar{z})\big].$$

Next, suppose that A is an algebraically closed set such that $\mathbb{R}$ is definable over A. For every $\pi \in \operatorname{Aut}\mathbb{M}$, it follows that

$$\begin{aligned}
\pi \restriction A = \mathrm{id}_A \quad &\Rightarrow \quad \pi[\mathbb{R}] = \mathbb{R} \\
&\Rightarrow \quad \varphi(\bar{x};\pi(\bar{b}'))^{\mathbb{M}} = \varphi(\bar{x};\bar{b}')^{\mathbb{M}}, \quad \text{for all } \bar{b}' \in [\bar{b}]_\chi \\
&\Rightarrow \quad \pi[\bar{b}]_\chi = [\bar{b}]_\chi \\
&\Rightarrow \quad \pi[\psi(\bar{x};\bar{c})^{\mathbb{M}}] = \psi(\bar{x};\bar{c})^{\mathbb{M}}, \quad \text{for all } \bar{c} \in C \\
&\Rightarrow \quad \pi[C] = C\,.
\end{aligned}$$

Since C is finite, it follows that every tuple $\bar{c} \in C$ has finitely many conjugates over A. Consequently, Theorem 1.6 implies that $\bigcup C \subseteq \operatorname{acl}(A)$.

(3) ⇒ (2) Let $\mathbb{R}$ be a parameter-definable relation. By assumption, there exists a least algebraically closed set U over which $\mathbb{R}$ is definable. Hence, we can apply Lemma 5.2 to obtain a weak canonical parameter $\bar{c} \subseteq U$ of $\mathbb{R}$.

(2) ⇒ (5) Let $e \in \mathbb{M}^{\mathrm{eq}}_\chi$ be an imaginary element and $\mathbb{E} := p_\chi^{-1}(e)$ its equivalence class. By assumption, $\mathbb{E}$ has a weak canonical definition $\psi(\bar{x};\bar{c})$. Obviously, we may assume that $\bar{c}$ is a finite tuple. Set

$$C := \{\, \bar{c}' \mid \psi(\bar{x};\bar{c}')^{\mathbb{M}} = \mathbb{E}\,\}\,.$$

For an automorphism $\pi \in \operatorname{Aut}\mathbb{M}^{\mathrm{eq}}$, it follows that

$$\begin{aligned}
\pi(e) = e \quad &\text{iff} \quad \pi[\mathbb{E}] = \mathbb{E} \\
&\text{iff} \quad \psi(\bar{x};\pi(\bar{c}))^{\mathbb{M}} = \psi(\bar{x};\bar{c})^{\mathbb{M}}, \quad \text{for all } \bar{c} \in C \\
&\text{iff} \quad \pi[C] = C\,.
\end{aligned}$$

Hence, e is a Galois base of C. Therefore, it follows by Lemma 3.4 (b) that

$$\operatorname{dcl}^{\mathrm{eq}}(e) = \operatorname{dcl}^{\mathrm{eq}}(B)\,, \quad \text{for every Galois base } B \text{ of } C\,.$$

(5) ⇒ (4) Suppose that $C = \{\bar{c}_0, \dots, \bar{c}_n\}$ is a finite relation such that $\mathrm{dcl}^{\mathrm{eq}}(e) = \mathrm{dcl}^{\mathrm{eq}}(B)$, for every Galois base B of C. Since $\mathbb{M}^{\mathrm{eq}}$ has elimination of imaginaries, there exists a Galois base $B \subseteq \mathbb{M}^{\mathrm{eq}}$ of C. Consequently, Lemma 3.4 (a) implies that e is also a Galois base of C.

Let π be an automorphisms of $\mathbb{M}^{\mathrm{eq}}$. Then

$$\begin{aligned} \pi(\bar{c}_0 \dots \bar{c}_n) = \bar{c}_0 \dots \bar{c}_n \quad &\text{implies} \quad \pi[C] = C \\ &\text{implies} \quad \pi(e) = e\,. \end{aligned}$$

By Corollary 1.8, it follows that $e \in \mathrm{dcl}^{\mathrm{eq}}(\bar{c}_0 \dots \bar{c}_n)$. Similarly,

$$\begin{aligned} \pi(e) = e \quad &\text{implies} \quad \pi[C] = C \\ &\text{implies} \quad \pi(\bar{c}_0 \dots \bar{c}_n) = \bar{c}_{\sigma(0)} \dots \bar{c}_{\sigma(n)}\,, \\ &\phantom{\text{implies}} \quad \text{for some permutation } \sigma. \end{aligned}$$

Therefore, there are only finitely many conjugates of $\bar{c}_0 \dots \bar{c}_n$ over e. According to Theorem 1.6 this implies that $\bar{c}_0 \dots \bar{c}_n \subseteq \mathrm{acl}^{\mathrm{eq}}(e)$. □

In later chapters we will present several conditions implying that a theory has weak elimination of imaginaries. Here, we give only one example.

Lemma 5.6. *A theory T satisfying the following two conditions has weak elimination of imaginaries:*

- *There is no strictly decreasing sequence $A_0 \supset A_1 \supset \dots$ of sets of the form $A_i = \mathrm{acl}(B_i)$ where each B_i is finite.*
- *If A and B are algebraic closures of finite sets, then $\mathfrak{Aut}\,\mathbb{M}_{A \cap B}$ is generated by $\mathrm{Aut}\,\mathbb{M}_A \cup \mathrm{Aut}\,\mathbb{M}_B$.*

Proof. By Proposition 5.5 it is sufficient to show that, for every parameter-definable relation $\mathbb{R}$, there is a least algebraically closed set over which $\mathbb{R}$ is definable.

Hence, let $\mathbb{R}$ be parameter-definable. First, let us show that, if $\mathbb{R}$ is definable over two algebraically closed sets A and B of the form $A = \mathrm{acl}(A_0)$

and $B = \mathrm{acl}(B_0)$, for finite A_0 and B_0, then it is also definable over their intersection $A \cap B$. If $\mathbb{R}$ is definable over both A and B, Lemma 1.10 implies that

$$\mathrm{Aut}\,\mathbb{M}_A \cup \mathrm{Aut}\,\mathbb{M}_B \subseteq \mathrm{Aut}\langle\mathbb{M}, \mathbb{R}\rangle .$$

Consequently, the second condition implies that

$$\mathrm{Aut}\,\mathbb{M}_{A\cap B} = \langle\!\langle \mathrm{Aut}\,\mathbb{M}_A \cup \mathrm{Aut}\,\mathbb{M}_B \rangle\!\rangle \subseteq \mathrm{Aut}\langle\mathbb{M}, \mathbb{R}\rangle .$$

Hence, it follows by Lemma 1.10 that $\mathbb{R}$ is definable over $A \cap B$.

By the first condition, it therefore follows that there is a least algebraically closed set over which $\mathbb{R}$ is definable. □

The following property is what is missing from weak elimination of imaginaries in order to obtain full elimination of imaginaries.

Definition 5.7. A complete first-order theory T has *elimination of finite imaginaries* if every finite relation has a finite Galois base in $\mathbb{M}$.

As an example, we consider the theory of algebraically closed fields. We will show later in Corollary **??** that this theory actually has uniform elimination of imaginaries.

Lemma 5.8. *The theory of algebraically closed fields of characteristic p has elimination of finite imaginaries.*

Proof. Let $R = \{\bar{c}^0, \dots, \bar{c}^{n-1}\}$ be a finite relation consisting of m-tuples $\bar{c}^i = \langle c^i_0, \dots, c^i_{m-1}\rangle$. We define the polynomial

$$p(x, y_0, \dots, y_{m-1}) := \prod_{i<n} (x - c^i_0 y_0 - \dots - c^i_{m-1} y_{m-1}) .$$

Let Z be the set of roots of p. Then

$$\pi[Z] = Z \quad \text{iff} \quad \pi[R] = R , \quad \text{for every automorphism } \pi .$$

Since p is the only polynomial with set of roots Z, it follows that an automorphism fixes p if, and only if, it permutes R. Consequently, the coefficients of p form a Galois base of R. □

Proposition 5.9. *A theory T has elimination of imaginaries if, and only if, it has both, elimination of finite imaginaries and weak elimination of imaginaries.*

Proof. ($\Rightarrow$) Since every canonical parameter is a weak canonical parameter, elimination of imaginaries implies weak elimination of implies. Moreover, it follows by Proposition 4.7 (3) that every theory with elimination of imaginaries has elimination of finite imaginaries.

($\Leftarrow$) Let $e \in \mathbb{M}^{eq}$. By Proposition 5.5, there exists a finite set $C \subseteq \mathbb{M}^{\bar{s}}$ such that

$$\mathrm{dcl}^{eq}(e) = \mathrm{dcl}^{eq}(B)\,, \quad \text{for every Galois base } B \text{ of } C\,.$$

As T has elimination of finite imaginaries, the set C has a finite Galois base $B_0 \subseteq \mathbb{M}$. Hence,

$$\mathrm{dcl}^{eq}(e) = \mathrm{dcl}^{eq}(B_0)\,.$$

By Proposition 4.7, it follows that T has elimination of imaginaries. □

E3. *Prime models*

1. *Isolated types*

The usual way to construct structures in model theory consists in writing down an appropriate theory and proving that it is consistent. In particular, we can reconstruct from the elementary diagram of a structure the structure itself, or we can use it to obtain an elementary extension. If we want to construct rich models realising many types then, as we have seen in Chapter E1, this approach works well.

In the present chapter, on the other hand, we are interested in models realising few types. We start by studying those types that are unavoidable in the sense that they are realised in every model.

Definition 1.1. Let T be a theory.

(a) A formula φ *isolates* a type $\mathfrak{p}$ (w.r.t. T) if $\varphi \vDash \mathfrak{p}$ modulo T. We call a type $\mathfrak{p}$ over U *isolated* if it is isolated by a formula $\varphi(\bar{x}, \bar{c})$ with parameters $\bar{c} \subseteq U$. In particular, a complete type $\mathfrak{p} \in S^{\bar{s}}(U)$ is isolated if and only if $\langle\varphi\rangle = \{\mathfrak{p}\}$, i.e., $\mathfrak{p}$ is an isolated point in the topology of $S^{\bar{s}}(U)$.

(b) A structure $\mathfrak{A}$ is *atomic* if every realised type $\mathfrak{p} \in S^{<\omega}(\varnothing)$ is isolated. More generally, if $B, U \subseteq A$ then we call B *atomic over* U if only isolated types $\mathfrak{p} \in S^{<\omega}(U)$ are realised in B.

Lemma 1.2. *If $\mathfrak{p}$ is isolated by $\varphi(\bar{x})$ then $\mathfrak{p}$ is realised in every model of $T \cup \{\exists\bar{x}\varphi\}$.*

Lemma 1.3. *If $\bar{a} \subseteq \mathrm{acl}(U)$ then $\mathrm{tp}(\bar{a}/U)$ is isolated.*

Proof. Let $\mathfrak{M}$ be a model containing U. Since $\bar{a}$ is algebraic over U we can choose a formula $\varphi(\bar{x}, \bar{c})$ with parameters $\bar{c} \subseteq U$ such that $\mathfrak{M} \vDash \varphi(\bar{a}, \bar{c})$

and the set $\varphi(\bar{x}, \bar{c})^{\mathfrak{M}}$ is finite and of minimal size. We claim that this formula isolates $\mathrm{tp}(\bar{a}/U)$.

For a contradiction suppose that there is some formula $\psi(\bar{x}, \bar{d}) \in \mathrm{tp}(\bar{a}/U)$ such that $\varphi \not\models \psi$. Then we can find a tuple $\bar{b} \subseteq M$ with

$$\mathfrak{M} \models \varphi(\bar{b}, \bar{c}) \wedge \neg\psi(\bar{b}, \bar{d}) .$$

It follows that

$$[\varphi(\bar{x}, \bar{c}) \wedge \psi(\bar{x}, \bar{d})]^{\mathfrak{M}} \subseteq \varphi(\bar{x}, \bar{c})^{\mathfrak{M}} \setminus \{\bar{b}\} \subset \varphi(\bar{x}, \bar{c})^{\mathfrak{M}},$$

in contradiction to our choice of φ. □

Lemma 1.4. *Every isolated type* $\mathfrak{p} \in S^{\bar{s}}(U)$ *is definable over a finite subset* $U_0 \subseteq U$.

Proof. Let $\varphi(\bar{x}, \bar{c})$ be a formula over U isolating $\mathfrak{p}$. We claim that $\mathfrak{p}$ is definable over $U_0 := \bar{c}$. Let $\psi(\bar{x}, \bar{y})$ be a formula and $\bar{b} \subseteq U$. Then we have

$$\begin{aligned} \psi(\bar{x}, \bar{b}) \in \mathfrak{p} \quad &\text{iff} \quad T(U) \cup \{\varphi(\bar{x}, \bar{c})\} \models \psi(\bar{x}, \bar{b}) \\ &\text{iff} \quad T(U) \models \forall \bar{x}[\varphi(\bar{x}, \bar{c}) \rightarrow \psi(\bar{x}, \bar{b})] . \end{aligned}$$

Consequently, $\delta_\psi(\bar{y}) := \forall \bar{x}[\varphi(\bar{x}, \bar{c}) \rightarrow \psi(\bar{x}, \bar{y})]$ is a ψ-definition of $\mathfrak{p}$ over U_0. □

Lemma 1.5. $\mathrm{tp}(\bar{a}\bar{b}/U)$ *is isolated if and only if the types* $\mathrm{tp}(\bar{a}/U)$ *and* $\mathrm{tp}(\bar{b}/U \cup \bar{a})$ *are isolated.*

Proof. ($\Leftarrow$) If $\varphi(\bar{x})$ isolates $\mathrm{tp}(\bar{a}/U)$ and $\psi(\bar{y}, \bar{a})$ isolates $\mathrm{tp}(\bar{b}/U \cup \bar{a})$ then the formula $\varphi(\bar{x}) \wedge \psi(\bar{y}, \bar{x})$ isolates $\mathrm{tp}(\bar{a}\bar{b}/U)$.

($\Rightarrow$) Let $\varphi(\bar{x}, \bar{y})$ be a formula isolating $\mathrm{tp}(\bar{a}\bar{b}/U)$. Then the formula $\varphi(\bar{a}, \bar{y})$ isolates $\mathrm{tp}(\bar{b}/U \cup \bar{a})$. Furthermore, we claim that $\exists \bar{y}' \varphi(\bar{x}, \bar{y})$ isolates $\mathrm{tp}(\bar{a}/U)$ where $\bar{y}' \subseteq \bar{y}$ is the finite tuple of those variables that actually appear in φ. Suppose that $\exists \bar{y}' \varphi \in \mathrm{tp}(\bar{c}/U)$. Then there is some tuple $\bar{d}$ with $\varphi \in \mathrm{tp}(\bar{c}\bar{d}/U)$. Consequently, $\mathrm{tp}(\bar{c}\bar{d}/U) = \mathrm{tp}(\bar{a}\bar{b}/U)$ and $\mathrm{tp}(\bar{c}/U) = \mathrm{tp}(\bar{a}/U)$. □

We conclude this section with a collection of basic facts about atomic models.

Lemma 1.6. *If A is atomic over U and $\bar{a} \in A^{<\omega}$ then A is atomic over $U \cup \bar{a}$.*

Proof. For every finite tuple $\bar{b} \in A^{<\omega}$ we know that $\mathrm{tp}(\bar{a}\bar{b}/U)$ is isolated. By Lemma 1.5 it follows that $\mathrm{tp}(\bar{b}/U \cup \bar{a})$ is also isolated. □

Lemma 1.7. *Let $A \subseteq B \subseteq C$. If C is atomic over B and B is atomic over A then C is atomic over A.*

Proof. Let $\bar{c} \subseteq C$ and suppose that $\mathrm{tp}(\bar{c}/B)$ is isolated by $\varphi(\bar{x}, \bar{b})$. Fix some formula $\psi(\bar{y}, \bar{a})$ isolating $\mathrm{tp}(\bar{b}/A)$. We claim that $\mathrm{tp}(\bar{c}/A)$ is isolated by the formula $\chi := \exists \bar{y}[\varphi(\bar{x}, \bar{y}) \wedge \psi(\bar{y}, \bar{a})]$.

Suppose that $\chi \in \mathrm{tp}(\bar{d}/A)$. Then there is some tuple $\bar{e}$ with

$$\varphi(\bar{d}, \bar{e}), \psi(\bar{e}, \bar{a}) \in \mathrm{tp}(\bar{d}\bar{e}/A)\,.$$

Consequently, we have $\mathrm{tp}(\bar{e}/A) = \mathrm{tp}(\bar{b}/A)$ and there exists an A-automorphism π with $\pi(\bar{e}) = \bar{b}$. Let $\bar{d}' := \pi(\bar{d})$. Then $\mathrm{tp}(\bar{d}'/\bar{b}) = \mathrm{tp}(\bar{d}/\bar{e})$ and $\varphi(\bar{x}, \bar{b}) \in \mathrm{tp}(\bar{d}'/\bar{b})$ implies that $\mathrm{tp}(\bar{d}'/B) = \mathrm{tp}(\bar{c}/B)$. It follows that

$$\mathrm{tp}(\bar{d}/A) = \mathrm{tp}(\bar{d}'/A) = \mathrm{tp}(\bar{c}/A)\,. \qquad \square$$

The following two remarks follow immediately from the definition of an atomic model.

Lemma 1.8. (a) *Every elementary substructure of an atomic model is atomic.*

(b) *The union of an elementary chain of atomic models is atomic.*

2. The Omitting Types Theorem

We have seen in Section C2.4 how to build structures from a given set of formulae. In order to find structures realising only certain types we take a closer look at this construction. First, let us determine a minimal set of sorts a model has to realise.

Lemma 2.1. *Let Σ be an S-sorted signature and $T \subseteq \mathrm{FO}^0[\Sigma]$ a first-order theory. There exists a minimal set $S_0 \subseteq S$ such that T has a model $\mathfrak{A}$ with*

$$A_s \neq \varnothing \quad \text{iff} \quad s \in S_0 .$$

Proof. Let $\mathcal{S}$ be the class of all sets $S_0 \subseteq S$ such that T has a model $\mathfrak{A}$ with $A = \bigcup_{s \in S_0} A_s$. It is sufficient to show that the partial order $\langle \mathcal{S}, \supseteq \rangle$ is inductively ordered. Let $(S_i)_{i \in I}$ be a decreasing sequence of sets $S_i \in \mathcal{S}$ and set $S_\infty := \bigcap_i S_i$. We claim that $S_\infty \in \mathcal{S}$. Let

$$\Phi := T \cup \{ \eta_s \mid s \in S \smallsetminus S_\infty \} ,$$

where $\eta_s := \neg \exists x_s (x_s = x_s)$ states that there are no elements of sort s. Every model of Φ witnesses that $S_\infty \in \mathcal{S}$.

To prove that Φ is satisfiable let $\Phi_0 \subseteq \Phi$ be finite. Then there are sorts $s_0, \dots, s_n \in S_\infty$ such that

$$\Phi_0 \subset T \cup \{ \eta_{s_0}, \dots, \eta_{s_n} \} .$$

Hence, we can find some index $i \in I$ with $s_0, \dots, s_n \in S \smallsetminus S_i$. By assumption there is some S_i-sorted model $\mathfrak{A}$ of T. It follows that $\mathfrak{A} \vDash \Phi_0$. □

We have seen in Section C2.4 how to construct Herbrand models from Hintikka sets. To refine this construction we introduce a special kind of Hintikka set called a *Henkin set.*

Definition 2.2. Let $\Phi \subseteq \mathrm{FO}^0[\Sigma]$ be a set of sentences and $C \subseteq \Sigma$ a set of constant symbols.

(a) Φ has the *Henkin property* with respect to C if, for every formula $\varphi(x) \in \mathrm{FO}^1[\Sigma]$, there is some constant $c \in C$ such that

$$\exists x \varphi(x) \rightarrow \varphi(c) \in \Phi .$$

(b) We say that Φ is a *Henkin set* for a set $\Phi_0 \subseteq \mathrm{FO}^0[\Sigma]$ with respect to C if $\Phi_0 \subseteq \Phi$, Φ is complete, and Φ has the Henkin property with respect to C.

Lemma 2.3. *Every Henkin set is a Hintikka set.*

Corollary 2.4. *Every Henkin set Φ with respect to C has a Herbrand model $\mathfrak{H}$ where every element is denoted by some constant from C.*

Proof. We have seen in Lemma C2.4.6 that Φ has a Herbrand model $\mathfrak{H}$ where every element is denoted by some term. Since Φ is a Hintikka set, we can find, for every term t a constant $c \in C$ with

$$\exists x(x = t) \to c = t \in \Phi .$$

Therefore, every element is denoted by some constant in C. □

The class of all Henkin sets is in one-to-one correspondence with the class of all Herbrand models. In the next lemma we prove that this class forms a co-meagre set in the type topology.

Lemma 2.5. *Suppose that Σ is a countable signature, $T \subseteq \mathrm{FO}^{\circ}[\Sigma]$ a theory, and, for every sort s, let C_s be a countably infinite set of constant symbols of sort s with $C_s \cap \Sigma = \varnothing$. Set $C := \bigcup_s C_s$ and*

$$S^{\circ}_C(T) := S(\mathrm{FO}^{\circ}[\Sigma_C]/T) .$$

(a) *The complement of the set*

$$H(T) := \{\, \mathfrak{p} \in S^{\circ}_C(T) \mid \mathfrak{p} \text{ is a Henkin set for } T \text{ w.r.t. } C \,\}$$

is meagre in $S^{\circ}_C(T)$.

(b) *If $\bar{s}$ is a finite tuple of sorts and $\Phi \subseteq \mathrm{FO}^{\bar{s}}[\Sigma]$ is a set such that $\langle \Phi \rangle_{S^{\bar{s}}(T)}$ is nowhere dense then the complement of*

$$O(\Phi) := \{\, \mathfrak{p} \in S^{\circ}_C(T) \mid \text{for every } \bar{c} \in C^{<\omega}, \text{ there is some } \varphi \in \Phi \text{ with } \neg\varphi(\bar{c}) \in \mathfrak{p} \,\}$$

is meagre in $S^{\circ}_C(T)$.

Proof. (a) We have

$$H(T) = \bigcap_{\varphi \in \mathrm{FO}^1[\Sigma_C]} H_\varphi \quad \text{where} \quad H_\varphi = \bigcup_{c \in C} \langle \exists x \varphi(x) \to \varphi(c) \rangle_{S^\circ_C(T)} .$$

Since $\mathrm{FO}^1[\Sigma_C]$ is countable, we can show that the complement of $H(T)$ is meagre by proving that the complement of each H_φ is nowhere dense. Because H_φ is open, it is sufficient to show that its complement has empty interior, that is, that H_φ is dense.

Let $\langle \psi \rangle_{S^\circ_C(T)}$ be a nonempty basic open set and fix some model

$$\mathfrak{M} \vDash T \cup \{\psi\} .$$

Choose some element $a \in M$ with

$$\mathfrak{M} \vDash \exists x \varphi(x) \to \varphi(a) .$$

Let $D \subseteq C$ be the set of constant symbols appearing in ψ or φ. This set is finite and we have

$$\mathfrak{M}|_{\Sigma_D} \vDash T \cup \{\psi, \exists x \varphi(x) \to \varphi(a)\} .$$

Fix some constant symbol $c \in C \smallsetminus D$ of the same sort as a and let $\mathfrak{N}$ be a Σ_C-expansion of $\mathfrak{M}|_{\Sigma_D}$ with $c^{\mathfrak{N}} = a$. Then

$$\mathfrak{N} \vDash T \cup \{\psi, \exists x \varphi(x) \to \varphi(c)\} .$$

Hence, $\mathrm{Th}(\mathfrak{N}) \in \langle \psi \rangle_{S^\circ_C(T)} \cap H_\varphi \neq \emptyset$.

(b) We have

$$O(\Phi) = \bigcap_{\bar{c} \in C^{<\omega}} O_{\bar{c}} \quad \text{where} \quad O_{\bar{c}} = \bigcup_{\varphi \in \Phi} \langle \neg \varphi(\bar{c}) \rangle_{S^\circ_C(T)} .$$

As above it is sufficient to prove that each set $O_{\bar{c}}$ is dense. Consider a nonempty basic open set $\langle \psi(\bar{c}, \bar{d}) \rangle_{S^\circ_C(T)}$ where $\psi \in \mathrm{FO}[\Sigma]$ and $\bar{d} \subseteq C \smallsetminus \bar{c}$. Fix some model $\mathfrak{M} \vDash T \cup \{\psi(\bar{c}, \bar{d})\}$. Then

$$\langle \mathfrak{M}|_\Sigma, \bar{c} \rangle \vDash T \cup \{\exists \bar{y} \psi(\bar{x}, \bar{y})\} .$$

Hence, $\langle \exists \bar{y}\psi(\bar{x}, \bar{y}) \rangle_{S^{\bar{s}}(T)} \neq \varnothing$. Since $\langle \Phi \rangle_{S^{\bar{s}}(T)}$ is nowhere dense it follows that

$$\langle \exists \bar{y}\psi(\bar{x}, \bar{y}) \rangle_{S^{\bar{s}}(T)} \setminus \langle \Phi \rangle_{S^{\bar{s}}(T)} \neq \varnothing .$$

Fix some model $\langle \mathfrak{N}_0, \bar{a} \rangle$ with

$$\mathrm{Th}(\mathfrak{N}_0, \bar{a}) \in \langle \exists \bar{y}\psi(\bar{x}, \bar{y}) \rangle_{S^{\bar{s}}(T)} \setminus \langle \Phi \rangle_{S^{\bar{s}}(T)} .$$

There is some formula $\varphi \in \Phi$ such that

$$\mathfrak{N}_0 \not\models \varphi(\bar{a}) .$$

Furthermore, we can find a tuple $\bar{b} \subseteq N_0$ with

$$\mathfrak{N}_0 \models \psi(\bar{a}, \bar{b}) .$$

Let $\mathfrak{N}$ be a Σ_C-expansion of $\mathfrak{N}_0$ with $\bar{c}^{\mathfrak{N}} = \bar{a}$ and $\bar{d}^{\mathfrak{N}} = \bar{b}$. Then we have

$$\mathrm{Th}(\mathfrak{N}) \in \langle \psi \rangle_{S^0_C(T)} \cap O_{\bar{c}} \neq \varnothing . \qquad \square$$

After these preparations we can prove that every meagre set of types is omitted in some model.

Theorem 2.6 (Omitting Types Theorem). *Let Σ be a countable S-sorted signature and $T \subseteq \mathrm{FO}[\Sigma]$ a countable first-order theory. For every $\bar{s} \in S^{<\omega}$, let $X_{\bar{s}} \subseteq S^{\bar{s}}(T)$ be a meagre set of types. There exists a model of T that omits every type in $\bigcup_{\bar{s}} X_{\bar{s}}$.*

Proof. For every sort s, fix a countably infinite set C_s of constant symbols disjoint from Σ. Each set $X_{\bar{s}}$ can be written as $X_{\bar{s}} = \bigcup_{n<\omega} X^n_{\bar{s}}$, where $X^n_{\bar{s}}$ is nowhere dense. Let $\Phi^n_{\bar{s}}$ be a set of formulae such that $\langle \Phi^n_{\bar{s}} \rangle = \mathrm{cl}(X^n_{\bar{s}})$. By the preceding lemma, we know that

$$Y := H(T) \cap \bigcap_{\bar{s} \in S^{<\omega}} \bigcap_{n<\omega} O(\Phi^n_{\bar{s}})$$

is a countable intersection of sets whose complement is meagre. Hence, the complement of Y is meagre. By Theorem B5.5.8 it follows that Y itself is also dense. Fix some type $\mathfrak{p} \in Y$.

By Corollary 2.4, there exists a Herbrand model $\mathfrak{H}$ of $\mathfrak{p}$ where every element is denoted by some constant in C. If $\bar{a} \in H^{\bar{s}}$ is a finite tuple denoted by the constants $\bar{c} \subseteq C$ then we have

$$\operatorname{tp}(\bar{a}) = \{ \varphi(\bar{x}) \mid \varphi(\bar{c}) \in \mathfrak{p} \} \notin X_{\bar{s}} .$$

Hence, no tuple in $\mathfrak{H}$ realises a type in $X_{\bar{s}}$. □

Corollary 2.7. *Let Σ be a countable signature and $T \subseteq \mathrm{FO}[\Sigma]$ a first-order theory. Let $\mathfrak{p}_n$, $n < \omega$, be a sequence of non-isolated partial types over T. There exists a model of T that omits every $\mathfrak{p}_n$, $n < \omega$.*

Let us give a simple example showing that the Omitting Types Theorem fails for uncountable theories.

Example. Let $\Sigma := \{ c_i \mid i < \omega_1 \} \cup \{ d_n \mid n < \omega \}$ be a signature of constant symbols and let

$$T := \{ c_i \neq c_k \mid i \neq k \} \cup \{ d_i \neq d_k \mid i \neq k \}$$

be the theory stating that the values of the c_i are distinct and that the values of the d_n are distinct. Consider the partial 1-type

$$\Phi := \{ x \neq d_n \mid n < \omega \} .$$

This type is not isolated since there is no formula $\varphi(x)$ implying that x is different from all constants d_n. On the other hand, every model of T has uncountably many elements and, therefore, realises Φ.

Theorem 2.8. *Let T be a countable complete theory with infinite models. There exists a family $(\mathfrak{M}_\xi)_{\xi < 2^{\aleph_0}}$ of models of T such that every complete type that is realised in at least two of the models is isolated.*

Proof. For every sort s, fix a countably infinite set C_s of constant symbols disjoint from Σ. Set $C := \dot\bigcup_s C_s$ and let $(\varphi_n)_n$ be an enumeration of $\mathrm{FO}^1[\Sigma_C]$. We fix an enumeration $\langle u_n, \bar{c}^n, \bar{d}^n \rangle_{n<\omega}$ of all triples in $2^{<\omega} \times C^{<\omega} \times C^{<\omega}$ such that $\bar{c}^n$ and $\bar{d}^n$ have the same length and the same sorts. We assume that the enumeration has been chosen such that every triple appears infinitely often in the sequence.

We construct an increasing chain $T_0 \subseteq T_1 \subseteq \dots$ of finite trees $T_n \subseteq 2^{<\omega}$ and, for each $w \in 2^{<\omega}$, we define a finite set $\Phi_w \subseteq \mathrm{FO}^0[\Sigma_C]$ of formulae such that $\Phi_u \subseteq \Phi_w$, for $u \leq w$.

We start with $T_0 := \{\langle\rangle\}$ and $\Phi_{\langle\rangle} := \varnothing$. For the inductive step, suppose that we have already defined T_n and Φ_w, for $w \in T_n$. To define T_{n+1} we distinguish two cases. If $u_n \notin T_n$ then we simply set

$$T_{n+1} := \{ w0 \mid w \text{ a leaf of } T_n \},$$

and, for every leaf w of T_n,

$$\Phi_{w0} := \Phi_w \cup \{\exists x \varphi_n \to \varphi_n(c)\},$$

where $c \in C$ is some new constant symbol not appearing in any formula of Φ_w.

It remains to consider the case that $u_n \in T_n$. Let $v_0, \dots, v_{l-1}$ be an enumeration of all leaves v of T_n with $u_n \leq v$, and let $w_0, \dots, w_{m-1}$ be an enumeration of all leaves w with $u_n \nleq w$. We define sets

$$\begin{aligned} \Phi_{w_i} &= \Psi^i_{-1} \subseteq \Psi^i_0 \subseteq \dots \subseteq \Psi^i_{l-1}, \quad && \text{for } i < m, \\ \Phi_{v_j} &= \Theta^j_{-1} \subseteq \Theta^j_0 \subseteq \dots \subseteq \Theta^j_{m-1}, \quad && \text{for } j < l, \end{aligned}$$

as follows. We start with $\Psi^i_{-1} := \Phi_{w_i}$ and $\Theta^j_{-1} := \Phi_{v_j}$. Suppose that we have already defined Ψ^i_j and Θ^j_i, for all pairs $\langle i, j\rangle$ lexicographically less than $\langle i_0, j_0\rangle$. To define $\Psi^{i_0}_{j_0}$ and $\Theta^{j_0}_{i_0}$ we set

$$\psi(\bar{c}^n, \bar{e}) := \bigwedge \Psi^{i_0}_{j_0-1} \quad \text{and} \quad \vartheta(\bar{d}^n, \bar{f}) := \bigwedge \Theta^{j_0}_{i_0-1},$$

where $\bar{e} \subseteq C$ contains all constants in $\Psi^{i_0}_{j_0-1}$ different from $\bar{c}^n$, and $\bar{f} \subseteq C$ contains all constants in $\Theta^{j_0}_{i_0-1}$ different from $\bar{d}^n$. If $\langle \exists \bar{y} \psi(\bar{x}, \bar{y}) \rangle_{S(T)}$ is a singleton then we set

$$\Psi^{i_0}_{j_0} := \Psi^{i_0}_{j_0-1} \quad \text{and} \quad \Theta^{j_0}_{i_0} := \Theta^{j_0}_{i_0-1} .$$

Otherwise, we choose some type $\mathfrak{q} \in \langle \exists \bar{y} \vartheta(\bar{x}, \bar{y}) \rangle_{S(T)}$. By assumption, we can find a type $\mathfrak{p} \in \langle \exists \bar{y} \psi(\bar{x}, \bar{y}) \rangle_{S(T)}$ different from $\mathfrak{q}$. We fix some formula $\eta(\bar{x}) \in \mathfrak{p} \setminus \mathfrak{q}$ and set

$$\Psi^{i_0}_{j_0} := \Psi^{i_0}_{j_0-1} \cup \{\eta(\bar{c}^n)\} \quad \text{and} \quad \Theta^{j_0}_{i_0} := \Theta^{j_0}_{i_0-1} \cup \{\neg\eta(\bar{d}^n)\} .$$

Having defined all Ψ^i_j and Θ^j_i we set

$$\Phi'_{w_i} := \Psi^i_{l-1} \cup \{\exists x \varphi_n \to \varphi_n(c)\} ,$$
$$\Phi'_{v_j} := \Theta^j_{m-1} \cup \{\exists x \varphi_n \to \varphi_n(c)\} ,$$

where $c \in C$ is some constant not appearing in any set Ψ^i_j or Θ^j_i. Let $z_0, \dots, z_{k-1}$ be an enumeration of all leaves z of T_n such that the set $\langle \Phi'_z \rangle_{S(T(C))}$ contains at least two types, and let $u_0, \dots, u_{r-1}$ be an enumeration of all other leaves of T_n. We define

$$T_{n+1} := T_n \cup \{ z_i b \mid i < k,\ b \in [2] \} \cup \{ u_i 0 \mid i < r \} ,$$

and $\Phi_{u_i 0} := \Phi'_{u_i}$, for $i < r$. For each $i < k$, we chose distinct types $\mathfrak{p}_i, \mathfrak{q}_i \in \langle \Phi'_{z_i} \rangle_{S(T(C))}$ and some formula $\eta_i \in \mathfrak{p}_i \setminus \mathfrak{q}_i$. Then we set

$$\Phi_{z_i 0} := \Phi'_{z_i} \cup \{\neg\eta_i\} \quad \text{and} \quad \Phi_{z_i 1} := \Phi'_{z_i} \cup \{\eta_i\} .$$

This completes the construction of T_{n+1}. To define the models $\mathfrak{M}_\xi$ let $T_\omega := \bigcup_n T_n$. A sequence $\beta \in 2^\omega$ is a *branch* of T_ω if $\beta \upharpoonright n \in T_\omega$, for all $n < \omega$. For each branch β of T_ω, we define a sequence $\beta^* \in 2^{<\omega}$ as follows. Let

$$I := \{ n < \omega \mid (\beta \upharpoonright n)0 \in T_\omega \text{ and } (\beta \upharpoonright n)1 \in T_\omega \} ,$$

and let $n_0 < n_1 < \dots$ be an enumeration of I. We define $\beta^* \in 2^{|I|}$ by

$$\beta^*(i) := \beta(n_i)\,, \quad \text{for } i < |I|\,.$$

For each $\xi \in 2^\omega$, there is a unique branch β_ξ with $\beta^*_\xi \leq \xi$. We define

$$\Psi_\xi := \bigcup_{n<\omega} \Phi_{\beta_\xi \restriction n}\,.$$

It follows by compactness that each set Ψ_ξ is satisfiable. Furthermore, the above construction ensures that each of these sets has the Henkin property with respect to C. Hence, we can use Corollary 2.4 to find a Herbrand model $\mathfrak{M}_\xi$ of Ψ_ξ.

It remains to prove that every type realised in two different models is isolated. Suppose that

$$\operatorname{tp}(\bar c/\mathfrak{M}_\xi) = \operatorname{tp}(\bar d/\mathfrak{M}_\zeta) \quad \text{where} \quad \xi \neq \zeta\,.$$

If β^*_ξ is finite then $\langle \Phi_{\beta^*_\xi} \rangle_{S(T(C))} = \{\mathfrak{p}\}$ is a singleton and every type realised in $\mathfrak{M}_\xi \vDash \Phi_{\beta^*_\xi}$ is isolated. Similarly, if β^*_ζ is finite then $\operatorname{tp}(\bar d/\mathfrak{M}_\zeta)$ is isolated.

Hence, suppose that β^*_ξ and β^*_ζ are both infinite. Then there is some $n < \omega$ such that

$$\bar c^n = \bar c\,, \quad \bar d^n = \bar d\,, \quad u_n \in T_n\,, \quad \text{and} \quad \beta_\xi \sqcap \beta_\zeta \prec u_n \prec \beta_\zeta\,.$$

Let w be the leaf of T_n with $w \prec \beta_\xi$ and let v be the leaf with $v \prec \beta_\zeta$. By construction of T_{n+1} it follows that either there is a formula isolating $\operatorname{tp}(\bar c/\mathfrak{M}_\xi)$, or there is some formula $\eta(\bar c) \in \Phi'_w \subseteq \Psi_\xi$ with $\neg\eta(\bar d) \in \Phi'_v \subseteq \Psi_\zeta$. In the first case we are done, whereas in the second case we obtain $\operatorname{tp}(\bar c/\mathfrak{M}_\xi) \neq \operatorname{tp}(\bar d/\mathfrak{M}_\zeta)$, a contradiction. □

3. *Prime and atomic models*

Not every theory has atomic models, but for countable signatures we can use the Omitting Types Theorem to construct such models.

Theorem 3.1. *Let T be a countable complete theory. If $S^{\bar{s}}(T)$ is countable, for all finite tuples $\bar{s}$, then there exists a countable atomic model of T.*

Proof. For every $\bar{s}$, there are at most countably many non-isolated $\bar{s}$-types. Consequently, they form a meagre set and we can use the Omitting Types Theorem to find a model of T that realises none of them. □

Lemma 3.2. *Let T be a countable complete theory. If $|S^{\bar{s}}(T)| < 2^{\aleph_0}$, for all finite $\bar{s}$, then T has an atomic model over A, for every finite set A of parameters.*

Proof. By Corollary B5.7.5, it follows that each type space $S^{\bar{s}}(T)$ is countable. Let $\bar{a}$ be an enumeration of A. Since $\mathrm{tp}(\bar{b}/\bar{a})$ is determined by $\mathrm{tp}(\bar{b}\bar{a})$ it follows that $S^{\bar{s}}(A)$ is also countable. Hence, according to the preceding theorem $T(A)$ has an atomic model. □

If the type space is too large, atomic models might not exist.

Example. Consider the theory $T := \mathrm{Th}(\mathfrak{C})$ where $\mathfrak{C} := \langle 2^\omega, (P_n)_{n<\omega} \rangle$ and

$$P_n := \{ \alpha \in 2^\omega \mid \alpha(n) = 1 \} .$$

As we have seen in the example on page 534, the type space $S^1(T)$ is homeomorphic to the Cantor discontinuum 2^ω, which does not contain isolated points. Consequently, no type is isolated and T does not have atomic models.

Theorem 3.3. *Let T be a countable complete first-order theory. There exists an atomic model of T if, and only if, the set of isolated $\bar{s}$-types is dense in $S^{\bar{s}}(T)$, for every finite $\bar{s}$.*

Proof. Let $X \subseteq S^{\bar{s}}(T)$ be the set of all isolated $\bar{s}$-types. If T has an atomic model $\mathfrak{M}$ then X is the set of types realised in $\mathfrak{M}$. By Lemma C3.2.6 it follows that X is dense. Conversely, if X is dense then its complement $Y_{\bar{s}} := S^{\bar{s}}(T) \setminus X$ is closed and has empty interior. By the Omitting Types Theorem, there exists a model $\mathfrak{M}$ of T omitting all types in $\bigcup_{\bar{s}} Y_{\bar{s}}$. This model is atomic. □

Corollary 3.4. *Let T be a countable complete theory. If*

$$\mathrm{rk}_{\mathrm{CB}}(S^n(T)) < \infty\,, \quad \text{for all } n < \omega\,,$$

then T has an atomic model.

Proof. Immediately by Theorem 3.3 and Proposition B5.5.12. □

Intuitively, an atomic model is the opposite of a saturated one. The next lemma shows that these models also behave in the opposite way with respect to the relation $\sqsubseteq^{\aleph_0}_{\mathrm{FO}}$.

Lemma 3.5. (a) *If $\mathfrak{A}$ is atomic then we have $\mathfrak{A} \sqsubseteq^{\aleph_0}_{\mathrm{FO}} \mathfrak{B}$, for all $\mathfrak{B} \equiv \mathfrak{A}$.*

(b) *If $\mathfrak{A}$ is a structure with countable signature such that $\mathfrak{A} \sqsubseteq^{\aleph_0}_{\mathrm{FO}} \mathfrak{B}$, for all $\mathfrak{B} \equiv \mathfrak{A}$, then $\mathfrak{A}$ is atomic.*

Proof. (a) Suppose that

$$\langle \mathfrak{A}, \bar{a} \rangle \equiv_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b} \rangle\,.$$

We have to prove the forth property. Let $c \in A$ and choose some formula $\varphi(\bar{x}, y)$ isolating $\mathfrak{p} := \mathrm{tp}(\bar{a}c/\mathfrak{A})$. Then

$$\mathfrak{A} \vDash \exists y \varphi(\bar{a}, y) \quad \text{implies} \quad \mathfrak{B} \vDash \exists y \varphi(\bar{b}, y)\,.$$

Consequently, there exists some $d \in B$ such that $\mathfrak{B} \vDash \varphi(\bar{b}, d)$. It follows that $\mathrm{tp}(\bar{b}d/\mathfrak{B}) = \mathfrak{p}$ and, hence,

$$\langle \mathfrak{A}, \bar{a}c \rangle \equiv_{\mathrm{FO}} \langle \mathfrak{B}, \bar{b}d \rangle\,.$$

(b) Suppose that $\mathfrak{A}$ contains a finite tuple $\bar{a} \subseteq A$ whose type $\mathrm{tp}(\bar{a})$ is not isolated. By the Omitting Types Theorem there is a structure $\mathfrak{B} \equiv \mathfrak{A}$ omitting $\mathrm{tp}(\bar{a})$. If $\mathfrak{A} \sqsubseteq^{\aleph_0}_{\mathrm{FO}} \mathfrak{B}$ then there would be some tuple $\bar{b} \subseteq B$ such that $\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{B}, \bar{b} \rangle$. Consequently, $\mathrm{tp}(\bar{b}/\mathfrak{B}) = \mathrm{tp}(\bar{a}/\mathfrak{A})$ would be realised in $\mathfrak{B}$. Contradiction. □

Corollary 3.6. *If $\mathfrak{A} \equiv \mathfrak{B}$ are atomic then $\mathfrak{A} \equiv^{\aleph_0}_{\mathrm{FO}} \mathfrak{B}$.*

Corollary 3.7. *Every atomic model is $\aleph_0$-homogeneous.*

Proof. By the preceding corollary we have $\mathfrak{A} \equiv_{\mathrm{FO}}^{\aleph_0} \mathfrak{A}$, for every atomic structure $\mathfrak{A}$. □

If a countable theory T has atomic models then it has a unique countable one. Furthermore, this countable atomic model can be embedded into every other model of T.

Definition 3.8. A structure $\mathfrak{A}$ is a *prime model* of a theory T if, for ever model $\mathfrak{B} \vDash T$, there exists an elementary embedding $\mathfrak{A} \to \mathfrak{B}$. Similarly, we say that $\mathfrak{A}$ is *prime over* a set $U \subseteq A$ if it is a prime model of $T(U)$.

Example. $\mathfrak{N} = \langle \mathbb{N}, +, \cdot, 0, 1 \rangle$ is a prime model of arithmetic.

Remark. Only complete theories can have prime models.

Lemma 3.9. *If $\mathfrak{M}$ is a structure with $M = \mathrm{acl}(\varnothing)$ then $\mathfrak{M}$ is prime.*

Exercise 3.1. Prove the preceding lemma.

Lemma 3.10. *Every prime model with a countable signature is atomic.*

Proof. Let $\mathfrak{M}$ be a model of a theory T that realises a non-isolated type $\mathfrak{p}$. By the Omitting Types Theorem, there exists some model $\mathfrak{N} \vDash T$ in which $\mathfrak{p}$ is not realised. Therefore, there exists no embedding $\mathfrak{M} \to \mathfrak{N}$ and $\mathfrak{M}$ cannot be prime. □

Lemma 3.11. *Every countable atomic model is prime.*

Proof. Let $\mathfrak{A}$ be a countable atomic model and suppose that $\mathfrak{B} \equiv \mathfrak{A}$. Let $(a_i)_{i<\omega}$ be an enumeration of A. Since $\mathfrak{A} \sqsubseteq_{\mathrm{FO}}^{\aleph_0} \mathfrak{B}$ we can find, by Lemma C4.4.9, an enumeration $(b_i)_{i<\omega}$ such that

$$\langle \mathfrak{A}, (a_i)_{i<n} \rangle \equiv_{\mathrm{FO}} \langle \mathfrak{B}, (b_i)_{i<n} \rangle , \qquad \text{for all } n < \omega .$$

Let $p_n : (a_i)_{i<n} \mapsto (b_i)_{i<n} \in I_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ be the corresponding partial isomorphisms. Since $I_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$ is $\aleph_1$-complete we have $p := \bigcup_n p_n \in I_{\mathrm{FO}}(\mathfrak{A}, \mathfrak{B})$. As $\mathrm{dom}\, p = A$ it follows that p is the desired elementary embedding of $\mathfrak{A}$ into $\mathfrak{B}$. □

The next theorem summarises the relation between prime and atomic models.

Theorem 3.12. *Let T be a countable complete theory.*

(a) *Every prime model of T is countable and atomic.*

(b) *Every countable atomic model of T is prime.*

(c) *T has a prime model if and only if it has an atomic model.*

(d) *All prime models of T are isomorphic.*

Proof. (a) and (b) were proved in Lemmas 3.10 and 3.11, respectively.

(c) By (a), every prime model is atomic. Conversely, if T has an atomic model then it also has a countable one, by the theorem of Löwenheim and Skolem. Hence, the claim follows by (b).

(d) If $\mathfrak{A}$ and $\mathfrak{B}$ are prime models of T then we have $\mathfrak{A} \cong_{\mathrm{FO}}^{\aleph_0} \mathfrak{B}$, by (a) and Corollary 3.6. Since $\mathfrak{A}$ and $\mathfrak{B}$ are countable, Lemma C4.4.10 implies that $\mathfrak{A} \cong \mathfrak{B}$. □

4. Constructible models

For uncountable signatures we cannot use the Omitting Types Theorem to construct prime models. In this section we present an alternative way to obtain such models.

Definition 4.1. Let $\mathfrak{M}$ be a structure and $A, U \subseteq M$.

(a) A *construction* of A over U is an enumeration $(a_i)_{i<\gamma}$ of A such that

$$\mathrm{tp}(a_\alpha/U \cup a[<\alpha]) \text{ is isolated}, \quad \text{for all } \alpha < \gamma,$$

where $a[<\alpha] := \{ a_i \mid i < \alpha \}$.

(b) If there exists a construction of A over U we say that A is *constructible* over U.

Example. Let T_{eq} be the theory of all infinite structures with empty signature. This theory has exactly one model of every infinite cardinality.

The countable model $\mathfrak{M}_{\aleph_0}$ of T_{eq} is constructible. If $(a_n)_{n<\omega}$ is an enumeration of $M_{\aleph_0}$ then $\mathrm{tp}(a_n/a_0 \dots a_{n-1})$ is isolated by the formula

$$x \neq a_0 \wedge \dots \wedge x \neq a_{n-1} .$$

Every uncountable model $\mathfrak{M}$ of T_{eq} is not constructible since, for every enumeration $(a_\alpha)_{\alpha<\gamma}$ of M, the type $\mathrm{tp}(a_\omega/a[<\omega])$ is not isolated.

We start by showing that constructible models are prime and atomic.

Lemma 4.2. *If $A \subseteq M$ is constructible over U then A is atomic over U.*

Proof. Let $(a_\alpha)_{\alpha<\gamma}$ be a construction of A over U. We prove by induction on α that $a[<\alpha]$ is atomic over U. For $\alpha = 0$ there is nothing to do. If α is a limit ordinal then any finite tuple in $a[<\alpha] = \bigcup_{\beta<\alpha} a[<\beta]$ belongs to some $a[<\beta]$ with $\beta < \alpha$. Hence, the claim follows immediately by inductive hypothesis.

For the inductive step, note that $a[<\alpha+1] = a[<\alpha] \cup \{a_\alpha\}$ is atomic over $U \cup a[<\alpha]$ and $U \cup a[<\alpha]$ is atomic over U. By Lemma 1.7, it follows that $a[<\alpha+1]$ is atomic over U. □

Proposition 4.3. *Let $\mathfrak{M}$ be a model of a complete theory T and let $U \subseteq M$ be a set such that M is constructible over U.*

(a) *$\mathfrak{M}$ is a prime model over U.*

(b) $|M| \leq |U| \oplus |T|$.

Proof. (a) Let $(a_\alpha)_{\alpha<\gamma}$ be a construction of M over U. Suppose that $\mathfrak{N}$ is a model of $T(U)$. We construct a sequence $(b_\alpha)_{\alpha<\gamma}$ as follows. Suppose that b_i has already been defined for all $i < \alpha$. Since the type $\mathrm{tp}(a_\alpha/U \cup a[<\alpha])$ is isolated, there exists some element $b_\alpha \in N$ with

$$b_\alpha b[<\alpha] \equiv_U a_\alpha a[<\alpha] .$$

The mapping $a_\alpha \mapsto b_\alpha$ is the desired elementary embedding $\mathfrak{M} \to \mathfrak{N}$.

(b) By the Theorem of Löwenheim and Skolem, $T(U)$ has a model $\mathfrak{N}$ of size $|N| \leq |U| \oplus |T|$. By (a), there exists an embedding $\mathfrak{M} \to \mathfrak{N}$. Consequently, $|M| \leq |N| \leq |U| \oplus |T|$. ☐

Our next aim is to prove that constructible models are unique, up to isomorphism.

Definition 4.4. Let $(a_\alpha)_{\alpha<\gamma}$ be a construction of A over U. A set $C \subseteq A$ is *closed* (w.r.t. this construction) if, for every $\alpha < \gamma$ with $a_\alpha \in C$, the type $\mathrm{tp}(a_\alpha/U \cup a[<\alpha])$ is isolated by some formula $\varphi(x; \bar{c})$ with parameters $\bar{c} \subseteq U \cup (C \cap a[<\alpha])$.

Lemma 4.5. *Let $(a_\alpha)_{\alpha<\gamma}$ be a construction of A over U.*

(a) *If $C, D \subseteq A$ are closed, then so is $C \cup D$.*

(b) *Every element $a \in A$ is contained in a finite closed set $C \subseteq A$.*

(c) *Every closed subset of A is constructible.*

Proof. (a) is immediate.

(b) By induction on $\alpha < \gamma$, we construct a finite closed set C_α containing a_α. For $\alpha = 0$, we can set $C_0 := \{a_0\}$ since $\mathrm{tp}(a_0/U)$ is isolated by some formula with parameters in U. For the inductive step, suppose that we have already defined C_i, for all $i < \alpha$. Fix a formula $\varphi(x; \bar{c})$ with parameters $\bar{c} \subseteq U \cup a[<\alpha]$ isolating $\mathrm{tp}(a_\alpha/U \cup a[<\alpha])$. Let $I := \{\, i < \alpha \mid a_i \in \bar{c} \,\}$. The set

$$C_\alpha := \{a_\alpha\} \cup \bigcup_{i \in I} C_i$$

is finite and closed.

(c) Let $(a_\alpha)_{\alpha<\gamma}$ be a construction of A over U, $C \subseteq A$ a closed set, and set $C_{<\alpha} := C \cap a[<\alpha]$. For $a_\alpha \in C$, the type $\mathrm{tp}(a_\alpha/U \cup a[<\alpha])$ is isolated by some formula $\varphi_\alpha(x, \bar{c})$ with $\bar{c} \subseteq U \cup (C \cap a[<\alpha]) = U \cup C_{<\alpha}$. Consequently, this formula also isolates the type $\mathrm{tp}(a_\alpha/U \cup C_{<\alpha})$. Hence, $\mathrm{tp}(a_\alpha/U \cup C_{<\alpha})$ is isolated, for all $a_\alpha \in C$, and we obtain a construction of C by omitting form $(a_\alpha)_{\alpha<\gamma}$ all elements that are not in C. ☐

Lemma 4.6. *Let* $(a_\alpha)_{\alpha<\gamma}$ *be a construction of* A *over* U, C *a closed subset of* A, $\bar c$ *an enumeration of* C, *and, for every* $a_\alpha \in C$, *let* $\varphi_\alpha(x_\alpha;\bar b_\alpha)$ *be a formula isolating* $\mathrm{tp}(a_\alpha/U \cup a[<\alpha])$. *Then*

$$T(U) \cup \{\, \varphi_\alpha(x_\alpha;\bar b_\alpha) \mid a_\alpha \in C \,\} \vDash \mathrm{tp}(\bar c/U)\,.$$

Proof. Note that $C_{<\alpha} := C \cap a[<\alpha]$ is closed. Hence, we can prove the claim by induction on α. For $\alpha = 0$ we have $\mathrm{tp}(\langle\rangle/U) = T(U)$. If α is a limit ordinal then the claim follows by inductive hypothesis since every formula refers only to finitely many elements of $C_{<\alpha}$. For the successor step, suppose that $\bar c = \bar c' a_\alpha$ where $\bar c'$ is an enumeration of $C_{<\alpha}$. By inductive hypothesis, we know that

$$T(U) \cup \{\, \varphi_i(x_i;\bar b_i) \mid i < \alpha,\ a_i \in C \,\} \vDash \mathrm{tp}(\bar c'/U)\,.$$

Furthermore,

$$T(U) \cup \{\varphi_\alpha(x_\alpha;\bar b_\alpha)\} \vDash \mathrm{tp}(a_\alpha/U \cup a[<\alpha]) \vDash \mathrm{tp}(a_\alpha/U \cup \bar c')\,.$$

Combining these two implications, the claim follows. □

Proposition 4.7. *Let* C *be a closed subset of a constructible set* A. *Then* A *is constructible over* C.

Proof. We start by showing that A is atomic over C. Let $A_0 \subseteq A$ be finite. By Lemma 4.5 (b), we can find a finite closed set D containing A_0. For $X \subseteq A$, set

$$\Phi(X) := \{\, \varphi_\beta(x_\beta;\bar b_\beta) \mid a_\beta \in X \,\}\,,$$

where $\varphi_\beta(x_\beta;\bar b_\beta)$ is some formula isolating $\mathrm{tp}(a_\beta/a[<\beta])$. According to Lemma 4.6 we have

$$T \cup \Phi(\bar b) \vDash \mathrm{tp}(\bar b)\,, \quad \text{for every closed set } \bar b \subseteq A\,.$$

In particular, we have

$$T \cup \Phi(C \cup D) \vDash \mathrm{tp}(\bar c\bar d)\,,$$

where $\bar{c}$ is an enumeration of C and $\bar{d}$ one of D. As $\Phi(C) \subseteq T(C)$, it follows that

$$T(C) \cup \Phi(D) \vDash \mathrm{tp}(\bar{d}/C) .$$

Hence, $\mathrm{tp}(\bar{d}/C)$ is isolated by the formula $\bigwedge \Phi(D)$. In particular, the type of A_0 over C is isolated.

To conclude the proof, let $(a_\alpha)_{\alpha<\gamma}$ be a construction of A. We prove that it is also a construction over C. Let $\alpha < \gamma$. Since $a[<\alpha]$ is closed, so is $C \cup a[<\alpha]$. By the first part of the proof, it follows that a_α is atomic over $C \cup a[<\alpha]$. ☐

Lemma 4.8. *If $(a_\alpha)_{\alpha<\gamma}$ is a construction of A over U then it is also a construction of A over $U \cup C$, for every finite subset $C \subseteq A$.*

Proof. By Lemma 4.2, A is atomic over $U \cup a[<\alpha]$, for every $\alpha < \gamma$. In particular, $C \cup \{a_\alpha\}$ is atomic over $U \cup a[<\alpha]$. By Lemma 1.5, it follows that a_α is atomic over $U \cup a[<\alpha] \cup C$. ☐

To prove the uniqueness of constructible models, we employ back-and-forth arguments.

Definition 4.9. Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures such that A and B are constructible over $\varnothing$. We define

$$I_{\mathrm{cl}}(\mathfrak{A},\mathfrak{B}) := \{\, p \in I_{\mathrm{FO}}(\mathfrak{A},\mathfrak{B}) \mid \mathrm{dom}\, p \text{ and } \mathrm{rng}\, p \text{ are closed} \,\} .$$

Lemma 4.10. *Suppose that $\mathfrak{A}$ and $\mathfrak{B}$ are structures where A and B are constructible over $\varnothing$. Then $I_{\mathrm{cl}}(\mathfrak{A},\mathfrak{B})$ is $\aleph_1$-bounded and it has the back-and-forth property with respect to itself.*

Proof. By symmetry, we only consider the forth property. Let $\bar{a} \mapsto \bar{b} \in I_{\mathrm{cl}}(\mathfrak{A},\mathfrak{B})$ and $x \in A$. By induction on n, we construct finite tuples $\bar{c}_n \subseteq A$ and $\bar{d}_n \subseteq B$ such that $\bar{a}\bar{c}_0\bar{c}_1 \cdots \mapsto \bar{b}\bar{d}_0\bar{d}_1 \cdots \in I_{\mathrm{cl}}(\mathfrak{A},\mathfrak{B})$, $x \in \bar{c}_0$, and

$$\langle \mathfrak{A}, \bar{a}\bar{c}_0 \dots \bar{c}_{n-1} \rangle \equiv \langle \mathfrak{B}, \bar{b}\bar{d}_0 \dots \bar{d}_{n-1} \rangle , \quad \text{for all } n < \omega .$$

We start with some finite closed set $\bar{c}_0$ containing x. For the inductive step, suppose that we have already defined $\bar{c}_0, \dots, \bar{c}_n$ and $\bar{d}_0, \dots, \bar{d}_{n-1}$ such that

$$\langle \mathfrak{A}, \bar{a}\bar{c}_0 \dots \bar{c}_{n-1} \rangle \equiv \langle \mathfrak{B}, \bar{b}\bar{d}_0 \dots \bar{d}_{n-1} \rangle .$$

Since $\mathfrak{A}$ is atomic over $\bar{a}$, we know that the type $\operatorname{tp}(\bar{c}_0 \dots \bar{c}_{n-1}\bar{c}/\bar{a})$ is isolated. By Lemma 1.5, it follows that the type $\operatorname{tp}(\bar{c}_n/\bar{a}\bar{c}_0 \dots \bar{c}_{n-1})$ is also isolated. As

$$\langle \mathfrak{A}, \bar{a}\bar{c}_0 \dots \bar{c}_{n-1} \rangle \equiv \langle \mathfrak{B}, \bar{b}\bar{d}_0 \dots \bar{d}_{n-1} \rangle ,$$

we can therefore find some tuple $\bar{d}_n \subseteq B$ with

$$\langle \mathfrak{A}, \bar{a}\bar{c}_0 \dots \bar{c}_{n-1}\bar{c}_n \rangle \equiv \langle \mathfrak{B}, \bar{b}\bar{d}_0 \dots \bar{d}_{n-1}\bar{d}_n \rangle .$$

If $\bar{b}\bar{d}_0 \dots \bar{d}_n$ is closed then we can stop. Otherwise, let $\bar{d}_{n+1}$ be a finite closed set containing $\bar{d}_n$. Again, since $\bar{b}\bar{d}_0 \dots \bar{d}_{n-1}$ is closed and the type $\operatorname{tp}(\bar{d}_{n+1}/\bar{b}\bar{d}_0 \dots \bar{d}_n)$ is isolated, we can find a tuple $\bar{c}_{n+1} \subseteq A$ such that

$$\langle \mathfrak{A}, \bar{a}\bar{c}_0 \dots \bar{c}_n\bar{c}_{n+1} \rangle \equiv \langle \mathfrak{B}, \bar{b}\bar{d}_0 \dots \bar{d}_n\bar{d}_{n+1} \rangle .$$

If $\bar{a}\bar{c}_0 \dots \bar{c}_{n+1}$ is closed we stop. Otherwise, choose a finite closed set $\bar{c}_{n+2}$ containing $\bar{c}_{n+1}$ and repeat the construction. □

Theorem 4.11 (Ressayre). *All constructible models of a complete theory T are isomorphic and strongly $\aleph_0$-homogeneous.*

Proof. Let $\mathfrak{A}$ and $\mathfrak{B}$ be constructible models of T. First, we show that $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic. Since constructible models are prime, it follows that we can embed $\mathfrak{A}$ into $\mathfrak{B}$ and vice versa. Hence, $\mathfrak{A}$ and $\mathfrak{B}$ have the same cardinality κ. It follows by Lemma 4.10 that $I_{\mathrm{cl}}(\mathfrak{A}, \mathfrak{B}) : \mathfrak{A} \cong_{\mathrm{iso}}^{\kappa \oplus \aleph_1} \mathfrak{B}$. Consequently, Lemma C4.4.10 implies that $\mathfrak{A} \cong \mathfrak{B}$.

It remains to show that $\mathfrak{A}$ is strongly $\aleph_0$-homogeneous. Suppose that

$$\langle \mathfrak{A}, \bar{a} \rangle \equiv \langle \mathfrak{A}, \bar{b} \rangle ,$$

for finite tuples $\bar{a}, \bar{b} \subseteq A$. By Lemma 4.8, these two expansions of $\mathfrak{A}$ are constructible models of the complete theory $T(\bar{a})$. As we have just shown, this implies that they are isomorphic. Hence, there is an automorphism of $\mathfrak{A}$ mapping $\bar{a}$ to $\bar{b}$. □

We apply these tools to show that $\aleph_0$-stable theories have prime models over all sets of parameters.

Lemma 4.12. *Let T be a totally transcendental theory and U a set of parameters. Then the isolated types are dense in $S^{\bar{s}}(U)$.*

Proof. Since $\mathrm{rk}_{\mathrm{CB}}(S^{\bar{s}}(U)) < \infty$ the statement follows from Proposition B5.5.12 (d). □

Proposition 4.13. *Let T be a totally transcendental theory. For every model $\mathfrak{M}$ of T and all parameters $U \subseteq M$, there exists an elementary substructure $\mathfrak{A} \preceq \mathfrak{M}$ such that A is constructible over U. In particular, $\mathfrak{A}$ is a prime model over U and atomic over U.*

Proof. By induction on α, we construct a sequence $(a_\alpha)_{\alpha<\gamma}$ of elements of M as follows. Suppose that we have already defined $(a_i)_{i<\alpha}$. If there is some $b \in M$ such that $\mathrm{tp}(b/U \cup a[<\alpha])$ is isolated then we select one such element and set $a_\alpha := b$. Otherwise, we stop the construction.

Let $A := a[<\gamma]$ be the set of all elements chosen. Clearly, $U \subseteq A$ and $(a_\alpha)_{\alpha<\gamma}$ is a construction of A over U. Hence, it remains to show that $\mathfrak{A} \preceq \mathfrak{M}$ where $\mathfrak{A}$ is the substructure induced by A.

We apply the Tarski-Vaught Test. Suppose that

$$\mathfrak{M} \vDash \varphi(\bar{b}, c), \quad \text{for } \bar{b} \subseteq A \text{ and } c \in M.$$

By Lemma 4.12, there exists an isolated type $\mathfrak{p} \in \langle \varphi(\bar{b}, y) \rangle \subseteq S^1(A)$. Let $d \in M$ be an element realising $\mathfrak{p}$. Since $\mathfrak{p} = \mathrm{tp}(d/A)$ is isolated, it follows by choice of $a[<\gamma]$ that $d \in a[<\gamma] \subseteq A$. Thus, we have found an element $d \in A$ with $\mathfrak{M} \vDash \varphi(\bar{b}, d)$. □

Combining the preceding proposition with Theorem 4.11, we obtain the following result.

Theorem 4.14. *Let T be a totally transcendental theory and let U be a set of parameters. There exists a prime model over U that is also atomic over U. Furthermore, all prime models over U are isomorphic over U.*

Corollary 4.15. *Let T be a totally transcendental theory and let U be a set of parameters. Every model that is prime over U is also atomic over U.*

E4. $\aleph_0$-categorical theories

1. $\aleph_0$-categorical theories and automorphisms

Model theory investigates axiomatisable classes of structures. One of the most basic question one can ask is how many structures of a given cardinality such a class contains.

Definition 1.1. A class $\mathcal{K}$ is *κ-categorical* if, up to isomorphism, it contains exactly one structure of size κ. Similarly, we call a theory T *κ-categorical* if $\mathrm{Mod}(T)$ is κ-categorical.

Example. (a) According to Theorem C4.1.5, the theory of open dense linear orders is $\aleph_0$-categorical.

(b) We have seen in Corollary B6.5.30 that the theory ACF_p of algebraically closed fields of characteristic p is κ-categorical for all uncountable cardinals κ. It has $\aleph_0$ different models of size $\aleph_0$. Hence, it is not $\aleph_0$-categorical.

(c) By Theorem D1.4.8, the same holds for the theory of divisible torsion-free abelian groups.

In this chapter we study $\aleph_0$-categorical theories. We start by showing that, for models of such theories, there is a tight relationship between definable relations and automorphisms. Recall that the automorphism group $\mathrm{Aut}\,\mathfrak{M}$ of a structure $\mathfrak{M}$ is *oligomorphic* if, for every finite tuple $\bar{s}$ of sorts, there are only finitely many orbits of $\mathrm{Aut}\,\mathfrak{M}$ on the set $M^{\bar{s}}$.

Theorem 1.2 (Engeler, Ryll-Nardzewski, Svenonius). *Let T be a countable complete theory with infinite models. The following statements are equivalent:*

(1) *T is $\aleph_0$-categorical.*

(2) $\operatorname{Aut}\mathfrak{M}$ *is oligomorphic, for every countable model $\mathfrak{M}$ of T.*

(3) *T has a countable model $\mathfrak{M}$ such that* $\operatorname{Aut}\mathfrak{M}$ *is oligomorphic.*

(4) *There exists a countable model $\mathfrak{M} \vDash T$ in which, for every finite tuple of sorts $\bar{s}$, only finitely many $\bar{s}$-types (over $\varnothing$) are realised.*

(5) $|S^{\bar{s}}(T)| < \aleph_0$*, for all finite $\bar{s}$.*

(6) *For all finite sets $\bar{x}$ of variables, there are only finitely many formulae $\varphi(\bar{x})$ with free variables $\bar{x}$ that are pairwise non-equivalent modulo T.*

(7) *Every type $\mathfrak{p} \in S^{<\omega}(T)$ is isolated.*

(8) *T has a model that is atomic and $\aleph_0$-saturated.*

(9) *Every model of T is atomic.*

(10) *Every model of T is $\aleph_0$-saturated.*

(11) $\mathfrak{A} \cong^{\aleph_0}_{\mathrm{FO}} \mathfrak{B}$*, for all models $\mathfrak{A}$ and $\mathfrak{B}$ of T.*

(12) $\mathfrak{A} \cong_\infty \mathfrak{B}$*, for all models $\mathfrak{A}$ and $\mathfrak{B}$ of T.*

Proof. (5) ⇒ (6) If $\langle\varphi\rangle = \langle\psi\rangle$ then $\varphi \equiv \psi$ modulo T. If $|S^{\bar{s}}(T)| = k < \aleph_0$ then there are at most 2^k sets of the form $\langle\varphi\rangle$ and, hence, at most that many non-equivalent formulae.

(6) ⇒ (7) For all finite tuples of sorts $\bar{s}$, fix a tuple of variables $\bar{x}$ of sort $\bar{s}$ and a maximal family $\Phi_{\bar{s}}$ of pairwise non-equivalent formulae with free variables $\bar{x}$. For $\mathfrak{p} \in S^{\bar{s}}(T)$, let

$$\psi_{\mathfrak{p}} := \bigwedge \{ \varphi \in \Phi_{\bar{s}} \mid \mathfrak{p} \in \langle\varphi\rangle \} .$$

Then $T \cup \{\psi_{\mathfrak{p}}\} \vDash \mathfrak{p}$ and $\mathfrak{p}$ is isolated.

(7) ⇒ (5) If every type in $S^{\bar{s}}(T)$ is isolated then $S^{\bar{s}}(T)$ is finite, by Lemma B5.5.10.

(7) ⇒ (9) Each model can only realise isolated types since there are no non-isolated ones.

(9) ⇒ (8) Every consistent theory has $\aleph_0$-saturated models.

(8) ⇒ (7) If there is a non-isolated type $\mathfrak{p} \in S^{<\omega}(T)$ then it is realised in all $\aleph_0$-saturated models. Consequently, none of them can be atomic.

(7) ⇒ (10) Suppose that $\mathfrak{M} \vDash T$ is a model, $\bar{a} \in M^m$ a finite tuple, and $\mathfrak{p} \in S^n(\bar{a})$. There is an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ in which $\mathfrak{p}$ is realised by some tuple $\bar{c} \in N^n$. Set $\mathfrak{q} := \operatorname{tp}(\bar{a}\bar{c}/\mathfrak{N})$. Then $\mathfrak{q} \in S^{m+n}(T)$ and, by hypothesis, there is some formula $\varphi(\bar{x}, \bar{y})$ isolating $\mathfrak{q}$. Let $\psi(\bar{x})$ be the formula isolating $\mathfrak{r} := \operatorname{tp}(\bar{a}/\mathfrak{M})$. We claim that

$$T \vDash \psi(\bar{x}) \to \exists \bar{y} \varphi(\bar{x}, \bar{y}) .$$

Then it follows that $\mathfrak{M} \vDash \exists \bar{y} \varphi(\bar{a}, \bar{y})$ and we can find some tuple $\bar{b} \in M^n$ realising $\mathfrak{p}$.

It remains to prove the claim. For a contradiction, suppose it does not hold. Since $\mathfrak{r}$ is complete it follows that $\neg \exists \bar{y} \varphi \in \mathfrak{r}$ and, therefore,

$$T \vDash \psi(\bar{x}) \to \forall \bar{y} \neg \varphi(\bar{x}, \bar{y}) .$$

On the other hand, $\mathfrak{r} \subseteq \mathfrak{q}$ implies that

$$T \vDash \varphi(\bar{x}, \bar{y}) \to \psi(\bar{x}) .$$

Consequently, $T \cup \{\varphi(\bar{x}, \bar{y})\}$ is inconsistent. But this contradicts the fact that $\mathfrak{q} \in S^{m+n}(T)$.

(10) ⇒ (11) follows from Corollary E1.2.3.

(11) ⇒ (12) immediately, since $\mathfrak{A} \cong^{\aleph_0}_{\mathrm{FO}} \mathfrak{B}$ implies $\mathfrak{A} \cong_\infty \mathfrak{B}$.

(12) ⇒ (1) Since T is a countable theory with infinite models it follows that T has a model of cardinality $\aleph_0$. Furthermore, by (12) and Lemma C4.4.10, all such models are isomorphic.

(1) ⇒ (7) Suppose that there exists a type $\mathfrak{p} \in S^{<\omega}(T)$ that is not isolated. T has a model $\mathfrak{A}$ in which $\mathfrak{p}$ is not realised, and it has a model $\mathfrak{B}$ in which $\mathfrak{p}$ is realised. By the Theorem of Löwenheim and Skolem, we can assume that $|A| = |B| = \aleph_0$. Since $\mathfrak{A} \ncong \mathfrak{B}$ T cannot be $\aleph_0$-categorical.

(5) ⇒ (2) Let $\mathfrak{A}$ be a countable model of T and let $\mathfrak{p} \in S^n(T)$. We claim that all tuples realising $\mathfrak{p}$ are in the same orbit of $\operatorname{Aut} \mathfrak{A}$. Hence, the number of orbits is bounded by the number of types which, by (5), is finite.

Suppose that $\bar{a}, \bar{b} \in A^n$ realise $\mathfrak{p}$. We have already seen that (5) implies (11). Hence, we have $\mathfrak{A} \cong_{\mathrm{FO}}^{\aleph_0} \mathfrak{A}$, and $\bar{a} \mapsto \bar{b} \in I_{\mathrm{FO}}^{\aleph_0}(\mathfrak{A}, \mathfrak{A})$ implies that $\langle \mathfrak{A}, \bar{a} \rangle \cong_{\mathrm{FO}}^{\aleph_0} \langle \mathfrak{A}, \bar{b} \rangle$. By Corollary E1.2.3, it follows that there exists an automorphism π with $\pi(\bar{a}) = \bar{b}$.

(2) $\Rightarrow$ (3) is trivial since T is satisfiable.

(3) $\Rightarrow$ (4) We have $\mathrm{tp}(\pi\bar{a}) = \mathrm{tp}(\bar{a})$, for all $\pi \in \mathrm{Aut}\,\mathfrak{M}$. Hence, the number of realised types is bounded by the number of orbits.

(4) $\Rightarrow$ (5) Fix a countable model $\mathfrak{M} \vDash T$ in which only finitely many $\bar{s}$-types are realised, for all finite $\bar{s}$. For a given $\bar{s}$, let $\mathfrak{p}_0, \dots, \mathfrak{p}_{k-1}$ be an enumeration of these $\bar{s}$-types. By Lemma C3.2.6, the set $\{\mathfrak{p}_0, \dots, \mathfrak{p}_{k-1}\}$ is dense in $S^{\bar{s}}(T)$. Consequently, it follows by Lemma B5.5.10 that $S^{\bar{s}}(T)$ is finite. □

Let us also mention a necessary condition for $\aleph_0$-categoricity that deals with the size of the algebraic closure of finite sets.

Lemma 1.3. *Let T be a countable $\aleph_0$-categorical theory with finitely many sorts. There exists a function $s : \omega \to \omega$ such that, for every model $\mathfrak{M}$ of T and every finite set $U \subseteq M$ of parameters, we have*

$$|\mathrm{acl}(U)| \leq s(|U|) .$$

In particular, $\mathrm{acl}(U)$ *is finite for finite sets U.*

Proof. Let $n := |U|$. By Theorem 1.2, $S^{n+1}(T)$ is finite. Let $\mathfrak{p}_0, \dots, \mathfrak{p}_{k-1}$ be an enumeration of $S^{n+1}(T)$ and set

$$I := \Big\{ i < k \;\Big|\; \text{there are } \varphi(x, \bar{y}) \in \mathfrak{p}_i \text{ and } m < \omega \text{ such that } T \vDash \neg\exists^m x \varphi(x, \bar{y}) \Big\} .$$

For $i \in I$, let $m_i < \omega$ be the least number such that

$$\neg\exists^{m_i} x \varphi(x, \bar{y}) \in \mathfrak{p}_i , \quad \text{for some formula } \varphi(x, \bar{y}) .$$

We set $s(n) := \sum_{i \in I} m_i$. Let $a \in \mathrm{acl}(U)$ and let $\bar{b} \in M^n$ be an enumeration of U. The tuple $a\bar{b}$ realises some type $\mathfrak{p}_i$ with $i \in I$. Since there are at most m_i elements c such that $c\bar{b}$ realises $\mathfrak{p}_i$, it follows that

$$|\mathrm{acl}(U)| \leq \sum_{i \in I} m_i = s(n)\,. \qquad \square$$

As an application, we consider fields and groups.

Lemma 1.4. *No infinite field has an $\aleph_0$-categorical theory.*

Proof. Let $\mathfrak{K}$ be an infinite field. By compactness, there exists an elementary extension $\mathfrak{K}_+ \succeq \mathfrak{K}$ that contains a transcendental element c. The algebraic closure $\mathrm{acl}(c)$ is infinite since it contains the elements $c, c^2, c^3, \dots$, which are all distinct. By Lemma 1.3, it follows that $\mathrm{Th}(\mathfrak{K})$ is not $\aleph_0$-categorical. $\square$

Lemma 1.5. *Let $\mathfrak{G}$ be an infinite group.*

(a) *If $\mathrm{Th}(\mathfrak{G})$ is $\aleph_0$-categorical then $\mathfrak{G}$ is locally finite and there exists a number $n < \omega$ such that $g^n = 1$, for all $g \in G$.*

(b) *Conversely, if $\mathfrak{G}$ is abelian and there exists a number n as in* (a), *then $\mathrm{Th}(\mathfrak{G})$ is $\aleph_0$-categorical.*

Proof. (a) Fix an element $g \in G$ and let $s : \omega \to \omega$ be the function from Lemma 1.3. Since $g^n \in \mathrm{acl}(g)$, for all $n < \omega$, and $|\mathrm{acl}(g)| \leq s(1)$, there is some $n < s(1)$ such that $g^{s(1)} = g^n$. Consequently, $g^{s(1)-n} = 1$. Setting $m := s(1)!$ it follows that $g^m = 1$ for all $g \in G$.

(b) Let $\mathfrak{G}$ be a countable abelian group such that $g^n = 1$, for all $g \in G$. There are prime numbers $p_0, \dots, p_{m-1}$, numbers $k_0, \dots, k_{m-1} < \omega$, and cardinals $\lambda_0, \dots, \lambda_{m-1} \leq \aleph_0$ such that

$$\mathfrak{G} \cong \bigoplus_{i<m} (\mathbb{Z}/p_i^{k_i}\mathbb{Z})^{(\lambda_i)}.$$

Set $q_i := p_i^{k_i}$. Note that, for $\lambda_i < \aleph_0$, the group $(\mathbb{Z}/q_i\mathbb{Z})^{(\lambda_i)}$ has

$$r_i := p_i^{\lambda_i k_i} - p_i^{\lambda_i(k_i-1)}$$

elements of order exactly q_i, and, for each element $g \in (\mathbb{Z}/q_i\mathbb{Z})^{(\lambda_i)}$ of order less than q_i, there exists some element h such that $g = h^{p_i}$.

It follows that $\mathfrak{G}$ satisfies the following formula:

- the axioms of an abelian group;
- $\forall x(x^{q_0 \cdots q_{m-1}} = 1)$
- for each $i < m$ such that $\lambda_i < \aleph_0$, the statement that there are exactly r_i elements of order exactly q_i that cannot be written in the form h^{p_i}, for some $h \in G$;
- for each $i < m$ such that $\lambda_i = \aleph_0$, the statement that there are infinitely many elements of order exactly q_i that cannot be written in the form h^{p_i}, for some $h \in G$.

Furthermore, every countable structure $\mathfrak{H}$ satisfying these formulae is isomorphic to $\mathfrak{G}$. Consequently, $\mathrm{Th}(\mathfrak{G})$ is $\aleph_0$-categorical. □

Having characterised the countable theories with exactly one countable model we turn to countable theories with several countable models.

Lemma 1.6. *If T is a countable complete theory with less than $2^{\aleph_0}$ countable models, up to isomorphism, then $|S^{\bar{s}}(T)| \leq \aleph_0$, for all finite $\bar{s}$.*

Proof. Assume that $S^{\bar{s}}(T)$ is uncountable. Then we have $|S^{\bar{s}}(T)| = 2^{\aleph_0}$, by Corollary B5.7.5. Each type $\mathfrak{p} \in S^{\bar{s}}(T)$ is realised in some countable model of T. Since each countable model of T realises only countably many types it follows that T has $2^{\aleph_0}$ models. □

Surprisingly there are no theories with exactly two countable models.

Theorem 1.7. *Let T be a countable complete theory. If T is not $\aleph_0$-categorical then it has at least 3 countable models.*

Proof. If there is a finite tuple $\bar{s}$ of sorts such that $S^{\bar{s}}(T)$ is uncountable then it follows by Lemma 1.6 that T has uncountably many countable models. Hence, we may assume that $S^{\bar{s}}(T)$ is countable, for all $\bar{s}$. By Theorem E3.3.1 and Proposition E1.2.15, it follows that T has a prime model $\mathfrak{A}$ and a countable saturated model $\mathfrak{B}$. If T is not $\aleph_0$-categorical then there is some $\bar{s}$ such that $S^{\bar{s}}(T)$ is infinite and there exists a non-isolated type $\mathfrak{p} \in S^{\bar{s}}(T)$. This type is realised in $\mathfrak{B}$ but not in $\mathfrak{A}$ which implies that $\mathfrak{A} \ncong \mathfrak{B}$.

Let $\bar{a} \in B^{\bar{s}}$ be a tuple of type $\mathfrak{p}$. We know that, for some $k < \omega$, there are infinitely many pairwise non-equivalent formulae with free variables $x_0, \dots, x_{k-1}$. These formulae are still non-equivalent modulo the theory $\mathrm{Th}(\mathfrak{B}_{\bar{a}})$. Hence, $\mathrm{Th}(\mathfrak{B}_{\bar{a}})$ is not $\aleph_0$-categorical and there exists a prime model $\mathfrak{C}$ of this theory. We have $\mathfrak{C} \not\cong \mathfrak{A}$ since $\mathfrak{p}$ is realised in $\mathfrak{C}$. As $\mathfrak{C}$ is not $\aleph_0$-saturated there is a non-isolated type $\mathfrak{q} \in S^{<\omega}(\bar{a})$. Since $\mathfrak{B}$ realises $\mathfrak{q}$ and $\mathfrak{C}$ does not, it follows that $\mathfrak{C} \not\cong \mathfrak{B}$. Thus, we have found three non-isomorphic models $\mathfrak{A}$, $\mathfrak{B}$, $\mathfrak{C}$. □

Lemma 1.8. *There is a countable complete theory T which has exactly three countable models.*

Proof. Let T be the theory of open dense linear orders augmented by the sentences $c_i < c_k$, for all $i < k < \omega$. This theory is complete, it admits quantifier elimination, and the only non-isolated type $\mathfrak{p}$ is the one containing all formulae $x > c_i$, $i < \omega$. There are three models.

(i) The prime model is $\mathfrak{M}_0 \cong \langle \mathbb{Q}, <, (n)_{n<\omega} \rangle$ where the type $\mathfrak{p}$ is not realised since the sequence $(c_i)_i$ is unbounded.

(ii) In $\mathfrak{M}_1 \cong \left\langle \mathbb{Q}, <, \left((1 + \frac{1}{n})^n\right)_{n<\omega} \right\rangle$ the sequence $(c_i)_i$ is bounded but it has no least upper bound.

(iii) In $\mathfrak{M}_2 \cong \left\langle \mathbb{Q}, <, \left(-\frac{1}{n}\right)_{n<\omega} \right\rangle$ the sequence $(c_i)_i$ has a least upper bound. □

Exercise 1.1. For every $3 < n < \omega$, find a countable complete first-order theory with exactly n models.

All possibilities for the number of countable models of a countable theory are listed in the following theorem. Each of them is realised by some theory. The question of whether there are really countable theories with exactly $\aleph_1$ countable models was open for a long time. An affirmative answer was recently given by Knight.

Theorem 1.9 (Morley). *The number of nonisomorphic countably infinite models of a countable complete theory is either a finite number $n \neq 2$, or it is one of $\aleph_0$, $\aleph_1$, or $2^{\aleph_0}$.*

We will not give the complete proof of this result. The next lemma characterises those theories with at most $\aleph_1$ countable models. Morley has shown that all theories that do not satisfy the conditions of the lemma have $2^{\aleph_0}$ countable models.

Lemma 1.10. *Let T be a countable complete theory and let $\mathcal{K}$ be the class of all countable models of T. If we have*

$$|\mathcal{K}/{\equiv_\alpha}| \leq \aleph_0\,, \quad \textit{for every } \alpha < \omega_1\,,$$

then, up to isomorphism, T has at most $\aleph_1$ countable models.

Proof. For $\mathfrak{A} \in \mathcal{K}$, let $\chi(\mathfrak{A}) := \langle \alpha, [\mathfrak{A}]_\alpha \rangle$ where α is the Scott height of $\mathfrak{A}$ and $[\mathfrak{A}]_\alpha \in \mathcal{K}/{\equiv_{\alpha+\omega}}$ is the $\equiv_{\alpha+\omega}$-class of $\mathfrak{A}$. By Corollary C4.4.11, it follows that we have

$$\chi(\mathfrak{A}) = \chi(\mathfrak{B}) \quad \text{iff} \quad \mathfrak{A} \cong \mathfrak{B}\,, \quad \text{for all } \mathfrak{A}, \mathfrak{B} \in \mathcal{K}\,.$$

Consequently, the number of countable models of T is at most

$$|\operatorname{rng} \chi| \leq \aleph_1 \otimes \sup \{\, |\mathcal{K}/{\equiv_\alpha}| \mid \alpha < \omega_1 \,\} \leq \aleph_1 \otimes \aleph_1 = \aleph_1\,. \qquad \square$$

We conclude this section by an investigation of definable relations in countable models of $\aleph_0$-categorical theories.

Lemma 1.11. *Let $\mathfrak{M}$ be a countable model of a countable $\aleph_0$-categorical theory T.*

(a) *Let $\bar{s}$ be a finite tuple of sorts. A relation $R \subseteq M^{\bar{s}}$ is definable in $\mathfrak{M}$ if and only if*

$$\pi[R] = R\,, \quad \textit{for all } \pi \in \operatorname{Aut} \mathfrak{M}\,.$$

(b) *A partial function $f : M_s \to M_t$ is definable in $\mathfrak{M}$ if and only if*

$$\pi \circ f = f \circ \pi\,, \quad \textit{for all } \pi \in \operatorname{Aut} \mathfrak{M}\,.$$

Proof. (a) For the nontrivial direction suppose that $\pi[R] = R$ for all automorphisms π. Since T is $\aleph_0$-categorical there are only finitely many orbits of $\operatorname{Aut}\mathfrak{M}$ on $M^{\bar{s}}$. Hence, R is a finite union of such orbits and it is sufficient to prove that every orbit S is definable.

Fix some tuple $\bar{a} \in S$. We have seen in Theorem 1.2 that $\mathfrak{M}$ is saturated. Hence, it follows by Lemma E1.4.2 and Proposition E1.4.7 that $\mathfrak{M}$ is strongly homogeneous. Consequently, $\operatorname{tp}(\bar{a}) = \operatorname{tp}(\bar{b})$ implies that there is some automorphism π mapping $\bar{a}$ to $\bar{b}$. It follows that

$$S = \{ \bar{b} \in M^{\bar{s}} \mid \operatorname{tp}(\bar{b}) = \operatorname{tp}(\bar{a}) \}.$$

Since every type is isolated there is some formula $\varphi(\bar{x})$ with

$$\mathfrak{M} \models \varphi(\bar{b}) \quad \text{iff} \quad \operatorname{tp}(\bar{b}) = \operatorname{tp}(\bar{a}).$$

It follows that $S = \varphi^{\mathfrak{M}}$.

(b) By (a), a function f is definable if and only if it is invariant under automorphisms, i.e., if and only if

$$b = f(a) \quad \text{iff} \quad \pi(b) = f(\pi(a)), \quad \text{for all } \pi \in \operatorname{Aut}\mathfrak{M}.$$

We can rewrite this condition as $\pi(f(a)) = f(\pi(a))$. □

We can use these results to relate interpretations and automorphism groups.

Definition 1.12. (a) Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures. $\mathfrak{B}$ is *definable* in $\mathfrak{A}$ if it is isomorphic to a structure $\mathfrak{C}$ each domain C_s of which is a definable subset of A such that all relations $R^{\mathfrak{C}}$ and functions $f^{\mathfrak{C}}$ are definable in $\mathfrak{A}$. We call $\mathfrak{A}$ and $\mathfrak{B}$ *bidefinable* if each of them is definable in the other one and the corresponding isomorphisms are inverses of each other.

Definition 1.13. Suppose that $\mathfrak{G}$ and $\mathfrak{H}$ are permutation groups with actions $\alpha : \mathfrak{G} \to \mathfrak{Sym}\,\Omega$ and $\beta : \mathfrak{H} \to \mathfrak{Sym}\,\Delta$, respectively.

(a) A *morphism* $\mathfrak{G} \to \mathfrak{H}$ (or, more precisely, $\alpha \to \beta$) is a pair $\langle h, i \rangle$ where $h : \mathfrak{G} \to \mathfrak{H}$ is a group homomorphism and $i : \Delta \to \Omega$ is a function such that

$$\alpha(g) \circ i = i \circ \beta(h(g)), \quad \text{for all } g \in G.$$

(b) An *embedding* of permutation groups is a morphism $\langle h, i \rangle : \mathfrak{G} \to \mathfrak{H}$ where h and i are both injective.

Theorem 1.14. *Let $\mathfrak{A}$ be a countable model of a countable $\aleph_0$-categorical theory. A structure $\mathfrak{B}$ is definable in $\mathfrak{A}$ if and only if there exists an embedding $\mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathfrak{B}$.*

Proof. The claim follows from Lemma 1.11. If $\mathfrak{B}$ is definable in $\mathfrak{A}$ then every relation $R^{\mathfrak{B}}$ of $\mathfrak{B}$ is closed under $\mathfrak{Aut}\,\mathfrak{A}$. This implies that every automorphism of $\mathfrak{A}$ is also an automorphism of $\mathfrak{B}$. Conversely, each relation $R^{\mathfrak{B}}$ of $\mathfrak{B}$ is closed under all automorphisms of $\mathfrak{B}$. If $\mathfrak{Aut}\,\mathfrak{A} \leq \mathfrak{Aut}\,\mathfrak{B}$ then it also closed under all automorphisms of $\mathfrak{A}$ and, hence, it is definable in $\mathfrak{A}$. ☐

Corollary 1.15. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be countable models of countable $\aleph_0$-categorical theories. Then $\mathfrak{A}$ and $\mathfrak{B}$ are bidefinable if and only if $\mathfrak{Aut}\,\mathfrak{A}$ and $\mathfrak{Aut}\,\mathfrak{B}$ are isomorphic as permutation groups.*

Corollary 1.16. *Let $\mathfrak{A}$ be a countable model of a countable $\aleph_0$-categorical theory. If $\mathfrak{B}$ is a structure with countable signature that is definable in $\mathfrak{A}$ then $\mathrm{Th}(\mathfrak{B})$ is also $\aleph_0$-categorical.*

Proof. If $\mathfrak{Aut}\,\mathfrak{A}$ is oligomorphic and $\mathfrak{Aut}\,\mathfrak{B} \geq \mathfrak{Aut}\,\mathfrak{A}$ then $\mathfrak{Aut}\,\mathfrak{B}$ is also oligomorphic. ☐

A similar characterisation holds for interpretations. Recall that every structure interpretable in $\mathfrak{M}$ can be seen as a definable substructure of $\mathfrak{M}^{\mathrm{eq}}$.

Definition 1.17. Let $\mathcal{I} = \langle \alpha, (\delta_s)_{s \in S}, (\varepsilon_s)_{s \in S}, (\varphi_\xi)_{\xi \in \Gamma} \rangle$ be a first-order interpretation and $\pi : \mathfrak{A} \to \mathfrak{B}$ an isomorphism.

(a) We denote by $\pi^{\mathrm{eq}} : \mathfrak{A}^{\mathrm{eq}} \to \mathfrak{B}^{\mathrm{eq}}$ the unique isomorphism with $\pi^{\mathrm{eq}} \restriction A = \pi$.

(b) Set $\mathfrak{C} := \mathcal{I}(\mathfrak{A})$. For every sort s, the coordinate map of $\mathcal{I}$ induces a bijection $\mathcal{I}_s : D_s \to C_s$ where

$$D_s := \{ [\bar{a}]_{\varepsilon_s} \mid \bar{a} \in \delta_s^{\mathfrak{A}} \} \subseteq A^{\mathrm{eq}}_{\varepsilon_s} .$$

(c) We define

$$\pi^{\mathcal{I}} := \bigcup_s \mathcal{I}_s \circ \pi^{\mathrm{eq}} \circ \mathcal{I}_s^{-1} ,$$

where s ranges over all sorts of $\mathcal{I}(\mathfrak{A})$. We denote the induced map on automorphism groups by $\mathrm{Aut}\,\mathcal{I} : \mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathcal{I}(\mathfrak{A}) : \pi \mapsto \pi^{\mathcal{I}}$.

Lemma 1.18. *Let $\mathcal{I}$ be a first-order interpretation. $\pi^{\mathcal{I}} : \mathcal{I}(\mathfrak{A}) \to \mathcal{I}(\mathfrak{B})$ is an isomorphism, for every isomorphism $\pi : \mathfrak{A} \to \mathfrak{B}$.*

Lemma 1.19. *Every isomorphism $h : \mathfrak{A} \to \mathfrak{B}$ induces an isomorphism $\mathfrak{Aut}\,h : \mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathfrak{B}$ where*

$$(\mathfrak{Aut}\,h)(\pi) := h \circ \pi \circ h^{-1}.$$

Lemma 1.20. *For every first-order interpretation $\mathcal{I}$, the map $\mathfrak{Aut}\,\mathcal{I}$ is a continuous homomorphism*

$$\mathfrak{Aut}\,\mathcal{I} : \mathfrak{Aut}\,\mathfrak{M} \to \mathfrak{Aut}\,\mathcal{I}(\mathfrak{M}) .$$

Proof. It is straightforward to verify that $\mathfrak{Aut}\,\mathcal{I} : \mathfrak{Aut}\,\mathfrak{M} \to \mathfrak{Aut}\,\mathcal{I}(\mathfrak{M})$ is a homomorphism. To see that it is continuous let $S \subseteq \mathrm{Aut}\,\mathcal{I}(\mathfrak{M})$ be a basic open neighbourhood of 1. Then there is some finite tuple $\bar{a}$ in $\mathcal{I}(\mathfrak{M})$ such that

$$S = (\mathrm{Aut}\,\mathcal{I}(\mathfrak{M}))_{(\bar{a})} .$$

Suppose that the sorts of $\bar{a}$ are $\bar{s}$. We fix elements $c_i \in D_{s_i}$ with $\mathcal{I}(c_i) = a_i$. There are finite tuples $\bar{c}_i^* \subseteq M$ such that

$$\mathrm{dcl}^{\mathrm{eq}}(c_i) = \mathrm{dcl}^{\mathrm{eq}}(\bar{c}_i^*) .$$

Setting $S' := (\mathfrak{Aut}\,\mathcal{I})^{-1}[S]$ we have

$$\begin{aligned} \pi \in S' \quad &\text{iff} \quad \mathfrak{Aut}\,\mathcal{I}(\pi)(\bar{a}) = \bar{a} \\ &\text{iff} \quad (\mathcal{I}_{s_i} \circ \pi^{\mathrm{eq}} \circ \mathcal{I}_{s_i}^{-1})(a_i) = a_i\,, \quad \text{for all } i \\ &\text{iff} \quad \pi^{\mathrm{eq}}(c_i) = c_i\,, \quad \text{for all } i \\ &\text{iff} \quad \pi(\bar{c}_i^*) = \bar{c}_i^*\,, \quad \text{for all } i\,. \end{aligned}$$

Consequently, $S' = (\mathrm{Aut}\,\mathfrak{M})_{(\bar{c}_0^* \dots \bar{c}_{m-1}^*)}$ is open. □

Let us call a function $f : M \to M$ *definable* in the structure $\mathfrak{M}$ if each restriction $f \restriction M_s$ is definable, where s ranges over all sorts of $\mathfrak{M}$.

Lemma 1.21. *Let $\varphi : \mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathfrak{B}$ be a continuous homomorphism and suppose that $\mathfrak{A}$ is a countable model of an $\aleph_0$-categorical theory. The following statements are equivalent:*

(1) $\varphi = \mathrm{Aut}\,\pi \circ \mathrm{Aut}\,\mathcal{I}$, *for some interpretation $\mathcal{I}$ and some isomorphism* $\pi : \mathcal{I}(\mathfrak{A}) \to \mathfrak{B}$.

(2) *The subgroup* $\mathrm{rng}\,\varphi \leq \mathfrak{Aut}\,\mathfrak{B}$ *is oligomorphic.*

Proof. (1) ⇒ (2) For every finite tuple $\bar{s}$ of sorts and every orbit S of $\mathrm{rng}\,\varphi$ on $B^{\bar{s}}$, we introduce a new relation R_S of type $\bar{s}$ containing all tuples in the orbit S. Let $\mathfrak{B}^+$ be the expansion of $\mathfrak{B}$ by all these relations R_S. Every automorphism $\sigma \in \mathrm{rng}\,\varphi$ is still an automorphism of the expansion $\mathfrak{B}^+$. Hence, $\mathrm{rng}\,\varphi \leq \mathfrak{Aut}\,\mathfrak{B}^+$. We claim that $\mathrm{rng}\,\varphi$ and $\mathfrak{Aut}\,\mathfrak{B}^+$ have the same orbits.

Since $\mathrm{rng}\,\varphi \leq \mathfrak{Aut}\,\mathfrak{B}^+$ it is sufficient to check that tuples $\bar{a}, \bar{b} \in B^{\bar{s}}$ in different orbits of $\mathrm{rng}\,\varphi$ belong to different orbits of $\mathfrak{Aut}\,\mathfrak{B}^+$. Let S and S' be the orbits under $\mathrm{rng}\,\varphi$ of $\bar{a}$ and $\bar{b}$, respectively. Then $\bar{a} \in R_S$ and $\bar{b} \in R_{S'}$. If $S \neq S'$ then R_S and $R_{S'}$ are disjoint and there is no automorphism of $\mathfrak{B}^+$ mapping $\bar{a}$ to $\bar{b}$.

Consequently, $\mathrm{rng}\,\varphi$ and $\mathfrak{Aut}\,\mathfrak{B}^+$ have the same orbits. To prove (2) it is therefore sufficient to show that $\mathfrak{Aut}\,\mathfrak{B}^+$ is oligomorphic. For a contradiction, suppose that some set $B^{\bar{s}}$ contains tuples $\bar{b}^n$, $n < \omega$, from pairwise distinct orbits. Fix tuples $\bar{a}^n \subseteq A$ such that $(\pi \circ \mathcal{I})(\bar{a}^n) = \bar{b}^n$.

Since $\mathfrak{A}$ is $\aleph_0$-categorical there are indices $k < n$ such that $\bar{a}^k$ and $\bar{a}^n$ belong to the same orbit under $\mathfrak{Aut}\,\mathfrak{A}$. Fix an automorphism $\sigma \in \mathfrak{Aut}\,\mathfrak{A}$ with $\sigma(\bar{a}^k) = \bar{a}^n$. Then

$$\begin{aligned}\varphi(\sigma)(\bar{b}^k) &= (\operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I})(\sigma)(\bar{b}^k) \\ &= (\pi \circ \mathcal{I}_{\bar{s}} \circ \sigma^{\mathrm{eq}} \circ \mathcal{I}_{\bar{s}}^{-1} \circ \pi^{-1})(\bar{b}^k) \\ &= (\pi \circ \mathcal{I}_{\bar{s}} \circ \sigma^{\mathrm{eq}})(\bar{a}^k) \\ &= (\pi \circ \mathcal{I}_{\bar{s}})(\bar{a}^n) = \bar{b}^n.\end{aligned}$$

Hence, the automorphism $\varphi(\sigma)$ maps $\bar{b}^k$ to $\bar{b}^n$. Contradiction.

(2) ⇒ (1) Let $\mathfrak{G} := \mathfrak{Aut}\,\mathfrak{A}$ and $\mathfrak{H} := \mathfrak{Aut}\,\mathfrak{B}$. For each sort s, fix representatives $b^s_0, b^s_1, \dots$ of the orbits of B_s under $\operatorname{rng}\varphi$. The stabiliser $\mathfrak{H}_{(b^s_i)}$ of b^s_i is a basic open neighbourhood of 1 in $\mathfrak{H}$. Since φ is continuous we can find, for each b^s_i, a basic open neighbourhood U^s_i of 1 in $\mathfrak{G}$ with

$$U^s_i \subseteq \varphi^{-1}[\mathfrak{H}_{(b^s_i)}]\,.$$

Every such neighbourhood is of the form $U^s_i = \mathfrak{G}_{(\bar{a}^s_i)}$, for some $\bar{a}^s_i \subseteq A$. Let O^s_i be the orbit of $\bar{a}^s_i$. We define a map $\pi^s_i : O^s_i \to B_s$ by

$$\pi^s_i(\sigma(\bar{a}^s_i)) := \varphi(\sigma)(b^s_i)\,, \quad \text{for } \sigma \in \mathfrak{G}\,.$$

It follows that $\operatorname{rng}\pi^s_i$ is the orbit of b^s_i under $\operatorname{rng}\varphi$. Note that $\ker \pi^s_i$ is invariant under automorphisms since

$$\pi^s_i(\sigma_0(\bar{a}^s_i)) = \pi^s_i(\sigma_1(\bar{a}^s_i))$$

implies

$$\begin{aligned}\pi^s_i((\rho \circ \sigma_0)(\bar{a}^s_i)) &= \varphi(\rho \circ \sigma_0)(b^s_i) \\ &= \varphi(\rho \circ \sigma_1)(b^s_i) = \pi^s_i((\rho \circ \sigma_1)(\bar{a}^s_i))\,.\end{aligned}$$

By Lemma 1.11 it follows that $\ker \pi^s_i$ is definable. We obtain a definable subset $U^s_i := O^s_i / \ker \pi^s_i \subseteq A^{\mathrm{eq}}$ and an injective function

$$\pi^s : \bigcup_i U^s_i \to B_s\,.$$

This map is also surjective since its range contains every orbit of B_s under $\operatorname{rng}\varphi$. Setting $\pi := \bigcup_s \pi^s$ we obtain a bijection $\pi : \bigcup_s U_s \to B$. We claim that this bijection is an isomorphism between $\mathfrak{B}$ and a structure of the form $\mathcal{I}(\mathfrak{A})$, for a suitable interpretation $\mathcal{I}$.

If R is a definable relation in $\mathfrak{B}$ then its preimage $\pi^{-1}[R]$ is invariant under automorphisms. Hence, $\pi^{-1}[R]$ is definable in $\mathfrak{A}^{\mathrm{eq}}$. It follows that there exists an interpretation $\mathcal{I}$ such that $\pi : \mathcal{I}(\mathfrak{A}) \cong \mathfrak{B}$.

It remains to check that $\varphi = \operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I}$. For every $\sigma \in \mathfrak{G}$ we have

$$\begin{aligned}(\operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I})(\sigma)(b^s_i) &= (\pi \circ \mathcal{I}_s \circ \sigma^{\mathrm{eq}} \circ \mathcal{I}_s^{-1} \circ \pi^{-1})(b^s_i) \\ &= (\pi \circ \mathcal{I}_s \circ \sigma^{\mathrm{eq}})(\bar{a}^s_i) \\ &= (\pi \circ \mathcal{I}_s)(\sigma(\bar{a}^s_i)) = \varphi(\sigma)(b^s_i) \qquad \square\end{aligned}$$

Corollary 1.22. *Let Σ and Γ be countable signatures and $\mathcal{I}$ a first-order interpretation from Σ to Γ. If $\mathfrak{A}$ is a countable Σ-structure with $\aleph_0$-categorical theory then the theory of $\mathcal{I}(\mathfrak{A})$ is also $\aleph_0$-categorical.*

Proof. $\operatorname{Aut}\mathcal{I} : \mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathcal{I}(\mathfrak{A})$ is a continuous homomorphism. By the preceding lemma it follows that $\operatorname{rng}(\operatorname{Aut}\mathcal{I})$ is oligomorphic. Since $\operatorname{rng}(\operatorname{Aut}\mathcal{I}) \leq \mathfrak{Aut}\,\mathcal{I}(\mathfrak{A})$ it follows that $\mathfrak{Aut}\,\mathcal{I}(\mathfrak{A})$ is also oligomorphic. $\square$

Definition 1.23. Let $\mathfrak{M}$ be a structure and suppose that $\mathcal{I}$ and $\mathcal{J}$ are interpretations such that there exists an isomorphism $\pi : \mathcal{I}(\mathfrak{M}) \cong \mathcal{J}(\mathfrak{M})$. We call $\mathcal{I}$ and $\mathcal{J}$ *homotopic* (via π) if there exists a definable function $\rho : M \to M$ such that $\pi \circ \mathcal{I} = \mathcal{J} \circ \rho$.

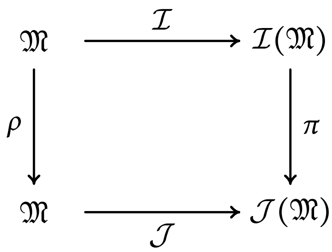

Lemma 1.24. *Let $\mathfrak{M}$ be a countable structure with $\aleph_0$-categorical theory and suppose that $\mathcal{I}$ and $\mathcal{J}$ are interpretations with $\mathcal{I}(\mathfrak{M}) \cong \mathcal{J}(\mathfrak{M})$. Let*

$\pi : \mathcal{I}(\mathfrak{M}) \to \mathcal{J}(\mathfrak{M})$ be an isomorphism. Then $\mathcal{I}$ and $\mathcal{J}$ are homotopic via π if and only if $\operatorname{Aut}\mathcal{J} = \operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I}$.

Proof. ($\Rightarrow$) Let $\rho : M \to M$ be a definable function such that $\pi \circ \mathcal{I} = \mathcal{J} \circ \rho$. For every element b of $\mathcal{J}(\mathfrak{M})$ and every automorphism $\sigma \in \operatorname{Aut}\mathfrak{M}$, we have

$$\begin{aligned}(\operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I})(\sigma)(b) &= (\pi \circ \mathcal{I}_s \circ \sigma^{\mathrm{eq}} \circ \mathcal{I}_s^{-1} \circ \pi^{-1})(b) \\ &= (\mathcal{J}_s \circ \rho \circ \sigma^{\mathrm{eq}} \circ \rho^{-1} \circ \mathcal{J}_s^{-1})(b) \\ &= (\mathcal{J}_s \circ \sigma^{\mathrm{eq}} \circ \rho \circ \rho^{-1} \circ \mathcal{J}_s^{-1})(b) \\ &= (\operatorname{Aut}\mathcal{J})(\sigma)(b)\end{aligned}$$

Hence, $\operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I} = \operatorname{Aut}\mathcal{J}$.

($\Leftarrow$) For $a \in M$, we define

$$\rho(a) := (\mathcal{J}_s^{-1} \circ \pi \circ \mathcal{I}_s)(a)\,.$$

We claim that ρ is definable. For $\sigma \in \operatorname{Aut}\mathfrak{M}$ and $a \in M$, we have

$$\begin{aligned}\rho(\sigma(a)) &= (\mathcal{J}_s^{-1} \circ \pi \circ \mathcal{I}_s \circ \sigma)(a) \\ &= (\mathcal{J}_s^{-1} \circ \pi \circ \mathcal{I}_s \circ \sigma \circ \mathcal{I}_s^{-1} \circ \pi^{-1} \circ \pi \circ \mathcal{I}_s)(a) \\ &= (\mathcal{J}_s^{-1} \circ (\operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I})(\sigma) \circ \pi \circ \mathcal{I}_s)(a) \\ &= (\mathcal{J}_s^{-1} \circ (\operatorname{Aut}\mathcal{J})(\sigma) \circ \pi \circ \mathcal{I}_s)(a) \\ &= (\sigma \circ \mathcal{J}_s^{-1} \circ \pi \circ \mathcal{I}_s)(a) \\ &= \sigma(\rho(a))\,.\end{aligned}$$

Hence, ρ is invariant under automorphisms and, thus, definable. □

Definition 1.25. Two structures $\mathfrak{A}$ and $\mathfrak{B}$ are *biinterpretable* if there exist first-order interpretations $\mathcal{I}, \mathcal{J}$ and isomorphisms $\pi : \mathcal{I}(\mathfrak{A}) \to \mathfrak{B}$ and $\rho : \mathcal{J}(\mathfrak{B}) \to \mathfrak{A}$ such that $\mathcal{J} \circ \mathcal{I}$ is homotopic to $\mathrm{id}_{\mathfrak{A}}$ via $\rho \circ \pi^{\mathcal{J}}$ and $\mathcal{I} \circ \mathcal{J}$ is homotopic to $\mathrm{id}_{\mathfrak{B}}$ via $\pi \circ \rho^{\mathcal{I}}$.

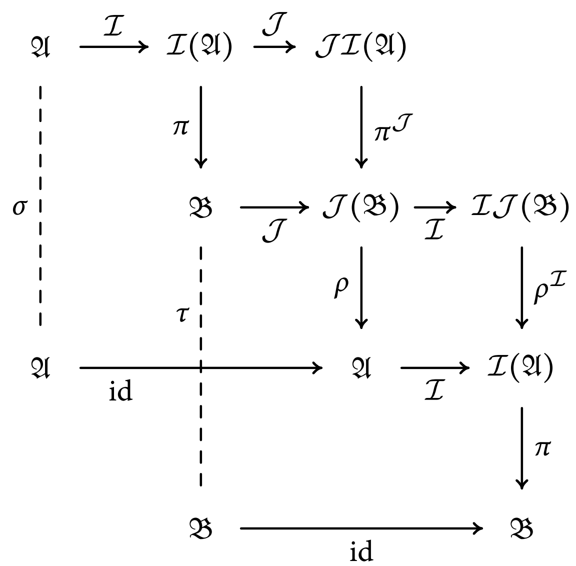

Theorem 1.26. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be countable models of countable $\aleph_0$-categorical theories. Then $\mathfrak{A}$ and $\mathfrak{B}$ are biinterpretable if and only if $\mathfrak{Aut}\,\mathfrak{A}$ and $\mathfrak{Aut}\,\mathfrak{B}$ are isomorphic as topological groups.*

Proof. ($\Rightarrow$) Let $\mathcal{I}, \mathcal{J}$ and π, ρ witness that $\mathfrak{A}$ and $\mathfrak{B}$ are biinterpretable. There exist definable maps $\sigma : A \to A$ and $\tau : B \to B$ such that

$$\rho \circ \pi^{\mathcal{J}} \circ \mathcal{J} \circ \mathcal{I} = \sigma \quad \text{and} \quad \pi \circ \rho^{\mathcal{I}} \circ \mathcal{I} \circ \mathcal{J} = \tau\,.$$

Set $\varphi := \operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I}$ and $\psi := \operatorname{Aut}\rho \circ \operatorname{Aut}\mathcal{J}$. Since σ and τ are definable we have

$$\operatorname{Aut}\sigma = \mathrm{id} \quad \text{and} \quad \operatorname{Aut}\tau = \mathrm{id}\,.$$

It follows that

$$\begin{aligned}
\varphi \circ \psi &= \operatorname{Aut}\rho \circ \operatorname{Aut}\mathcal{J} \circ \operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I} \\
&= \operatorname{Aut}(\rho \circ \mathcal{J} \circ \pi \circ \mathcal{I}) \\
&= \operatorname{Aut}(\rho \circ \pi^{\mathcal{J}} \circ \mathcal{J} \circ \mathcal{I}) \\
&= \operatorname{Aut}\sigma \\
&= \mathrm{id}\,,
\end{aligned}$$

and, analogously,

$$\psi \circ \varphi = \mathrm{id}\,.$$

Hence, $\psi = \varphi^{-1}$ and $\varphi : \mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathfrak{B}$ is the desired isomorphism.

($\Leftarrow$) Let $\varphi : \mathfrak{Aut}\,\mathfrak{A} \to \mathfrak{Aut}\,\mathfrak{B}$ be an isomorphism. Since $\operatorname{rng}\varphi = \mathfrak{Aut}\,\mathfrak{B}$ is oligomorphic it follows by Lemma 1.21 that $\varphi = \operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I}$, for some interpretation $\mathcal{I}$ and some isomorphism $\pi : \mathcal{I}(\mathfrak{A}) \to \mathfrak{B}$. Similarly, $\operatorname{rng}\varphi^{-1}$ is oligomorphic and we have $\varphi^{-1} = \operatorname{Aut}\rho \circ \operatorname{Aut}\mathcal{J}$, for some $\mathcal{J}$ and ρ. It follows that

$$\begin{aligned}\operatorname{Aut}(\rho \circ \mathcal{J} \circ \pi \circ \mathcal{I}) &= \operatorname{Aut}\rho \circ \operatorname{Aut}\mathcal{J} \circ \operatorname{Aut}\pi \circ \operatorname{Aut}\mathcal{I} \\ &= \varphi^{-1} \circ \varphi = \mathrm{id}\,.\end{aligned}$$

By Lemma 1.24, there exists a definable map $\sigma : A \to A$ such that

$$\pi \circ \mathcal{I} \circ \rho \circ \mathcal{J} = \sigma\,.$$

Analogously, we obtain a definable map $\tau : B \to B$ such that

$$\rho \circ \mathcal{J} \circ \pi \circ \mathcal{I} = \tau\,.$$

Hence, $\mathcal{J} \circ \mathcal{I}$ and id are homotopic via $\rho \circ \pi^{\mathcal{J}}$ and $\mathcal{I} \circ \mathcal{J}$ and id are homotopic via $\pi \circ \rho^{\mathcal{I}}$. □

2. *Back-and-forth arguments in accessible categories*

In the next section, we will prove a result about accessible categories using back-and-forth arguments. The necessary machinery for such arguments is developed in the present section. We start by generalising the notion of a partial isomorphism and the forth-property.

Definition 2.1. Let $\mathcal{C}$ be a category, $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ a class of objects, and $\mathfrak{a}, \mathfrak{b} \in \mathcal{C}$.

(a) A *partial morphism* from $\mathfrak{a}$ to $\mathfrak{b}$ is a pair $p = \langle f, f' \rangle$ of morphisms $f : \mathfrak{c} \to \mathfrak{a}$ and $f' : \mathfrak{c} \to \mathfrak{b}$, for some object $\mathfrak{c} \in \mathcal{C}$. We call $\mathfrak{a}$ the *domain* of p, $\mathfrak{b}$ its *codomain,* and $\mathfrak{c}$ is its *base.*

(b) Let $p = \langle f, f' \rangle$ and $q = \langle g, g' \rangle$ be partial morphisms with bases $\mathfrak{c}$ and $\mathfrak{d}$, respectively. A *morphism* $p \to q$ is a morphism $h : \mathfrak{c} \to \mathfrak{d}$ such that

$$f = g \circ h \quad \text{and} \quad f' = g' \circ h\,.$$

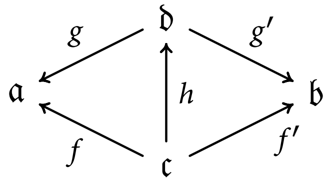

(c) We denote by $\mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ the category of all partial morphisms p from $\mathfrak{a}$ to $\mathfrak{b}$ whose base belongs to $\mathcal{K}$. If $\mathcal{K}$ is the class of all κ-presentable objects, we will write $\mathfrak{pMor}_{\kappa}(\mathfrak{a}, \mathfrak{b})$ instead.

(d) The *domain projection* is the functor

$$P : \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \to \mathfrak{Sub}_{\mathcal{K}}(\mathfrak{a})$$

that maps a partial morphism $p = \langle f, f' \rangle$ to its first component f and a morphism $h : \langle f, f' \rangle \to \langle g, g' \rangle$ of $\mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ to the underlying morphism $h : f \to g$ of $\mathfrak{Sub}_{\mathcal{K}}(\mathfrak{a})$.

Analogously, the *codomain projection* is the functor

$$Q : \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \to \mathfrak{Sub}_{\mathcal{K}}(\mathfrak{b})$$

mapping $\langle f, f' \rangle$ to f' and $h : \langle f, f' \rangle \to \langle g, g' \rangle$ to $h : f' \to g'$.

Finally, the *base projection* is the functor

$$B : \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \to \mathcal{C}$$

mapping a partial morphism p to its base and a morphism $h : p \to q$ to the corresponding morphism $h : B(p) \to B(q)$ between the bases.

Definition 2.2. Let $\mathcal{C}$ be a category, $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ a class of objects, and $\mathfrak{a}, \mathfrak{b} \in \mathcal{C}$.

(a) A set I of partial morphisms from $\mathfrak{a}$ to $\mathfrak{b}$ has the *forth property* with respect to $\mathcal{K}$ if, for every $p = \langle f, f' \rangle \in I$ with base $\mathfrak{c}$, every $\mathfrak{d} \in \mathcal{K}$, and every pair of morphisms $g : \mathfrak{d} \to \mathfrak{a}$ and $h : \mathfrak{c} \to \mathfrak{d}$ with $f = g \circ h$, there exists a morphism $g' : \mathfrak{d} \to \mathfrak{b}$ such that $\langle g, g' \rangle \in I$ and

$$h : \langle f, f' \rangle \to \langle g, g' \rangle .$$

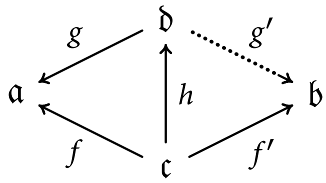

(b) We write

$$\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b} \quad :\text{iff} \quad \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \text{ is nonempty and it has the forth property with respect to } \mathcal{K}.$$

Furthermore, we write

$$\mathfrak{a} \sqsubseteq^{\kappa}_{\text{pres}} \mathfrak{b} \quad :\text{iff} \quad \mathfrak{a} \sqsubseteq_{\mathcal{K}_\kappa} \mathfrak{b} ,$$

where $\mathcal{K}_\kappa \subseteq \mathcal{C}$ is the class of all κ-presentable objects. The corresponding equivalence relations are

$$\mathfrak{a} \equiv_{\mathcal{K}} \mathfrak{b} \quad :\text{iff} \quad \mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b} \quad \text{and} \quad \mathfrak{b} \sqsubseteq_{\mathcal{K}} \mathfrak{a} ,$$
$$\mathfrak{a} \equiv^{\kappa}_{\text{pres}} \mathfrak{b} \quad :\text{iff} \quad \mathfrak{a} \sqsubseteq^{\kappa}_{\text{pres}} \mathfrak{b} \quad \text{and} \quad \mathfrak{b} \sqsubseteq^{\kappa}_{\text{pres}} \mathfrak{a} .$$

Remark. In the category $\mathfrak{Emb}(\Sigma)$ we have

$$\mathfrak{A} \sqsubseteq^{\kappa}_{\text{pres}} \mathfrak{B} \quad \text{iff} \quad \mathfrak{A} \sqsubseteq^{\kappa}_{0} \mathfrak{B} .$$

Note that, for an arbitrary category, the relation $\sqsubseteq_{\mathcal{K}}$ is not very well-behaved. For instance, in general it is neither reflexive nor transitive. The next lemma collects some basic properties that hold in every category.

Lemma 2.3. *Let $\mathcal{C}$ be a category and $\mathcal{K} \subseteq \mathcal{C}^{\text{obj}}$.*

(a) *If there exists a morphism $\varphi : \mathfrak{a}_0 \to \mathfrak{a}$, then*

$$\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b} \quad \textit{implies} \quad \mathfrak{a}_0 \sqsubseteq_{\mathcal{K}} \mathfrak{b} .$$

(b) *If $\mathfrak{a} \in \mathcal{K}$ and $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$, then there exists a morphism $\mathfrak{a} \to \mathfrak{b}$.*

(c) *If $\mathfrak{a}, \mathfrak{b} \in \mathcal{K}$ and $\mathfrak{a} \equiv_{\mathcal{K}} \mathfrak{b}$, then $\mathfrak{a} \cong \mathfrak{b}$.*

Proof. (a) Let $\langle f, f' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}_0, \mathfrak{b})$ be a partial morphism with base $\mathfrak{c}$ and let $h : \mathfrak{c} \to \mathfrak{d}$ and $g : \mathfrak{d} \to \mathfrak{a}_0$ be morphisms with $f = g \circ h$ and $\mathfrak{d} \in \mathcal{K}$. Then $\langle \varphi \circ f, f' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ and $h : \mathfrak{c} \to \mathfrak{d}$ and $\varphi \circ g : \mathfrak{d} \to \mathfrak{a}$ are morphisms such that $\varphi \circ f = \varphi \circ g \circ h$ and $\mathfrak{d} \in \mathcal{K}$. Consequently, $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$ implies that there exists a morphism $g' : \mathfrak{d} \to \mathfrak{b}$ such that

$$\langle \varphi \circ g, g' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \quad \text{and} \quad h : \langle \varphi \circ f, f' \rangle \to \langle \varphi \circ g, g' \rangle .$$

It follows that $\langle g, g' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}_0, \mathfrak{b})$ and $h : \langle f, f' \rangle \to \langle g, g' \rangle$.

(b) As $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$, there exists a partial morphism $\langle f, f' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$. Since $\mathfrak{a} \in \mathcal{K}$, we can use the forth-property to find a morphism $g : \mathfrak{a} \to \mathfrak{b}$ such that

$$\langle \mathrm{id}_{\mathfrak{a}}, g \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$$
$$\text{and} \quad f : \langle f, f' \rangle \to \langle \mathrm{id}_{\mathfrak{a}}, g \rangle .$$

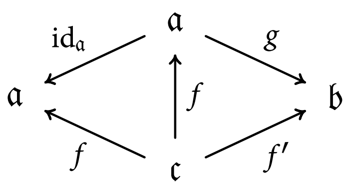

(c) As $a \equiv_{\mathcal{K}} \mathfrak{b}$, there exists a partial morphism $\langle f, f' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$. As in (b), we can use the forth-property to find a morphism $g : \mathfrak{a} \to \mathfrak{b}$ such that

$$\langle \mathrm{id}_{\mathfrak{a}}, g \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \quad \text{and} \quad f : \langle f, f' \rangle \to \langle \mathrm{id}_{\mathfrak{a}}, g \rangle .$$

Similarly, we can use the back-property to find a morphism $h : \mathfrak{b} \to \mathfrak{a}$ such that

$$\langle h, \mathrm{id}_{\mathfrak{b}} \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$$
and $\quad g : \langle \mathrm{id}_{\mathfrak{a}}, g \rangle \to \langle h, \mathrm{id}_{\mathfrak{b}} \rangle\,.$

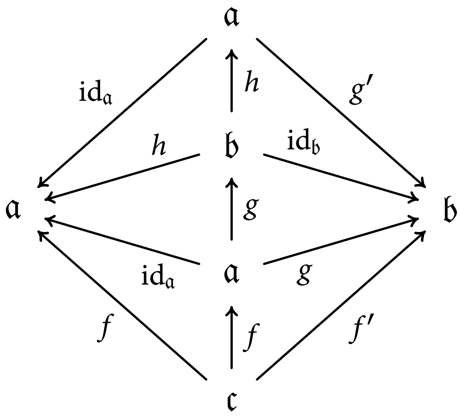

Using the forth-property again, we obtain a morphism $g' : \mathfrak{a} \to \mathfrak{b}$ such that

$$\langle \mathrm{id}_{\mathfrak{a}}, g' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$$
and $\quad h : \langle h, \mathrm{id}_{\mathfrak{b}} \rangle \to \langle \mathrm{id}_{\mathfrak{a}}, g' \rangle\,.$

In particular, $h \circ g = \mathrm{id}_{\mathfrak{a}}$ and $g' \circ h = \mathrm{id}_{\mathfrak{b}}$. By Lemma B1.3.4, it follows that $g = g'$ and $h : \mathfrak{b} \cong \mathfrak{a}$ is an isomorphism. □

Our goal is to generalise Lemma C4.4.10 to relations of the form $\sqsubseteq_{\mathcal{K}}$. We start with the forth-property.

Proposition 2.4. *Let κ be an infinite cardinal or $\kappa = \infty$, $\mathcal{C}$ a category with colimits of nonempty chains of length less than κ, and let $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ be a class of objects that is closed under colimits of nonempty chains of length less than κ. Let $D : \gamma \to \mathcal{K}$ be a chain of length $0 < \gamma \leq \kappa$ with limiting cocone $\mu \in \mathrm{Cone}(D, \mathfrak{a})$. Suppose that every morphism from some object in $\mathcal{K}$ to $\mathfrak{a}$ factorises essentially uniquely through μ.*

If $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$, then there exists a chain $E : \gamma \to \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ such that $D = B \circ E$, where B is the base projection functor.

Proof. By induction on $\alpha < \gamma$, we define morphisms $\nu_\alpha : D(\alpha) \to \mathfrak{b}$ such that

$$\langle \mu_\alpha, \nu_\alpha \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$$
and $\quad D(\alpha, \beta) : \langle \mu_\alpha, \nu_\alpha \rangle \to \langle \mu_\beta, \nu_\beta \rangle\,,$

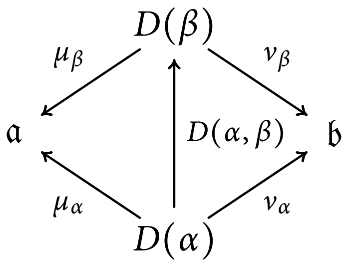

for $\alpha \leq \beta < \gamma$.

Then we can set

$$E(\alpha) := \langle \mu_\alpha, \nu_\alpha \rangle \quad \text{and} \quad E(\alpha, \beta) := D(\alpha, \beta), \quad \text{for } \alpha \leq \beta < \gamma .$$

For $\alpha = 0$, we define ν_α as follows. Since $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$, there exists a partial morphism $\langle f, f' \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$. Let $\mathfrak{c}$ be its base. By assumption on D, f factorises as $f = \mu_\alpha \circ f_0$, for some index $\alpha < \gamma$ and some morphism $f_0 : \mathfrak{c} \to D(\alpha)$. As $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$, there exists a morphism $\nu_\alpha : D(\alpha) \to \mathfrak{b}$ such that $\langle \mu_\alpha, \nu_\alpha \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ and

$$f_0 : \langle f, f' \rangle \to \langle \mu_\alpha, \nu_\alpha \rangle .$$

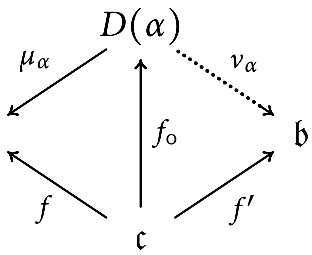

Setting $\nu_0 := \nu_\alpha \circ D(0, \alpha)$ we obtain the desired morphism $D(0) \to \mathfrak{b}$.

For the inductive step, suppose that we have already defined ν_α for all $\alpha < \beta$. Let λ^β be a limiting cocone from $D \upharpoonright \beta$ to some object $\mathfrak{d}_\beta$. As $\mathcal{K}$ is closed under colimits of chains of length β, we have $\mathfrak{d}_\beta \in \mathcal{K}$. Since $(\mu_\alpha)_{\alpha<\beta}$ and $(\nu_\alpha)_{\alpha<\beta}$ are cocones of $D \upharpoonright \beta$, there exist unique morphisms $\varphi : \mathfrak{d}_\beta \to \mathfrak{a}$ and $\varphi' : \mathfrak{d}_\beta \to \mathfrak{b}$ such that

$$(\mu_\alpha)_{\alpha<\beta} = \varphi * \lambda^\beta \quad \text{and} \quad (\nu_\alpha)_{\alpha<\beta} = \varphi' * \lambda^\beta .$$

Similarly, $(D(\alpha, \beta))_{\alpha<\beta}$ is a cocone from $D \upharpoonright \beta$ to $D(\beta)$ and there exists a unique morphism $\psi : \mathfrak{d}_\beta \to D(\beta)$ such that

$$(D(\alpha, \beta))_{\alpha<\beta} = \psi * \lambda^\beta .$$

Since

$$\mu_\beta \circ \psi \circ \lambda^\beta_\alpha = \mu_\beta \circ D(\alpha, \beta) = \mu_\alpha = \varphi \circ \lambda^\beta_\alpha , \quad \text{for all } \alpha < \beta ,$$

it follows by Lemma B3.4.2 that

$$\mu_\beta \circ \psi = \varphi \,.$$

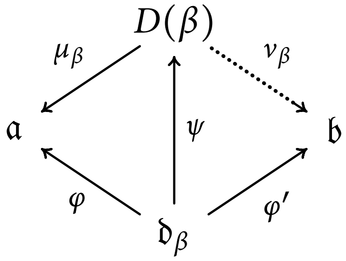

Therefore, $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$ implies that there exists a morphism $\nu_\beta : D(\beta) \to \mathfrak{b}$ such that

$$\langle \mu_\beta, \nu_\beta \rangle \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b}) \quad \text{and} \quad \psi : \langle \varphi, \varphi' \rangle \to \langle \mu_\beta, \nu_\beta \rangle \,.$$

For $\alpha < \beta$ it follows that $D(\alpha, \beta) = \psi \circ \lambda^\beta_\alpha$ is a morphism

$$D(\alpha, \beta) : \langle \mu_\alpha, \nu_\alpha \rangle \to \langle \mu_\beta, \nu_\beta \rangle \,. \qquad \square$$

Proposition 2.5. *Let κ be an infinite cardinal or $\kappa = \infty$, $\mathcal{C}$ a category with colimits of nonempty chains of length at most κ, and let $\mathcal{K} \subseteq \mathcal{C}^{\text{obj}}$ be a class of objects that is closed under colimits of nonempty chains of length less than κ. Let $D : \gamma^D \to \mathcal{K}$ and $E : \gamma^E \to \mathcal{K}$ be chains of length $0 < \gamma^D, \gamma^E \leq \kappa$ with limiting cocones $\lambda^D \in$ Cone$(D, \mathfrak{a})$ and $\lambda^E \in$ Cone$(E, \mathfrak{b})$. Suppose that every morphism from some object in $\mathcal{K}$ to $\mathfrak{a}$ or $\mathfrak{b}$ factorises essentially uniquely through, respectively, λ^D and λ^E.*

If $\mathfrak{a} \equiv_{\mathcal{K}} \mathfrak{b}$ and $p \in \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$, there exists a morphism $\varphi : p \to q$ of $\mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ such that $q = \langle g, g' \rangle$ consists of two epimorphisms.

Proof. By induction on the ordinals γ^D and γ^E, we construct a chain $F : \delta \to \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$, two links s and t from $B \circ F$ to D and E, respectively, and two increasing functions $\rho_0 : \gamma^D \to \delta$ and $\theta_0 : \gamma^E \to \delta$ such that

$$\begin{aligned} B(F(\alpha)) &= D(\rho(\alpha))\,, & s_\alpha &= \mathrm{id}_{D(\rho(\alpha))}\,, & &\text{for } \alpha \in \operatorname{rng} \rho_0\,, \\ B(F(\alpha)) &= E(\theta(\alpha))\,, & t_\alpha &= \mathrm{id}_{E(\theta(\alpha))}\,, & &\text{for } \alpha \in \operatorname{rng} \theta_0\,, \end{aligned}$$

where B is the base projection functor and ρ and θ are the index maps of s and t, respectively.

For $\gamma^D, \gamma^E = 0$, we start with $\delta := 1$ and $F(0) := p$. To define s and t, suppose that $p = \langle f, f' \rangle$. By assumption, f and f' factorise essentially uniquely through λ^D and λ^E, respectively. Let $f = \lambda^D_\alpha \circ f_0$ and $f' = \lambda^E_\beta \circ f'_0$ be the corresponding factorisations. We set $s_0 := f_0$ and $t_0 := f'_0$.

For the inductive step, suppose that, for the restrictions $D \upharpoonright \beta^D$ and $E \upharpoonright \beta^E$, we have already defined a chain $F : \delta \to \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ with $0 < \delta < \kappa$, links s and t from $B \circ F$ to $D \upharpoonright \beta^D$ and $E \upharpoonright \beta^E$, respectively, and increasing functions $\rho_0 : \beta^D \to \delta$ and $\theta_0 : \beta^E \to \delta$.

We will show how to extend these definitions to $D \upharpoonright \beta^D + 1$. (The extension to $E \upharpoonright \beta^E + 1$ works in the same way.) Let μ be a limiting cocone from $B \circ F$ to some object $\mathfrak{c}$. As $\mathcal{K}$ is closed under limits of chains of length $0 < \delta < \kappa$, it follows that $\mathfrak{c} \in \mathcal{K}$. Since $\lambda^D * s$ is a cocone of $B \circ F$, there exists a unique morphism $\varphi^D : \mathfrak{c} \to \mathfrak{a}$ such that $\lambda^D * s = \varphi^D * \mu$. In the same way, we obtain a unique morphism $\varphi^E : \mathfrak{c} \to \mathfrak{b}$ with $\lambda^E * t = \varphi^E * \mu$.

As $\mathfrak{c} \in \mathcal{K}$, there exists an essentially unique factorisation $\varphi^D = \lambda^D_\alpha \circ \varphi_0$, for some morphism $\varphi_0 : \mathfrak{c} \to D(\alpha)$ with $\alpha \geq \beta^D$. Since $\mathfrak{a} \sqsubseteq_{\mathcal{K}} \mathfrak{b}$, we can find a morphism $\psi : \mathfrak{c} \to \mathfrak{b}$ such that

$$\psi \circ \varphi_0 = \varphi^E .$$

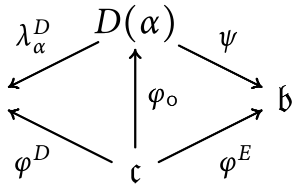

As $D(\alpha) \in \mathcal{K}$, there exists an essentially unique factorisation $\psi = \lambda^E_\beta \circ \psi_0$, for some morphism $\psi_0 : D(\alpha) \to E(\beta)$ with $\beta \geq \beta^D$. We set

$$F(\delta) := \langle \lambda^D_\alpha, \psi \rangle , \qquad F(i, \delta) := \varphi_0 \circ \mu_i , \quad \text{for } i < \delta ,$$
$$s_\delta := \mathrm{id}_{D(\alpha)} , \qquad \rho_0(\beta^D) := \alpha ,$$
$$t_\delta := \psi_0 .$$

Let us show that these morphisms have the desired properties. First, we check that the extension of s is a link from the extension of $B \circ F$ to D. For every $i < \delta$, it follows by choice of φ^D that

$$\lambda^D_\alpha \circ D(\rho(i), \alpha) \circ s_i = \lambda^D_{\rho(i)} \circ s_i = \varphi^D \circ \mu_i = \lambda^D_\alpha \circ \varphi_0 \circ \mu_i .$$

Since $B(F(i)) \in \mathcal{K}$, this morphism has an essentially unique factorisation through λ^D. Hence, the above two factorisations are a.p.-equivalent.

$$D(\varphi(i), \alpha) \circ s_i \curlywedge_D \varphi_0 \circ \mu_i \,.$$

By Lemma B3.5.3 (d), this implies that

$$s_i \curlywedge_D \varphi_0 \circ \mu_i = s_\delta \circ F(i, \delta)\,,$$

as desired.

We also have to check that the extension of t is a link. Let $i < \delta$. Then

$$\begin{aligned}\lambda^E_\beta \circ t_\delta \circ F(i, \delta) &= \lambda^E_\beta \circ \psi_0 \circ \varphi_0 \circ \mu_i \\ &= \psi \circ \varphi_0 \circ \mu_i = \varphi^E \circ \mu_i = \lambda^E_{\theta(i)} \circ t_i \,.\end{aligned}$$

Since $B(F(i)) \in \mathcal{K}$, this morphism has an essentially unique factorisation through λ^E. Hence, the above two factorisations are a.p.-equivalent.

$$t_\delta \circ F(i, \delta) \curlywedge_E t_i\,,$$

as desired.

Having defined $F : \delta \to \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$, we construct the desired partial morphism $q = \langle g, g' \rangle \in \mathfrak{pMor}(\mathfrak{a}, \mathfrak{b})$ as follows. Let λ^F be a limiting cocone from $B \circ F$ to some object $\mathfrak{c} \in \mathcal{C}$. Since $\lambda^D * s$ and $\lambda^E * t$ are cocones of F, there exist unique morphisms $g : \mathfrak{c} \to \mathfrak{a}$ and $g' : \mathfrak{c} \to \mathfrak{b}$ such that $\lambda^D * s = g * \lambda^F$ and $\lambda^E * t = g' * \lambda^F$. We claim that g and g' are epimorphisms. By symmetry, it is sufficient to give a proof for g. Hence, let $h, h' : \mathfrak{a} \to \mathfrak{d}$ be morphisms such that $h \circ g = h' \circ g$. For every $i < \gamma^D$, it follows that

$$\begin{aligned} h \circ \lambda_i^D &= h \circ \lambda^D_{\rho(\rho_0(i))} \circ D(i, \rho(\rho_0(i))) \\ &= h \circ \lambda^D_{\rho(\rho_0(i))} \circ s_{\rho_0(i)} \circ D(i, \rho(\rho_0(i))) \\ &= h \circ g \circ \lambda^F_{\rho_0(i)} \circ D(i, \rho(\rho_0(i))) \\ &= h' \circ g \circ \lambda^F_{\rho_0(i)} \circ D(i, \rho(\rho_0(i))) \\ &= h' \circ \lambda^D_{\rho(\rho_0(i))} \circ s_{\rho_0(i)} \circ D(i, \rho(\rho_0(i))) \\ &= h' \circ \lambda^D_{\rho(\rho_0(i))} \circ D(i, \rho(\rho_0(i))) \\ &= h' \circ \lambda_i^D . \end{aligned}$$

Consequently, Lemma B3.4.2 implies that $h = h'$.

Finally, note that $\lambda_0^F : B(F(0)) \to \mathfrak{c}$ is the desired morphism $p \to q$ since, by choice of g, g', s_0, t_0, we have

$$g \circ \lambda_0^F = \lambda^D_{\rho(0)} \circ s_0 = f \quad \text{and} \quad g' \circ \lambda_0^F = \lambda^E_{\theta(0)} \circ t_0 = f' . \qquad \square$$

The preceding two results are phrased in a quite general form. Their statements can be simplified significantly if we assume that the category is $\aleph_0$-accessible, all morphisms are monomorphisms, and all epimorphisms are isomorphisms. Since in the applications below we will mainly be working in $\mathfrak{Emb}(\Sigma)$ and similar categories where these assumptions are met, we record here the corresponding simplified versions. We start by proving that, under these assumptions, every object can be written as the colimit of a chain.

Lemma 2.6. *Let $\mathcal{C}$ be a category where every morphism is a monomorphism. For every κ-filtered diagram $D : \mathcal{I} \to \mathcal{C}$ of size λ that has a colimit, there exists a κ-directed diagram $E : \mathfrak{K} \to \mathcal{C}$ of size at most λ with*

$$\varinjlim E = \varinjlim D \quad \textit{and} \quad \operatorname{rng} E^{\mathrm{obj}} = \operatorname{rng} D^{\mathrm{obj}}.$$

Proof. Fix a limiting cocone $\mu \in \operatorname{Cone}(D, \mathfrak{a})$. For the index order $\mathfrak{K}$ of the diagram E, we choose the set $K := \mathcal{I}^{\mathrm{obj}}$ where we define the order by

$$\mathfrak{i} \leq \mathfrak{j} \quad :\text{iff} \quad \mathcal{I}(\mathfrak{i}, \mathfrak{j}) \neq \emptyset .$$

Since $\mathcal{I}$ is κ-filtered, this preorder is clearly κ-directed. We define the diagram E by setting

$$E^{\mathrm{obj}}(\mathfrak{i}) := D(\mathfrak{i})$$

$$\text{and} \quad E^{\mathrm{mor}}(\mathfrak{i},\mathfrak{j}) := D(f)\,, \quad \text{for an arbitrary } f \in \mathcal{I}(\mathfrak{i},\mathfrak{j})\,.$$

First, note that E is well-defined in the sense that the value of $E(\mathfrak{i},\mathfrak{j})$ does not depend on the choice of f: if $f, f' \in \mathcal{I}(\mathfrak{i},\mathfrak{j})$, then

$$\mu_{\mathfrak{j}} \circ D(f) = \mu_{\mathfrak{i}} = \mu_{\mathfrak{j}} \circ D(f') \quad \text{implies} \quad D(f) = D(f')\,,$$

as $\mu_{\mathfrak{j}}$ is a monomorphism. Furthermore, it follows immediately from the definition that $\operatorname{rng} E^{\mathrm{obj}} = \operatorname{rng} D^{\mathrm{obj}}$.

Hence, it remains to show that D and E have the same colimit. We will prove below that $\operatorname{Cone}(E,\mathfrak{b}) = \operatorname{Cone}(D,\mathfrak{b})$, for every $\mathfrak{b} \in \mathcal{C}$. Hence, the identity maps provide a natural isomorphism

$$\mathrm{id} : \operatorname{Cone}(D,-) \to \operatorname{Cone}(E,-)$$

and it follows by Lemma B3.4.3 that D and E have the same colimits.

To prove the claim, let $\nu \in \operatorname{Cone}(D,\mathfrak{b})$. For all $\mathfrak{i} \leq \mathfrak{j}$ and $f \in \mathcal{I}(\mathfrak{i},\mathfrak{j})$, it follows that

$$\nu_{\mathfrak{i}} = \nu_{\mathfrak{j}} \circ D(f) = \nu_{\mathfrak{j}} \circ E(\mathfrak{i},\mathfrak{j})\,.$$

Hence, $\nu \in \operatorname{Cone}(E,\mathfrak{b})$. Conversely, let $\nu \in \operatorname{Cone}(E,\mathfrak{b})$. For all $f : \mathfrak{i} \to \mathfrak{j}$ in $\mathcal{I}$, it follows that

$$\nu_{\mathfrak{i}} = \nu_{\mathfrak{j}} \circ E(\mathfrak{i},\mathfrak{j}) = \nu_{\mathfrak{j}} \circ D(f)\,.$$

Hence, $\nu \in \operatorname{Cone}(D,\mathfrak{b})$. □

Corollary 2.7. *Let $\mathcal{C}$ be an $\aleph_0$-accessible category where every morphism is a monomorphism. For every κ^+-presentable object $\mathfrak{a} \in \mathcal{C}$, there exists a chain $D : \kappa \to \mathcal{C}$ such that*

- $\varinjlim D = \mathfrak{a}$,
- *every object* $D(\alpha)$ *is* κ*-presentable and,*
- *for each* κ*-presentable object* $\mathfrak{b}$*, every morphism* $f : \mathfrak{b} \to \mathfrak{a}$ *factorises essentially uniquely through every limiting cocone from* D *to* $\mathfrak{a}$*.*

Proof. If $\mathfrak{a}$ is κ-presentable, we can take the constant diagram $D : \kappa \to \mathcal{C}$ where $D(\alpha) = \mathfrak{a}$ and $D(\alpha, \beta) = \mathrm{id}_{\mathfrak{a}}$, for all $\alpha \leq \beta < \kappa$. Hence, it remains to consider the case where $\mathfrak{a}$ is κ^+-presentable, but not κ-presentable. Then we can use Theorem B4.4.3 to find an $\aleph_0$-filtered diagram $E : \mathcal{I} \to \mathcal{C}$ of size at most κ with colimit $\mathfrak{a}$ such that every object $E(\mathfrak{i})$ is $\aleph_0$-presentable. We use Lemma 2.6 to construct a $\aleph_0$-directed diagram $F : \mathfrak{K} \to \mathcal{C}$ of size at most κ with $\varinjlim F = \mathfrak{a}$ such that every $F(\mathfrak{i})$ is $\aleph_0$-presentable. By Proposition B3.4.16, there exists a chain $D : \gamma \to \mathcal{C}$ of length $\gamma \leq |K| \leq \kappa$ with colimit $\mathfrak{a}$ such that each object $D(\alpha)$ is a colimit of a directed diagram of size less than $|K|$. In particular, every $D(\alpha)$ is κ-presentable. As $\mathfrak{a}$ is not κ-presentable, it follows by Theorem B4.4.3 that $\gamma = \kappa$.

Finally, let $\mu \in \mathrm{Cone}(D, \mathfrak{a})$ be limiting. If κ is regular, the index order $\langle \kappa, \leq \rangle$ of D is κ-directed and every morphism $f : \mathfrak{b} \to \mathfrak{a}$ from a κ-presentable object $\mathfrak{b}$ to $\mathfrak{a}$ factorises essentially uniquely through μ. Hence, suppose that κ is singular. Then it follows by Lemma B4.1.4 that an object is κ-presentable if, and only if, it is κ^+-presentable. This contradicts our assumption that $\mathfrak{a}$ is κ^+-presentable but not κ-presentable. □

In the following theorem let us state the special cases of Propositions 2.4 and 2.5 that we will need below.

Theorem 2.8. *Let* $\mathcal{C}$ *be an* $\aleph_0$*-accessible category where every morphism is a monomorphism and every epimorphism is an isomorphism.*

(a) *If* $\mathfrak{a} \in \mathcal{C}$ *is* κ^+*-presentable and* $\mathfrak{a} \sqsubseteq^{\kappa}_{\mathrm{pres}} \mathfrak{b}$*, then there exists a morphism* $f : \mathfrak{a} \to \mathfrak{b}$*.*

(b) *Let* $\mathfrak{a}, \mathfrak{b} \in \mathcal{C}$ *be* κ^+*-presentable objects with* $\mathfrak{a} \equiv^{\kappa}_{\mathrm{pres}} \mathfrak{b}$*. For every partial morphism* $p = \langle f, f' \rangle \in \mathfrak{pMor}_{\kappa}(\mathfrak{a}, \mathfrak{b})$*, there exists an isomorphism* $\pi : \mathfrak{a} \to \mathfrak{b}$ *with* $f' = \pi \circ f$*.*

Proof. We start by proving that $\mathcal{C}$ and the class $\mathcal{K}$ of all κ-presentable objects satisfy the conditions of Propositions 2.4 and 2.5. Clearly, being $\aleph_0$-accessible $\mathcal{C}$ has colimits of chains.

To show that $\mathcal{K}$ is closed under colimits of nonempty chains of length less than κ, let $F : \gamma \to \mathcal{K}$ be such a chain. As every object $F(i)$, for $i < \gamma$, is κ-presentable, it follows by Proposition B4.3.7 that the colimit of F is $(\kappa \oplus |\gamma|^+)$-presentable, i.e., κ-presentable.

(a) We can use Corollary 2.7 to express $\mathfrak{a}$ as the colimit of a chain $D : \kappa \to \mathcal{K}$ of the form required by Proposition 2.4. Consequently, we obtain a diagram $F : \kappa \to \mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ such that $D = B \circ F$. Let λ be a limiting cocone from D to $\mathfrak{a}$ and set $\mu_\alpha := Q(F(\alpha))$, for $\alpha < \kappa$, where Q is the codomain projection functor. Then $\mu := (\mu_\alpha)_{\alpha<\kappa}$ is a cocone from D to $\mathfrak{b}$. As λ is limiting, there exists a morphism $f : \mathfrak{a} \to \mathfrak{b}$ such that $\mu = f * \lambda$.

(b) We can use Corollary 2.7 to express $\mathfrak{a}$ and $\mathfrak{b}$ as colimits of chains $D : \kappa \to \mathcal{K}$ and $E : \kappa \to \mathcal{K}$ of the form required by Proposition 2.5. Therefore, we obtain a morphism $h : p \to q$ of $\mathfrak{pMor}_{\mathcal{K}}(\mathfrak{a}, \mathfrak{b})$ where $q = \langle g, g' \rangle$ consists of two isomorphisms. It follows that $\pi := g' \circ g^{-1}$ is the desired isomorphism between $\mathfrak{a}$ and $\mathfrak{b}$. □

3. *Fraïssé limits*

In this section we will present a method to construct structures with an $\aleph_0$-categorical theory. These structures will be approximated by a directed diagram of finitely generated substructures. Since this construction has further applications, we will present it in the general setting of an accessible category.

Ultrahomogeneous objects

As in the case of κ-saturated structures and atomic ones, we can characterise the maximal objects of the order $\sqsubseteq^{\kappa}_{\mathrm{pres}}$. For the category $\mathfrak{Emb}(\Sigma)$, these structures will have an $\aleph_0$-categorical theory.

Definition 3.1. Let $\mathcal{C}$ be a category. An object $\mathfrak{u} \in \mathcal{C}$ is *κ-ultrahomogeneous* if, for every κ-presentable object $\mathfrak{a}$ and all pairs of morphisms $f, f' : \mathfrak{a} \to \mathfrak{u}$, there exists an automorphism $\pi : \mathfrak{u} \to \mathfrak{u}$ with $f' = \pi \circ f$.

We call an object $\mathfrak{u}$ *ultrahomogeneous* if it is $\|\mathfrak{u}\|$-ultrahomogeneous.

Example. (a) The order $\langle \mathbb{Q}, \leq \rangle$ of the rationals is ultrahomogeneous in $\mathfrak{Emb}(\leq)$.

(b) Let $\langle \omega, p \rangle$ be the structure where $p(0) := 0$ and $p(n+1) := n$. This structure is ultrahomogeneous in $\mathfrak{Emb}(p)$ since no two distinct substructures are isomorphic.

(c) We have shown in Corollary B6.5.31 that algebraically closed fields are $\aleph_0$-ultrahomogeneous.

Exercise 3.1. Find a dense linear order that is not $\aleph_0$-ultrahomogeneous in $\mathfrak{Emb}(\leq)$. Can you find an open one?

One important parameter of an ultrahomogeneous structure is the class of its substructures.

Definition 3.2. Let $\mathcal{C}$ be a category, κ an infinite cardinal, and $\mathfrak{a} \in \mathcal{C}$. We denote by $\mathrm{Sub}_\kappa(\mathfrak{a})$ the class of all κ-presentable objects $\mathfrak{c} \in \mathcal{C}$ such that there exists a morphism $\mathfrak{c} \to \mathfrak{a}$.

For accessible categories this class is well-behaved.

Lemma 3.3. *Let $\mathcal{C}$ be a κ-accessible category.*

$$\mathfrak{a} \sqsubseteq^\kappa_{\mathrm{pres}} \mathfrak{b} \quad \textit{implies} \quad \mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{b}) .$$

Proof. Let $\mathfrak{c} \in \mathrm{Sub}_\kappa(\mathfrak{a})$ and let $g : \mathfrak{c} \to \mathfrak{a}$ be a corresponding morphism. Since $\mathfrak{a} \sqsubseteq^\kappa_{\mathrm{pres}} \mathfrak{b}$, there exists a partial morphism $\langle f, f' \rangle \in \mathfrak{pMor}_\kappa(\mathfrak{a}, \mathfrak{b})$. According to Proposition B4.4.12, the category $\mathfrak{Sub}_\kappa(\mathfrak{a})$ is κ-filtered. Therefore, there exist an object $h : \mathfrak{e} \to \mathfrak{a}$ of $\mathfrak{Sub}_\kappa(\mathfrak{a})$ and morphisms $\varphi : f \to h$ and $\psi : g \to h$. Since $\mathfrak{a} \sqsubseteq^\kappa_{\mathrm{pres}} \mathfrak{b}$, we can find a morphism $h' : \mathfrak{e} \to \mathfrak{b}$ such that $\langle h, h' \rangle \in \mathfrak{pMor}_\kappa(\mathfrak{a}, \mathfrak{b})$ and

$$\varphi : \langle f, f' \rangle \to \langle h, h' \rangle .$$

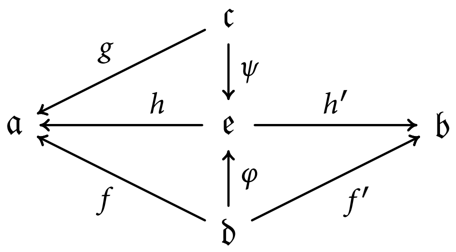

We obtain a morphism $h' \circ \psi : \mathfrak{c} \to \mathfrak{b}$ witnessing the fact that $\mathfrak{c} \in \mathrm{Sub}_\kappa(\mathfrak{b})$. □

Corollary 3.4. *Let $\mathcal{C}$ be an $\aleph_0$-accessible category where every morphism is a monomorphism, and let $\mathfrak{u}$ be κ-ultrahomogeneous. Then*

$$\mathfrak{a} \sqsubseteq^\kappa_{\mathrm{pres}} \mathfrak{u} \quad \text{iff} \quad \mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{u}) , \quad \textit{for all objects } \mathfrak{a} .$$

Proof. $(\Rightarrow)$ Since $\aleph_0$-accessible categories are κ-accessible, for all infinite cardinals κ, this direction follows from Lemma 3.3.

$(\Leftarrow)$ Let $p = \langle f, f' \rangle \in \mathfrak{pMor}_\kappa(\mathfrak{a}, \mathfrak{u})$ be a partial morphism with base $\mathfrak{c}$ and let $h : \mathfrak{c} \to \mathfrak{d}$ and $g : \mathfrak{d} \to \mathfrak{a}$ be morphisms with $g \circ h = f$ where $\mathfrak{d}$ is κ-presentable. Since $\mathfrak{d} \in \mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{u})$, there exists some morphism $g' : \mathfrak{d} \to \mathfrak{u}$. As $\mathfrak{u}$ is κ-ultrahomogeneous, we can find an automorphism $\pi : \mathfrak{u} \to \mathfrak{u}$ such that

$$f' = \pi \circ g' \circ h .$$

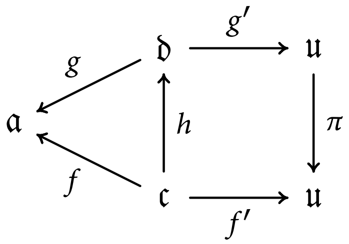

We obtain a partial morphism $q := \langle g, \pi \circ g' \rangle \in \mathfrak{pMor}_\kappa(\mathfrak{a}, \mathfrak{u})$ such that $h : p \to q$. □

The statement of the previous corollary can be used to characterise ultrahomogeneous objects.

Proposition 3.5. *Let $\mathcal{C}$ be an $\aleph_0$-accessible category where every morphism is a monomorphism and every epimorphism an isomorphism. For a κ^+-presentable object $\mathfrak{u} \in \mathcal{C}$, the following statements are equivalent:*

(1) *$\mathfrak{u}$ is κ-ultrahomogeneous.*

(2) $\mathfrak{a} \sqsubseteq^{\kappa}_{\mathrm{pres}} \mathfrak{u}$, *for all $\mathfrak{a} \in \mathcal{C}$ with* $\mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{u})$.

(3) $\mathfrak{u} \equiv^{\kappa}_{\mathrm{pres}} \mathfrak{u}$

Proof. (1) ⇒ (2) was already proved in Corollary 3.4 and (2) ⇒ (3) is trivial. Hence, it remains to prove (3) ⇒ (1). To show that $\mathfrak{u}$ is κ-ultrahomogeneous, consider morphisms $f, f' : \mathfrak{c} \to \mathfrak{u}$ with κ-presentable domain $\mathfrak{c}$. By assumption, we have $\mathfrak{u} \equiv^{\kappa}_{\mathrm{pres}} \mathfrak{u}$. Consequently, we can use Theorem 2.8 (b) to find an isomorphism $\pi : \mathfrak{u} \to \mathfrak{u}$ such that $f' = \pi \circ f$. □

Corollary 3.6. *Let $\mathcal{C}$ be an $\aleph_0$-accessible category where every morphism is a monomorphism and every epimorphism an isomorphism.*

(a) *Let $\mathfrak{u}, \mathfrak{v}$ be κ^+-presentable κ-ultrahomogeneous objects. Then*

$$\mathrm{Sub}_\kappa(\mathfrak{u}) = \mathrm{Sub}_\kappa(\mathfrak{v}) \quad \textit{implies} \quad \mathfrak{u} \cong \mathfrak{v}\,.$$

(b) *Let $\mathfrak{u}$ be κ-ultrahomogeneous and $\mathfrak{a}$ κ^+-presentable. Then*

$$\mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{u}) \quad \textit{implies} \quad \mathfrak{a} \in \mathrm{Sub}_{\kappa^+}(\mathfrak{u})\,.$$

Proof. (a) This follows by Theorem 2.8 (b) and Proposition 3.5.

(b) By Corollary 3.4, $\mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{b})$ implies $\mathfrak{a} \sqsubseteq^{\kappa}_{\mathrm{pres}} \mathfrak{u}$. Hence, the claim follows by Theorem 2.8 (a). □

We have claimed above that ultrahomogeneous structures in $\mathfrak{Emb}(\Sigma)$ have an $\aleph_0$-categorical theory. We start by showing that they are existentially closed.

Proposition 3.7. *Let $\mathfrak{U}$ be an $\aleph_0$-ultrahomogeneous structure in $\mathfrak{Emb}(\Sigma)$. Then $\mathfrak{U}$ is existentially closed in the class*

$$\mathcal{C} := \{\, \mathfrak{M} \in \mathrm{Str}[\Sigma] \mid \mathrm{Sub}_{\aleph_0}(\mathfrak{M}) \subseteq \mathrm{Sub}_{\aleph_0}(\mathfrak{U}) \,\}\,.$$

Proof. Suppose that $\mathfrak{U} \subseteq \mathfrak{M}$ for some structure $\mathfrak{M} \in \mathcal{C}$. Let $\varphi(\bar{x}, \bar{y})$ be a quantifier-free formula and $\bar{a} \subseteq U$ parameters such that

$$\mathfrak{M} \vDash \exists \bar{y} \varphi(\bar{a}, \bar{y}) .$$

We have to show that $\mathfrak{U} \vDash \exists \bar{y} \varphi(\bar{a}, \bar{y})$. Fix a tuple $\bar{b} \subseteq M$ with $\mathfrak{M} \vDash \varphi(\bar{a}, \bar{b})$. By Corollary 3.6 (b), there exists an embedding $h : \langle\!\langle \bar{a}\bar{b} \rangle\!\rangle_{\mathfrak{M}} \to \mathfrak{U}$. Since $\mathfrak{U}$ is $\aleph_0$-ultrahomogeneous and

$$\langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{U}} \cong \langle\!\langle h(\bar{a}) \rangle\!\rangle_{\mathfrak{U}}$$

we can find an automorphism π of $\mathfrak{U}$ with $\pi(h(\bar{a})) = \bar{a}$. Consequently,

$$\begin{aligned} \mathfrak{M} \vDash \varphi(\bar{a}, \bar{b}) \quad &\text{iff} \quad \langle\!\langle \bar{a}\bar{b} \rangle\!\rangle_{\mathfrak{M}} \vDash \varphi(\bar{a}, \bar{b}) \\ &\text{iff} \quad \mathfrak{U} \vDash \varphi(h(\bar{a}), h(\bar{b})) \\ &\text{iff} \quad \mathfrak{U} \vDash \varphi(\bar{a}, \pi(h(\bar{b}))) . \end{aligned}$$

Hence, $\mathfrak{U} \vDash \exists \bar{y} \varphi(\bar{a}, \bar{y})$. □

With slightly stronger assumptions we obtain $\aleph_0$-categoricity.

Proposition 3.8. *Let Σ be a finite relational signature and let $\mathfrak{U}$ be a countable ultrahomogeneous structure in $\mathfrak{Emb}(\Sigma)$. Then $\mathrm{Th}(\mathfrak{U})$ is $\aleph_0$-categorical.*

Proof. Note that, for every finite tuple $\bar{s}$ of sorts, there are only finitely many substructures $\langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{U}}$ of $\mathfrak{U}$ that are generated by a tuple $\bar{a} \in U^{\bar{s}}$ of sort $\bar{s}$. As $\mathfrak{U}$ is $\aleph_0$-ultrahomogeneous, it follows that any isomorphism between two such substructures extends to an isomorphism of $\mathfrak{U}$. Consequently, the automorphism group of $\mathfrak{U}$ is oligomorphic and it follows by Theorem 1.2 that $\mathrm{Th}(\mathfrak{U})$ is $\aleph_0$-categorical. □

Example. (a) We have seen above that $\langle \mathbb{Q}, \leq \rangle$ is $\aleph_0$-ultrahomogeneous. Consequently, it follows by the proposition that $\mathrm{Th}(\mathbb{Q}, \leq)$ is $\aleph_0$-categorical.

(b) That the restriction on the signature Σ is necessary, is shown by the example $\langle \omega, p \rangle$. We have seen above that this structures is $\aleph_0$-ultrahomogeneous, but its theory is not $\aleph_0$-categorical.

The theorems of Fraïssé

We have seen in Corollary 3.6 (a) that an ultrahomogeneous object $\mathfrak{u}$ is uniquely determined by the class $\mathrm{Sub}_\kappa(\mathfrak{u})$. Therefore it is worthwhile to characterise such classes. In the present section we will provide a characterisation in terms of the following properties.

Definition 3.9. Let $\mathcal{C}$ be a category, κ a cardinal, and $\mathcal{K} \subseteq \mathcal{C}$.

(a) The class $\mathcal{K}$ is *κ-hereditary* if

$$\mathfrak{a} \in \mathcal{K} \quad \text{implies} \quad \mathrm{Sub}_\kappa(\mathfrak{a}) \subseteq \mathcal{K}\,.$$

We call $\mathcal{K}$ *hereditary* if it is κ-hereditary, for all cardinals κ.

(b) $\mathcal{K}$ has the *κ-joint embedding property* if, for every set $X \subseteq \mathcal{K}$ of size $|X| < \kappa$, there exist an object $\mathfrak{c} \in \mathcal{K}$ and morphisms $\mathfrak{a} \to \mathfrak{c}$, for each $\mathfrak{a} \in X$.

(c) $\mathcal{K}$ has the *κ-amalgamation property* if, for every family of morphisms $f_i : \mathfrak{a} \to \mathfrak{b}_i$, $i < \gamma$, with $\mathfrak{a}, \mathfrak{b}_i \in \mathcal{K}$ and $\gamma < \kappa$, there exist an object $\mathfrak{c} \in \mathcal{K}$ and morphisms $g_i : \mathfrak{b}_i \to \mathfrak{c}$, $i < \gamma$, such that

$$g_i \circ f_i = g_k \circ f_k\,, \quad \text{for all } i, k < \gamma\,.$$

Remark. If the subcategory of $\mathcal{C}$ induced by a class $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ is κ-filtered, then Condition (F1) states that $\mathcal{K}$ has the κ-joint embedding property, and Lemma B4.1.2 implies that $\mathcal{K}$ has the κ-amalgamation property.

The converse fails in general. For instance, consider the class $\mathcal{K} \subseteq \mathfrak{Emb}(\Sigma)$ of all finitely generated structures. This class has the $\aleph_0$-joint embedding property and the $\aleph_0$-amalgamation property, but it is not $\aleph_0$-filtered: take finitely generated structures $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$ such that there are two different embeddings $f, g : \mathfrak{A} \to \mathfrak{B}$. Then $h \circ f \neq h \circ g$, for every embedding h.

Exercise 3.2. For a suitable signature Σ, find a class $\mathcal{K} \subseteq \mathfrak{Emb}(\Sigma)$ with the $\aleph_0$-amalgamation property that does not have the $\aleph_0$-joint embedding property.

Exercise 3.3. Suppose that the class $\mathcal{K}$ is closed under unions of chains of length less than κ. Prove that, if $\mathcal{K}$ has the $\aleph_0$-joint embedding property, it also has the κ-joint embedding property and that, if it has the $\aleph_0$-amalgamation property, it has the κ-amalgamation property.

Before giving a characterisation of classes of the form $\mathrm{Sub}_\kappa(\mathfrak{a})$, we start with a technical remark on such classes for κ-filtered colimits.

Lemma 3.10. *Let $\mathfrak{a}$ be the colimit of a κ-filtered diagram $D : \mathcal{I} \to \mathcal{C}$. Then*

$$\mathrm{Sub}_\kappa(\mathfrak{a}) = \bigcup_{\mathfrak{i} \in \mathcal{I}} \mathrm{Sub}_\kappa(D(\mathfrak{i}))\,.$$

Proof. Let $\lambda \in \mathrm{Cone}(D, \mathfrak{a})$ be a limiting cocone.

($\supseteq$) For every $\mathfrak{b} \in \mathrm{Sub}_\kappa(D(\mathfrak{i}))$, there is some morphism $f : \mathfrak{b} \to D(\mathfrak{i})$. Hence, $\lambda_\mathfrak{i} \circ f$ is a morphism $\mathfrak{b} \to \mathfrak{a}$.

($\subseteq$) Let $\mathfrak{b} \in \mathrm{Sub}_\kappa(\mathfrak{a})$ and let $f : \mathfrak{b} \to \mathfrak{a}$ be the corresponding morphism. Since $\mathfrak{b}$ is κ-presentable, we can find a morphism $f_0 : \mathfrak{b} \to D(\mathfrak{i})$, for some $\mathfrak{i} \in \mathcal{I}$, such that $f = \lambda_\mathfrak{i} \circ f_0$. Hence, $\mathfrak{b} \in \mathrm{Sub}_\kappa(D(\mathfrak{i}))$. □

Let us characterise when a class is of the form $\mathrm{Sub}_\kappa(\mathfrak{a})$, for an arbitrary object $\mathfrak{a}$. We start with an obvious necessary condition.

Proposition 3.11. *Let $\mathcal{C}$ be a κ-accessible category. For every object $\mathfrak{a} \in \mathcal{C}$, the class $\mathrm{Sub}_\kappa(\mathfrak{a})$ is κ-hereditary and it has the κ-joint embedding property.*

Proof. Clearly, if there are morphisms $\mathfrak{b} \to \mathfrak{a}$ and $\mathfrak{c} \to \mathfrak{b}$, there is also a morphism $\mathfrak{c} \to \mathfrak{a}$. Hence, $\mathrm{Sub}_\kappa(\mathfrak{b}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{a})$, for every $\mathfrak{b} \in \mathrm{Sub}_\kappa(\mathfrak{a})$.

Furthermore, we have shown in Proposition B4.4.12 that $\mathfrak{Sub}_\kappa(\mathfrak{a})$ is κ-filtered. This implies that $\mathrm{Sub}_\kappa(\mathfrak{a})$ has the κ-joint embedding property. □

The converse only holds for $\kappa = \aleph_0$ and if $\mathcal{K}$ is small enough.

Theorem 3.12 (Fraïssé). *Let $\mathcal{C}$ be an $\aleph_0$-accessible category and let $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ be a class of $\aleph_0$-presentable objects that, up to isomorphism, contains only countably many objects. If $\mathcal{K}$ is $\aleph_0$-hereditary and if it has the $\aleph_0$-joint embedding property, then $\mathcal{K} = \mathrm{Sub}_{\aleph_0}(\mathfrak{a})$, for some $\aleph_1$-presentable object $\mathfrak{a} \in \mathcal{C}$.*

Proof. Fix an enumeration $(\mathfrak{c}_n)_{n<\omega}$ of all objects in $\mathcal{K}$ up to isomorphism. We define a diagram $D : \omega \to \mathcal{K}$ by induction on n. Set $D(0) := \mathfrak{c}_0$. If $D(n)$ is already defined then, by the $\aleph_0$-joint embedding property, we can find an object $D(n+1) \in \mathcal{K}$ with morphisms $\mathfrak{c}_{n+1} \to D(n+1)$ and $f_n : D(n) \to D(n+1)$. Setting

$$D(i,k) := f_{k-1} \circ \dots \circ f_i \,, \quad \text{for } i < k < \omega \,,$$

we obtain a $\aleph_0$-directed diagram $D : \omega \to \mathcal{K}$. Let $\mathfrak{a}$ be its colimit. According to Proposition B4.3.7, $\mathfrak{a}$ is $\aleph_1$-presentable. Since $\mathcal{K}$ is $\aleph_0$-hereditary,

$$D(n) \in \mathcal{K} \quad \text{implies} \quad \mathrm{Sub}_{\aleph_0}(D(n)) \subseteq \mathcal{K} \,, \quad \text{for every } n < \omega \,.$$

By Lemma 3.10, it follows that $\mathrm{Sub}_{\aleph_0}(\mathfrak{a}) \subseteq \mathcal{K}$. Conversely, we have

$$\mathfrak{c}_n \in \mathrm{Sub}_{\aleph_0}(D(n)) \subseteq \mathrm{Sub}_{\aleph_0}(\mathfrak{a}) \,, \quad \text{for every } n < \omega \,.$$

Since $\mathrm{Sub}_{\aleph_0}(\mathfrak{a})$ is closed under isomorphisms, this implies that $\mathcal{K} \subseteq \mathrm{Sub}_{\aleph_0}(\mathfrak{a})$. □

For a given class $\mathcal{K}$ there may be several non-isomorphic objects $\mathfrak{a}$ such that $\mathcal{K} = \mathrm{Sub}_{\aleph_0}(\mathfrak{a})$. For instance, if $\mathcal{K} \subseteq \mathfrak{Emb}(\leq)$ is the class of all finite linear orders then $\mathcal{K} = \mathrm{Sub}_{\aleph_0}(\mathfrak{L})$, for every infinite linear order $\mathfrak{L}$. We are looking for an object $\mathfrak{a}$ with $\mathrm{Sub}_{\aleph_0}(\mathfrak{a}) = \mathcal{K}$ that is in a certain sense the most general one. As we have seen in Corollary 3.6, ultrahomogeneous objects $\mathfrak{u}$ are uniquely determined by $\mathrm{Sub}_\kappa(\mathfrak{u})$. Therefore, we can take ultrahomogeneity as the required additional property.

Definition 3.13. Let $\mathcal{C}$ be a category and $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$. An object $\mathfrak{f} \in \mathcal{C}$ is a *Fraïssé limit* of $\mathcal{K}$ if there is some cardinal κ such that $\mathfrak{f}$ is κ^+-presentable, κ-ultrahomogeneous, and $\mathrm{Sub}_\kappa(\mathfrak{f}) = \mathcal{K}$.

Example. $\langle \mathbb{Q}, \leq \rangle$ is the Fraïssé limit of the class of all finite linear orders in $\mathfrak{Emb}(\leq)$.

Before considering their existence, let us prove that Fraïssé limits are unique.

Proposition 3.14. *Let $\mathcal{C}$ be an $\aleph_0$-accessible category where every morphism is a monomorphism and every epimorphism an isomorphism. Up to isomorphism, a class $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ has at most one Fraïssé limit.*

Proof. Suppose that $\mathfrak{f}$ and $\mathfrak{g}$ are Fraïssé limits of $\mathcal{K}$. By definition, there are infinite cardinals κ and λ such that $\mathfrak{f}$ is κ^+-presentable and κ-ultrahomogeneous, $\mathfrak{g}$ is λ^+-presentable and λ-ultrahomogeneous, and

$$\mathrm{Sub}_\kappa(\mathfrak{f}) = \mathcal{K} = \mathrm{Sub}_\lambda(\mathfrak{g}) .$$

By symmetry, we may assume that $\kappa \leq \lambda$. As every object in $\mathrm{Sub}_\lambda(\mathfrak{g}) = \mathrm{Sub}_\kappa(\mathfrak{f})$ is κ-presentable, we have

$$\mathrm{Sub}_\kappa(\mathfrak{g}) = \mathrm{Sub}_\lambda(\mathfrak{g}) = \mathcal{K} = \mathrm{Sub}_\kappa(\mathfrak{f})$$

and it follows by Corollary 3.6 (b) that there exists a morphism $\mathfrak{f} \to \mathfrak{g}$. Consequently,

$$\mathrm{Sub}_\lambda(\mathfrak{f}) \subseteq \mathrm{Sub}_\lambda(\mathfrak{g}) = \mathcal{K} = \mathrm{Sub}_\kappa(\mathfrak{f}) \subseteq \mathrm{Sub}_\lambda(\mathfrak{f}) .$$

Hence, $\mathrm{Sub}_\lambda(\mathfrak{f}) = \mathrm{Sub}_\lambda(\mathfrak{g})$ and, if we can show that $\mathfrak{f}$ is λ-ultrahomogeneous, it will follow by by Corollary 3.6 (a) that $\mathfrak{f} \cong \mathfrak{g}$.

For λ-ultrahomogeneity of $\mathfrak{f}$, consider two morphisms $f, f' : \mathfrak{a} \to \mathfrak{f}$ with λ-presentable domain $\mathfrak{a}$. Then $\mathfrak{a} \in \mathrm{Sub}_\lambda(\mathfrak{f}) = \mathrm{Sub}_\lambda(\mathfrak{g}) = \mathrm{Sub}_\kappa(\mathfrak{g})$ implies that $\mathfrak{a}$ is even κ-presentable. Hence, we can use κ-ultrahomogeneity of $\mathfrak{f}$ to find the desired automorphism $\pi : \mathfrak{f} \to \mathfrak{f}$ with $f' = \pi \circ f$. □

Next, let us describe $\mathrm{Sub}_\kappa(\mathfrak{u})$ for a κ-ultrahomogeneous object $\mathfrak{u}$.

Lemma 3.15. *Let $\mathcal{C}$ be an $\aleph_0$-accessible category where every morphism is a monomorphism. If $\mathfrak{u} \in \mathcal{C}$ is κ-ultrahomogeneous then* $\mathrm{Sub}_\kappa(\mathfrak{u})$ *is κ-hereditary, closed under colimits of nonempty chains of length less than κ, and it has the κ-joint embedding property and the κ-amalgamation property.*

Proof. Note that every $\aleph_0$-accessible category is also κ-accessible. Therefore, it follows by Proposition 3.11 that the class $\mathrm{Sub}_\kappa(\mathfrak{u})$ is κ-hereditary and that it has the κ-joint embedding property. To check the κ-amalgamation property, let $f_i : \mathfrak{a} \to \mathfrak{b}_i$, $i < \gamma$, be a family of $\gamma < \kappa$ morphisms with $\mathfrak{a}, \mathfrak{b}_i \in \mathrm{Sub}_\kappa(\mathfrak{u})$. Fix morphisms $h_i : \mathfrak{b}_i \to \mathfrak{u}$, for $i < \gamma$. Since $\mathfrak{u}$ is κ-ultrahomogeneous, there exist automorphisms $\pi_i \in \mathrm{Aut}(\mathfrak{u})$ such that

$$\pi_i \circ h_i \circ f_i = h_0 \circ f_0 \,, \quad \text{for all } i < \gamma \,.$$

Consequently, $f_i : h_0 \circ f_0 \to \pi_i \circ h_i$ is a morphism of $\mathfrak{Sub}_\kappa(\mathfrak{u})$. We have seen in Proposition B4.4.12 that $\mathfrak{Sub}_\kappa(\mathfrak{u})$ is κ-filtered. Therefore, we can use Lemma B4.1.2 to find an object $g \in \mathfrak{Sub}_\kappa(\mathfrak{u})$ and morphisms

$$\varphi_i : \pi_i \circ h_i \to g \,, \quad \text{for } i < \gamma \,,$$

such that

$$\varphi_i \circ f_i = \varphi_k \circ f_k \,, \quad \text{for all } i, k < \gamma \,.$$

This family witnesses the κ-amalgamation property.

It remains to check that $\mathrm{Sub}_\kappa(\mathfrak{u})$ is closed under colimits of nonempty chains of length less than κ. Let $D : \gamma \to \mathrm{Sub}_\kappa(\mathfrak{u})$ be a chain of length $0 < \gamma < \kappa$. As $\mathcal{C}$ is $\aleph_0$-accessible, D has a colimit $\mathfrak{a}$ which, according to Theorem B4.4.3, is κ-presentable. Furthermore, Lemma 3.10 implies that

$$\mathrm{Sub}_\kappa(\mathfrak{a}) = \bigcup_{\alpha < \kappa} \mathrm{Sub}_\kappa(D(\alpha)) \subseteq \mathrm{Sub}_\kappa(\mathfrak{u}) \,.$$

Hence, it follows by Corollary 3.4 that $\mathfrak{a} \sqsubseteq^\kappa_{\mathrm{pres}} \mathfrak{u}$. Consequently, we can use Lemma 2.3 (b) to find a morphism $\mathfrak{a} \to \mathfrak{u}$. Thus, $\mathfrak{a} \in \mathrm{Sub}_\kappa(\mathfrak{u})$. □

The converse is given by the following theorem, which can be used to construct ultrahomogeneous structures by describing their class of substructures. Again we have to require $\mathcal{K}$ to be small enough.

Theorem 3.16 (Fraïssé). *Let κ be a regular cardinal, let $\mathcal{C}$ be an $\aleph_0$-accessible category where all morphisms are monomorphisms and all epimorphisms are isomorphisms, and let $\mathcal{K} \subseteq \mathcal{C}^{\mathrm{obj}}$ be a κ-hereditary class of κ-presentable objects that is closed under nonempty chains of length less* *than κ and that has the $\aleph_0$-joint embedding property and the $\aleph_0$-amalgamation property, and such that the full subcategory of $\mathcal{C}$ induced by $\mathcal{K}$ has a skeleton $\mathcal{K}_0$ with at most κ morphisms. Then $\mathcal{K}$ has a Fraïssé limit $\mathfrak{f}$.*

Proof. We will construct a diagram $D : \kappa \to \mathcal{K}_0$ satisfying the following condition:

($*$) If $f : \mathfrak{a} \to \mathfrak{b}$ and $g : \mathfrak{a} \to D(\alpha)$ are morphisms with $\mathfrak{a}, \mathfrak{b} \in \mathcal{K}_0$, there is some index $\beta > \alpha$ and a morphism $g' : \mathfrak{b} \to D(\beta)$ such that

$$g' \circ f = D(\alpha, \beta) \circ g . \qquad \begin{array}{ccc} \mathfrak{b} & \xrightarrow{g'} & D(\beta) \\ {\scriptstyle f}\uparrow & & \uparrow{\scriptstyle D(\alpha,\beta)} \\ \mathfrak{a} & \xrightarrow[g]{} & D(\alpha) \end{array}$$

Let $\mathfrak{f}$ be the colimit of this diagram. By Theorem B4.4.3, $\mathfrak{f}$ is κ^+-presentable, and Lemma 3.10 implies that $\mathrm{Sub}_\kappa(\mathfrak{f}) \subseteq \mathcal{K}$. Conversely, if $\mathfrak{a} \in \mathcal{K}$ then, by the $\aleph_0$-joint embedding property, there are an object $\mathfrak{b} \in \mathcal{K}$ and morphisms $h : \mathfrak{a} \to \mathfrak{b}$ and $f : D(0) \to \mathfrak{b}$. By ($*$), we can extend the identity morphism $\mathrm{id} : D(0) \to D(0)$ to a morphism $g' : \mathfrak{b} \to D(\alpha)$, for some $\alpha > 0$. Consequently, $\mathfrak{b} \in \mathrm{Sub}_\kappa(D(\alpha)) \subseteq \mathrm{Sub}_\kappa(\mathfrak{f})$ and $\mathfrak{a} \in \mathrm{Sub}_\kappa(\mathfrak{b}) \subseteq \mathrm{Sub}_\kappa(\mathfrak{f})$. It follows that $\mathcal{K} = \mathrm{Sub}_\kappa(\mathfrak{f})$.

To show that $\mathfrak{f}$ is ultrahomogeneous it is sufficient, by Proposition 3.5, to prove that $\mathfrak{f} \sqsubseteq^\kappa_{\mathrm{pres}} \mathfrak{f}$. Consider morphisms $f : \mathfrak{a} \to \mathfrak{f}$, $f' : \mathfrak{a} \to \mathfrak{f}$, $g : \mathfrak{b} \to \mathfrak{f}$, $h : \mathfrak{a} \to \mathfrak{b}$ such that $f = g \circ h$ and $\mathfrak{a}$ and $\mathfrak{b}$ are κ-presentable. As κ is regular, the order $\langle \kappa, \leq \rangle$ is κ-directed. Since $\mathfrak{a}$ is κ-presentable, there therefore exists an essentially unique factorisation $f' = \lambda_\alpha \circ f'_0$, for some index $\alpha < \kappa$, some morphism $f'_0 : \mathfrak{a} \to D(\alpha)$, and a limiting cocone λ from D to $\mathfrak{f}$. Hence, we can use ($*$) to find an index $\beta > \alpha$ and a morphism $g' : \mathfrak{b} \to D(\beta)$ such that

$$g' \circ h = D(\alpha, \beta) \circ f_0' .$$

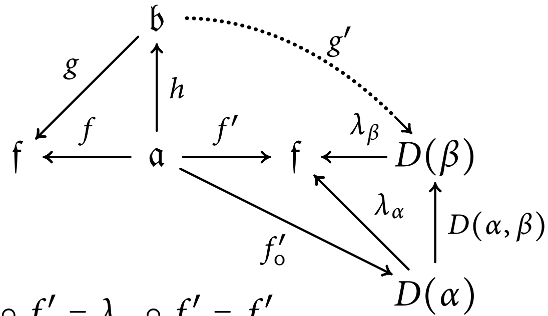

Since

$$\lambda_\beta \circ g' \circ h = \lambda_\beta \circ D(\alpha, \beta) \circ f_0' = \lambda_\alpha \circ f_0' = f' ,$$

it follows that $\langle g, \lambda_\beta \circ g' \rangle$ is a partial morphism with

$$h : \langle f, f' \rangle \to \langle g, \lambda_\beta \circ g' \rangle .$$

Consequently, $\mathfrak{f}$ is a Fraïssé limit of $\mathcal{K}$.

It remains to construct a chain $D : \kappa \to \mathcal{K}_0$ satisfying $(*)$. Choose a bijection $\pi : \kappa \times \kappa \to \kappa$ such that $\pi(\alpha, \beta) \geq \alpha$, for all $\alpha, \beta < \kappa$. (For instance, the bijection constructed in the proof of Theorem A4.3.8 has this property.) We construct $D(\alpha)$ by induction on α. We start with an arbitrary object $D(0) \in \mathcal{K}_0$. For the successor step, suppose that $D(\alpha)$ has already been defined. Fix a list of all pairs $\langle f_{\alpha\beta}, g_{\alpha\beta} \rangle$, for $\beta < \kappa$, where $f_{\alpha\beta} : \mathfrak{a}_{\alpha\beta} \to \mathfrak{b}_{\alpha\beta}$ is a morphism in $\mathcal{K}_0$ and $g_{\alpha\beta} : \mathfrak{a}_{\alpha\beta} \to D(\alpha)$ is an arbitrary morphism. Let $\langle \gamma, \beta \rangle := \pi^{-1}(\alpha)$. Note that we have chosen π such that $\gamma \leq \alpha$. By the $\aleph_0$-amalgamation property, we can find a structure $\mathfrak{c} \in \mathcal{K}$ and morphisms $h_{\gamma\beta} : \mathfrak{b}_{\gamma\beta} \to \mathfrak{c}$ and $h'_{\gamma\beta} : D(\alpha) \to \mathfrak{c}$ such that

$$h_{\gamma\beta} \circ f_{\gamma\beta} = h'_{\gamma\beta} \circ D(\gamma, \alpha) \circ g_{\gamma\beta} .$$

We set

$$D(\alpha + 1) := \mathfrak{c} \quad \text{and} \quad D(i, \alpha + 1) := h'_{\gamma\beta} \circ D(i, \alpha) , \quad \text{for } i \leq \alpha .$$

For the limit step, suppose that $D(\alpha)$ is already defined for all $\alpha < \delta$. Let $D(\delta) := \varinjlim (D \restriction \delta)$ and let λ be a corresponding limiting cocone. By assumption $D(\delta) \in \mathcal{K}_0$ and we can set $D(\alpha, \delta) := \lambda_\alpha$, for $\alpha < \delta$.

We claim that the diagram D defined this way satisfies Condition $(*)$. Let $f : \mathfrak{a} \to \mathfrak{b}$ and $g : \mathfrak{a} \to D(\alpha)$ be morphisms with $\mathfrak{a}, \mathfrak{b} \in \mathcal{K}_0$. Then

$\langle f, g\rangle = \langle f_{\alpha\beta}, g_{\alpha\beta}\rangle$, for some ordinal $\beta < \kappa$. Consequently, the morphism $h_{\alpha\beta} : \mathfrak{b}_{\alpha\beta} \to D(\pi(\alpha,\beta) + 1)$ chosen in the inductive step above satisfies

$$\begin{aligned} h_{\alpha\beta} \circ f_{\alpha\beta} &= h'_{\alpha\beta} \circ D(\alpha, \pi(\alpha,\beta)) \circ g_{\alpha\beta} \\ &= D(\alpha, \pi(\alpha,\beta) + 1) \circ g_{\alpha\beta} . \end{aligned}$$ □

Example. (a) Let $\mathcal{P} \subseteq \mathfrak{Emb}(E)$ be the class of all finite planar graphs. Clearly, $\mathcal{P}$ is hereditary. The class $\mathcal{P}$ does not have a Fraïssé limit since it does not have the $\aleph_0$-amalgamation property. Consider the following graphs:

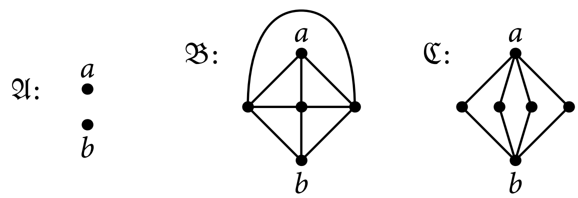

Let $f : \mathfrak{A} \to \mathfrak{B}$ and $g : \mathfrak{A} \to \mathfrak{C}$ be the embeddings with $a \mapsto a$ and $b \mapsto b$. There is no planar graph $\mathfrak{D}$ such that we can find embeddings $h : \mathfrak{B} \to \mathfrak{D}$ and $k : \mathfrak{C} \to \mathfrak{D}$ with $h \circ f = k \circ g$.

(b) Similarly we can show that the class $\mathcal{F} \subseteq \mathfrak{Emb}(E)$ of all finite acyclic graphs does not have the $\aleph_0$-amalgamation property. The counterexample is given by the graphs:

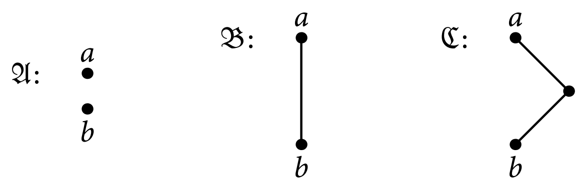

4. *Zero-one laws*

In this section we study Fraïssé limits by axiomatising their theories.

Definition 4.1. (a) Let $\mathfrak{M}$ be a structure. The *atomic type* of $\bar{a} \subseteq M$ is the set

$$\operatorname{atp}(\bar{a}) := \{ \varphi \mid \varphi \text{ a literal such that } \mathfrak{M} \vDash \varphi(\bar{a}) \} .$$

An *atomic n-type* $\mathfrak{p}$ is a set of the form $\mathfrak{p} = \operatorname{atp}(\bar{a})$, for $\bar{a} \in M^n$.

(b) Let $\mathfrak{p}$ be an atomic n-type and $\mathfrak{q}$ an atomic $(n+1)$-type such that $\mathfrak{p} \subseteq \mathfrak{q}$. The *extension axiom* associated with $\mathfrak{p}$ and $\mathfrak{q}$ is the sentence

$$\eta_{\mathfrak{pq}} := \forall \bar{x}[\mathfrak{p}(\bar{x}) \rightarrow \exists y \mathfrak{q}(\bar{x}, y)] .$$

(We write $\mathfrak{p}(\bar{x})$ for the formula $\bigwedge \mathfrak{p}$.)

(c) Let $\mathcal{K}$ be a hereditary class of finitely generated structures. We define

$$\Gamma_{\mathcal{K}} := \{ \operatorname{atp}(\bar{a}/\mathfrak{M}) \mid \bar{a} \text{ is a finite tuple generating } \mathfrak{M} \in \mathcal{K} \} ,$$

and $\quad T[\mathcal{K}] := \{ \eta_{\mathfrak{pq}} \mid \mathfrak{q} \in \Gamma_{\mathcal{K}} \} \cup \{ \forall \bar{x} \neg \mathfrak{p}(\bar{x}) \mid \mathfrak{p} \notin \Gamma_{\mathcal{K}} \}$

The set of all extension axioms over a signature Σ is $T_{\mathrm{ran}}[\Sigma] := T[\mathcal{C}]$, where $\mathcal{C}$ is the class of all finitely generated Σ-structures.

Remark. Note that, in general, $T[\mathcal{K}]$ is an infinitary theory. It is a first-order theory if the signature in question is finite and relational.

Example. An important example of a Fraïssé limit is the *random graph*, also called the *Rado graph*. It can be defined as follows. $\mathfrak{R} := \langle V, E \rangle$ where $V :=$ HF is the set of all hereditary finite sets and the edge relation is

$$E := \{ \langle a, b \rangle \mid a \in b \text{ or } b \in a \} .$$

This graph satisfies the following extension axiom: for every pair X, Y of finite disjoint sets of vertices, there exists some vertex $c \in V$ that is adjacent to every vertex in X, but not adjacent to any in Y. For a proof, note that, if $X = \{a_0, \dots, a_{m-1}\}$ and $Y = \{b_0, \dots, b_{n-1}\}$ then we can take $c := \{a_0, \dots, a_{m-1}, x\}$ where the set $x := \{b_0, \dots, b_{n-1}\}$ is needed to ensure that $c \notin b_i$.

Let us investigate the relationship between the theories $T[\mathcal{K}]$ and ultrahomogeneous structures.

Lemma 4.2. *If* $\mathfrak{U}$ *is ultrahomogeneous then* $\mathfrak{U} \vDash T[\mathrm{Sub}_{\aleph_0}(\mathfrak{U})]$.

Lemma 4.3. *If* $\mathfrak{A}, \mathfrak{B} \vDash T[\mathcal{K}]$ *then*

$$\mathfrak{A} \equiv_0 \mathfrak{B} \quad \textit{implies} \quad \mathfrak{A} \cong_0^{\aleph_0} \mathfrak{B} .$$

Proof. Since $\mathfrak{A} \equiv_0 \mathfrak{B}$ we have $\mathrm{pIso}_{\aleph_0}(\mathfrak{A}, \mathfrak{B}) \neq \varnothing$. To check the forth condition, let $\bar{a} \mapsto \bar{b} \in \mathrm{pIso}_{\aleph_0}(\mathfrak{A}, \mathfrak{B})$ and $c \in A$. Set $\mathfrak{p} := \mathrm{atp}(\bar{a})$ and $\mathfrak{q} := \mathrm{atp}(\bar{a}c)$. Then $\mathfrak{p} \subseteq \mathfrak{q}$ and $\mathfrak{q} \in \Gamma_{\mathcal{K}}$. Hence, $\eta_{\mathfrak{p}\mathfrak{q}} \in T[\mathcal{K}]$ and $\mathfrak{B} \vDash \eta_{\mathfrak{p}\mathfrak{q}}$. Since $\mathrm{atp}(\bar{b}) = \mathfrak{p}$ we can, therefore, find some $d \in B$ with $\mathrm{atp}(\bar{b}d) = \mathfrak{q}$. Consequently, $\bar{a}c \mapsto \bar{b}d \in \mathrm{pIso}_{\aleph_0}(\mathfrak{A}, \mathfrak{B})$. ☐

Corollary 4.4. *Every model of* $T[\mathcal{K}]$ *is ultrahomogeneous.*

It follows that the theories $T[\mathcal{K}]$ axiomatise Fraïssé limits.

Theorem 4.5. *Let* $\mathcal{K}$ *be a hereditary class of finitely generated structures containing a unique* 0*-generated structure* $\mathfrak{A}_0$. *A structure* $\mathfrak{F}$ *is the Fraïssé limit of* $\mathcal{K}$ *if and only if it is countable,* $\langle\!\langle \varnothing \rangle\!\rangle_{\mathfrak{F}} \cong \mathfrak{A}_0$, *and* $\mathfrak{F} \vDash T[\mathcal{K}]$.

Proof. $(\Rightarrow)$ A Fraïssé limit $\mathfrak{F}$ is countable by definition. Furthermore, $\mathrm{Sub}_{\aleph_0}(\mathfrak{F}) \subseteq \mathcal{K}$ implies that $\mathfrak{F} \vDash \forall \bar{x} \neg \mathfrak{p}(\bar{x})$, for all $\mathfrak{p} \notin \Gamma_{\mathcal{K}}$.

Finally, let $\eta_{\mathfrak{p}\mathfrak{q}} \in T[\mathcal{K}]$. Then $\mathfrak{q} \in \Gamma_{\mathcal{K}}$ and $\mathcal{K} \subseteq \mathrm{Sub}_{\aleph_0}(\mathfrak{F})$ implies that there is some tuple $\bar{c} \subseteq F$ with $\mathrm{atp}(\bar{c}) = \mathfrak{q}$. Since $\mathfrak{F}$ is ultrahomogeneous it follows that, for every tuple $\bar{a}$ with $\mathrm{atp}(\bar{a}) = \mathfrak{p}$, there is some element $b \in F$ such that $\mathrm{atp}(\bar{a}b) = \mathrm{atp}(\bar{c}) = \mathfrak{q}$. Hence, $\mathfrak{F} \vDash \eta_{\mathfrak{p}\mathfrak{q}}$.

$(\Leftarrow)$ By assumption, $\mathfrak{F}$ is countable, and we have shown in Corollary 4.4 that it is ultrahomogeneous. Furthermore, $\mathfrak{F} \vDash \forall \bar{x} \neg \mathfrak{p}(\bar{x})$, for $\mathfrak{p} \notin \Gamma_{\mathcal{K}}$ implies that $\mathrm{Sub}_{\aleph_0}(\mathfrak{F}) \subseteq \mathcal{K}$. Hence, it remains to show that $\mathcal{K} \subseteq \mathrm{Sub}_{\aleph_0}(\mathfrak{F})$. Let $\mathfrak{B} \in \mathcal{K}$ be generated by a finite tuple $\bar{b} = b_0 \dots b_{n-1}$. Note that $\langle\!\langle \varnothing \rangle\!\rangle_{\mathfrak{B}} \cong \mathfrak{A}_0 \cong \langle\!\langle \varnothing \rangle\!\rangle_{\mathfrak{F}} \subseteq \mathfrak{F}$. Since $\mathfrak{F}$ satisfies the needed extension axioms we can, therefore, use induction to find elements $a_0, \dots, a_{n-1} \in F$ such that

$$\langle\!\langle b_0 \dots b_{k-1} \rangle\!\rangle_{\mathfrak{B}} \cong \langle\!\langle a_0, \dots, a_{k-1} \rangle\!\rangle_{\mathfrak{F}} , \quad \text{for all } k \leq n .$$

Consequently, we have $\mathfrak{B} = \langle\!\langle \bar{b} \rangle\!\rangle_{\mathfrak{B}} \cong \langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{F}} \subseteq \mathfrak{F}$. ☐

Proposition 4.6. *$T[\mathcal{K}]$ admits quantifier elimination for* $\mathrm{FO}_{\infty\aleph_0}$.

Proof. This follows immediately from Theorem D1.2.9 and Lemma 4.3. ☐

Corollary 4.7. *Let $\mathcal{K}$ be a class of Σ-structures where the signature Σ is finite and relational. Then $T[\mathcal{K}]$ admits quantifier elimination for* FO.

Proof. Since $T[\mathcal{K}]$ is a first-order theory, the claim follows by Corollary D1.2.10. ☐

Corollary 4.8. *Let $\mathcal{K}$ be a class of Σ-structures where Σ is a finite, relational signature without* 0*-ary relations. Then $T[\mathcal{K}]$ is complete.*

Proof. Let $\varphi \in \mathrm{FO}^0[\Sigma]$. There exists a sentence $\psi \in \mathrm{QF}^0[\Sigma]$ such that $T[\mathcal{K}] \vDash \varphi \leftrightarrow \psi$. Since Σ is relational and it contains no 0-ary relations, the only quantifier-free sentences are true and false. If $\psi \equiv$ true then $T[\mathcal{K}] \vDash \varphi$ and if $\psi \equiv$ false then $T[\mathcal{K}] \vDash \neg\varphi$. ☐

The extension axioms have the surprising property that, asymptotically, they hold with probability 1 in every finite structure. Let us make this claim more precise.

Consider a finite signature Σ. For each finite number $n < \omega$, we count how many Σ-structures with universe $[n]$ satisfy a given sentence. Note that, for every n, there are only finitely many such structures.

Definition 4.9. For $\varphi, \psi \in \mathrm{FO}[\Sigma]$ we define

$$\kappa_n(\varphi) := |\{\, \mathfrak{M} \mid \mathfrak{M} \vDash \varphi,\ M = [n] \,\}|\,,$$

$$\Pr^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \varphi \mid \mathfrak{M} \vDash \psi \,] := \frac{\kappa_n(\varphi \wedge \psi)}{\kappa_n(\psi)}\,.$$

We use the shorthand $\Pr^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \varphi \,] := \Pr^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \varphi \mid \mathfrak{M} \vDash \text{true} \,]$.

Lemma 4.10. *Let Σ be a finite, relational signature without* 0*-ary relations. Then*

$$\lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \eta_{\mathfrak{pq}} \,] = 1\,, \quad \textit{for every } \eta_{\mathfrak{pq}} \in T_{\mathrm{ran}}[\Sigma]\,.$$

Proof. Suppose that $\mathfrak{p}$ is an m-type and $n > m$. Since Σ is finite there exists some constant $p \in (0, 1)$ such that

$$\mathrm{Pr}^n_{\mathfrak{M}}\big[\, \mathfrak{M} \nvDash \mathfrak{q}(0, \dots, m-1, m) \,\big|\, \mathfrak{M} \vDash \mathfrak{p}(0, \dots, m-1) \,\big] = p\,.$$

Hence,

$$\mathrm{Pr}^n_{\mathfrak{M}}\big[\, \mathfrak{M} \nvDash \exists x_m \mathfrak{q}(0, \dots, m-1, x_m) \,\big|\, \mathfrak{M} \vDash \mathfrak{p}(0, \dots, m-1) \,\big] \\ = p^{n-m},$$

which implies that $\mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \nvDash \eta_{\mathfrak{pq}} \,] \leq n^m k^{n-m}$. Since $p < 1$ we have

$$\lim_{n \to \infty} n^m k^{n-m} = 0\,,$$

and it follows that

$$\lim_{n \to \infty} \mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \eta_{\mathfrak{pq}} \,] \geq \lim_{n \to \infty} (1 - n^m k^{n-m}) = 1\,.$$ □

Lemma 4.11. *$T_{\mathrm{ran}}[\Sigma]$ is satisfiable, for every finite relational signature Σ without 0-ary relations.*

Proof. For a contradiction suppose that $T_{\mathrm{ran}}[\Sigma]$ is inconsistent. Then there exists a finite inconsistent set $\Phi \subseteq T_{\mathrm{ran}}[\Sigma]$. Suppose that $\Phi = \{\varphi_0, \dots, \varphi_{m-1}\}$. By the preceding lemma, we have

$$\lim_{n \to \infty} \mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \varphi_i \,] = 1\,, \quad \text{for all } i < m\,.$$

Therefore, there exists some number n such that

$$\mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \neg\varphi_i \,] < \frac{1}{m}\,.$$

It follows that

$$\mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \bigwedge \Phi \,] = 1 - \mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \bigvee_i \neg\varphi_i \,] \\ \geq 1 - \sum_i \mathrm{Pr}^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \neg\varphi_i \,] > 1 - m \cdot \frac{1}{m} = 0\,.$$

Consequently, Φ has a model of size n. Contradiction. □

Theorem 4.12 (Zero-One Law). *Let Σ be a finite, relational signature without 0-ary relations. For every sentence $\varphi \in \mathrm{FO}[\Sigma]$, we have*

$$\lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \varphi\,] \in \{0,1\}\,.$$

Proof. If $T_{\mathrm{ran}}[\Sigma] \vDash \varphi$ then there are axioms $\eta_{\mathfrak{p}_0\mathfrak{q}_0}, \dots, \eta_{\mathfrak{p}_k\mathfrak{q}_k} \in T_{\mathrm{ran}}[\Sigma]$ such that $\eta_{\mathfrak{p}_0\mathfrak{q}_0} \wedge \dots \wedge \eta_{\mathfrak{p}_k\mathfrak{q}_k} \vDash \varphi$. Hence, we have

$$\lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \varphi\,] \geq \lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \eta_{\mathfrak{p}_0\mathfrak{q}_0} \wedge \dots \wedge \eta_{\mathfrak{p}_k\mathfrak{q}_k}\,] = 1\,.$$

Now suppose that $T_{\mathrm{ran}}[\Sigma] \nvDash \varphi$. Since $T_{\mathrm{ran}}[\Sigma]$ is complete, we have $T_{\mathrm{ran}}[\Sigma] \vDash \neg\varphi$. By the first case, it follows that

$$\lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \varphi\,] = \lim_{n\to\infty} \big(1 - \Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \neg\varphi\,]\big) = 1 - 1 = 0\,. \qquad \square$$

Exercise 4.1. Prove that the theorem fails for signatures with 0-ary relations.

Lemma 4.13. *The Zero-One Law fails for signatures with functions.*

Proof. Let $\Sigma = \{f\}$ be a signature consisting just of a unary function symbol f, and define

$$\varphi := \forall x (fx \neq x)\,.$$

We have

$$\Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \varphi\,] = \frac{(n-1)^n}{n^n} = \big(1 - \tfrac{1}{n}\big)^n$$

which implies that

$$\lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\,\mathfrak{M} \vDash \varphi\,] = \lim_{n\to\infty} \big(1 - \tfrac{1}{n}\big)^n = \frac{1}{e}\,. \qquad \square$$

Lemma 4.14. *Let Σ be a finite relational signature. There exists no sentence $\varphi \in \mathrm{FO}[\Sigma]$ such that*

$$\mathfrak{M} \vDash \varphi \quad \text{iff} \quad |M| \text{ is even}\,, \qquad \text{for all finite } \Sigma\text{-structures } \mathfrak{M}\,.$$

Proof. $\lim_{n\to\infty} \Pr^n_{\mathfrak{M}}[\, \mathfrak{M} \vDash \varphi \,]$ does not exist in contradiction to the Zero-One Law. □

Remark. For every $n < \omega$, we can extend the Zero-One Law to the logic $\mathrm{FO}^{(n)}_{\infty\aleph_0}$ consisting of all $\mathrm{FO}_{\infty\aleph_0}$-formulae using at most n variables (both free and bound). Note that every FO(PFP)-formula can be translated to such a formula, for some suitable n. Hence, the Zero-One Law also holds for FO(LFP) and FO(PFP).

E5. *Indiscernible sequences*

1. *Ramsey Theory*

In this chapter we introduce some technical tools to study properties of sequences. This machinery is based on combinatorial results concerning colourings of linear orders.

Definition 1.1. (a) For a linear order I and a cardinal ν, we define

$$[I]^\nu := \{ \bar{\imath} \in I^\nu \mid \bar{\imath} \text{ is increasing} \}.$$

For an unordered set X we abuse notation by defining

$$[X]^\nu := \{ s \subseteq X \mid |s| = \nu \}.$$

(This is consistent with our convention of identifying sequences with their ranges.)

(b) Let $c : [A]^\nu \to \lambda$ be a function. A subset $C \subseteq A$ is *homogeneous* with respect to c if we have $c(\bar{a}) = c(\bar{a}')$, for all $\bar{a}, \bar{a}' \in [C]^\nu$.

(c) Let $\kappa, \lambda, \mu, \nu$ be cardinals. We write $\kappa \to (\mu)^\nu_\lambda$ if, for every set A of size $|A| \geq \kappa$ and each function $c : [A]^\nu \to \lambda$, there exists a homogeneous subset $C \subseteq A$ of size $|C| \geq \mu$.

Example. $6 \to (3)^2_2$ is equivalent to the statement that every undirected graph $\mathfrak{G} = \langle V, E \rangle$ with at least 6 elements contains a triangle or an independent set of size 3.

Exercise 1.1. Prove that $6 \to (3)^2_2$.

Let us start with the simplest case, that of unary colourings.

Theorem 1.2 (Pigeon Hole Principle). $\kappa \to (\kappa)^1_\lambda$, *for all infinite cardinals* κ *and every* $\lambda < \operatorname{cf} \kappa$.

Proof. Let A be a set of size $|A| = \kappa$ and suppose that $c : A \to \lambda$ is a function. We have to show that there is some $\alpha < \lambda$ with $|c^{-1}(\alpha)| = \kappa$. Suppose otherwise. Then $\lambda < \operatorname{cf} \kappa$ implies

$$|A| = \sum_{\alpha<\lambda} |c^{-1}(\alpha)| < \kappa .$$

A contradiction. □

The Theorem of Ramsey generalises the Pigeon Hole Principle to colourings of higher arities. We present two versions: one for infinite sets and one for finite sets.

Theorem 1.3 (Ramsey). $\aleph_0 \to (\aleph_0)^n_l$, *for all* $0 < n, l < \aleph_0$.

Proof. Let A be a set of size $|A| = \aleph_0$ and $c : [A]^n \to l$ a function. W.l.o.g. we may assume that $A = \omega$. By induction on n, we construct an infinite subset $C \subseteq \omega$ that is homogeneous with respect to c.

For $n = 1$ the claim follows from the Pigeon Hole Principle. Hence, we may assume that $n > 1$. In a first step, we define an infinite subset $B \subseteq \omega$ such that the value of $c(\bar b)$, for $\bar b \in [B]^n$, only depends on the minimal element b_0. For every $a \in \omega$, we define a function $c'_a : [\omega \setminus \{a\}]^{n-1} \to l$ by $c'_a(\bar b) := c(\bar b \cup \{a\})$. We construct an increasing sequence $a_0 < a_1 < \dots$ of elements and a decreasing sequence $A_0 \supseteq A_1 \supseteq \dots$ of subsets of ω as follows. We start with $a_0 := 0$ and $A_0 := \omega$. If a_i and A_i are already defined then we can use the inductive hypothesis to find an infinite subset $A_{i+1} \subseteq A_i \setminus \{a_0, \dots, a_i\}$. that is homogeneous with respect to c'_{a_i}. Let a_{i+1} be the minimal element of A_{i+1}.

Let $B := \{ a_i \mid i < \omega \}$ and set $k_i := c(a_i a_{i+1} \dots a_{i+n-1})$. Note that, for $i_0 < \dots < i_{n-1}$, we have $a_{i_1}, \dots, a_{i_{n-1}} \in A_{i_0+1}$. Hence, the above construction ensures that

$$\begin{aligned} c(a_{i_0} \dots a_{i_{n-1}}) &= c'_{a_{i_0}}(a_{i_1} \dots a_{i_{n-1}}) \\ &= c'_{a_{i_0}}(a_{i_0+1} \dots a_{i_0+n-1}) = c(a_{i_0} \dots a_{i_0+n-1}) = k_{i_0} . \end{aligned}$$

By the Pigeon Hole Principle, there exists an infinite subset $C \subseteq B$ such that $k_i = k_j$, for all $a_i, a_j \in C$. This set C is the desired homogeneous subset of ω. □

Example. Let $\langle P, \leq \rangle$ be an infinite partial order. We can use the Ramsey Theorem to prove that there exists an infinite set $C \subseteq P$ such that C is either linearly ordered or all elements of C are pairwise incomparable.

Let $c : [P]^2 \to 2$ be the function such that

$$c(\{a, b\}) := \begin{cases} 1 & \text{if } a \leq b \text{ or } b \leq a\,, \\ 0 & \text{otherwise}\,. \end{cases}$$

By the theorem there exists an infinite homogeneous set $C \subseteq P$. If we have $c(\{a, b\}) = 1$, for all $a, b \in C$, then C is a chain. Otherwise, all elements of C are pairwise incomparable.

The finite version of the Ramsey Theorem is as follows.

Theorem 1.4 (Ramsey). *For all $l, m, n < \aleph_0$, there exists a finite cardinal $k < \aleph_0$ such that $k \to (m)^n_l$.*

Proof. For a contradiction, suppose that there exists no finite k with $k \to (m)^n_l$. Let F_k be the set of all functions $c : [k]^n \to l$ such that there is no subset $C \subseteq [k]$ of size $|C| \geq m$ that is homogeneous with respect to c. It follows that each set F_k is finite and nonempty. Furthermore, $c \in F_{k+1}$ implies that $c \upharpoonright [k]^n \in F_k$. Hence, if we order the set $T := \bigcup_k F_k$ by inclusion then we obtain a tree $\langle T, \subseteq \rangle$. This tree is infinite and finitely branching. By the Lemma of Kőnig it therefore contains an infinite branch $(c_k)_{k<\omega}$ with $c_k \in F_k$. Set $c := \bigcup_k c_k$. Then c is a function $c : [\aleph_0]^n \to l$. By the infinite version of the Ramsey Theorem, there exists an infinite subset $C \subseteq \aleph_0$ that is homogeneous with respect to c. Fix a subset $Z \subseteq C$ of size $|Z| = m$ and let k be the maximal element of Z. It follows that Z is homogeneous with respect to c_{k+1}. A contradiction. □

Next, we consider the case of infinitely many colours and uncountable homogeneous sets. We start with a counterexample.

Lemma 1.5. $2^{\aleph_0} \nrightarrow (3)^2_{\aleph_0}$

Proof. Let $c : [2^{\aleph_0}]^2 \to \aleph_0$ be the function mapping a pair $\{f, g\}$ of distinct functions $f, g : \aleph_0 \to 2$ to the least number n with $f(n) \neq g(n)$. If $\{f, g, h\}$ were homogeneous with respect to c, we would have $f(n) \neq g(n)$, $f(n) \neq h(n)$, and $g(n) \neq h(h)$, for some n. Since $f(n), g(n), h(n) \in \{0, 1\}$ this is impossible. □

Theorem 1.6 (Erdős-Rado). *For all cardinals $\kappa \geq \aleph_0$ and $n < \aleph_0$,*

$$\beth_n(\kappa)^+ \to (\kappa^+)^{n+1}_\kappa .$$

Proof. We prove the claim by induction on n. By the Pigeon Hole Principle, we have $\kappa^+ \to (\kappa^+)^1_\kappa$. Hence, the claim holds for $n = 1$. For the inductive step, suppose we have already proved the theorem for n. Set $\lambda := \beth_{n+1}(\kappa)$ and $\mu := \beth_n(\kappa)$, and let $c : [\lambda^+]^{n+1} \to \kappa$ be a colouring.

As a first step we define an increasing sequence of ordinals $\beta_i < \lambda^+$, for $i < \kappa^+$, with the following property:

(∗) For every set $S \subseteq \beta_i$ of size $|S| \leq \mu$ and all ordinals $\gamma < \lambda^+$, there exists some ordinal $\eta < \beta_{i+1}$ such that

$$\eta \in S \quad \text{iff} \quad \gamma \in S ,$$
$$\text{and} \quad c(\bar{\alpha}\eta) = c(\bar{\alpha}\gamma) , \quad \text{for all } \bar{\alpha} \in S^n .$$

The ordinals β_i will be used as a measuring stick in the construction below. We define β_i by induction on i. Let $\beta_0 := 0$ and set $\beta_\delta := \sup_{i<\delta} \beta_i$, for limit ordinals δ. For the inductive step, we set

$$\beta_{i+1} := \sup \{ \eta(S, \gamma) \mid \gamma < \lambda^+, \ S \subseteq \beta_i \text{ with } |S| \leq \mu \} ,$$

where $\eta(S, \gamma)$ denotes the minimal ordinal η such that

$$\eta \in S \quad \text{iff} \quad \gamma \in S ,$$
$$\text{and} \quad c(\bar{\alpha}\eta) = c(\bar{\alpha}\gamma) , \quad \text{for all } \bar{\alpha} \in S^n .$$

Note that there are at most $|\beta_i|^\mu = \lambda^\mu = (2^\mu)^\mu = \lambda$ subsets of β_i of size $|S| \leq \mu$ and there are at most $\kappa^\mu = 2^\mu = \lambda$ functions $S \to \kappa$. Consequently, the supremum above is taken over a set of at most $\lambda \otimes \lambda = \lambda$ ordinals each of which is less than λ^+. Since λ^+ is regular it follows that the supremum β_{i+1} is less than λ^+.

Having defined the β_i we set $\beta^* := \sup_{i<\mu^+} \beta_i$ and we define ordinals $\alpha_i < \beta_{i+1}$, for $i < \mu^+$, such that $\alpha_i \neq \alpha_k$, for $i \neq k$, and

$$c(\alpha_{k_0}, \dots, \alpha_{k_{n-1}}, \alpha_i) = c(\alpha_{k_0}, \dots, \alpha_{k_{n-1}}, \beta^*),$$

for all $k_0, \dots, k_{n-1} < i$. We can find α_i by induction on i using property $(*)$ with $S = \{ \alpha_k \mid k < i \}$ and $\gamma := \beta^*$.

Define a colouring $c' : [\mu^+]^n \to \kappa$ by

$$c'(\bar{\imath}) := c(\alpha_{i_0} \dots \alpha_{i_{n-1}} \beta^*).$$

By inductive hypothesis, there exists a set $I \subseteq \mu^+$ of size $|I| \geq \kappa^+$ such that

$$c'(\bar{\imath}) = c'(\bar{k}), \quad \text{for all } \bar{\imath}, \bar{k} \in [I]^n.$$

Let $J := \{ \alpha_i \mid i \in I \}$. For $\bar{\gamma}, \bar{\eta} \in [J]^{n+1}$ it follows that

$$\begin{aligned} c(\gamma_0 \dots \gamma_{n-1}\gamma_n) &= c(\gamma_0 \dots \gamma_{n-1}\beta^*) \\ &= c(\eta_0 \dots \eta_{n-1}\beta^*) = c(\eta_0 \dots \eta_{n-1}\eta_n). \end{aligned}$$

Hence, J is the desired homogeneous subset of λ^+. □

2. *Ramsey Theory for trees*

So far, we have considered homogeneous subsets of linear orders. A special property of linear orders is that every subset again induces a linear order. When considering colourings of other structures this is no longer the case. In this section we prove variants of the Pigeon Hole Principle and the Theorem of Ramsey for trees where the homogeneous sets we obtain again induce trees. There are two kinds of tree structures we will be working with: trees of the form $\mathfrak{T}_*(\kappa^{<\alpha})$ are equipped with the tree-order $\leq$ and relations $\prec_p$ for the direction of the immediate successors, while trees $\mathfrak{T}_n(\kappa^{<\alpha})$ also have functions pf to compare the levels of elements.

Definition 2.1. Let κ be a cardinal and α an ordinal.

(a) We denote the tree order on $\kappa^{<\alpha}$ by $\leq$ and $\sqcap$ is the infimum operation with respect to $\leq$. For $\eta, \zeta \in \kappa^{<\alpha}$ and $p \in \kappa$, we further set

$$\eta \prec_p \zeta \quad :\text{iff} \quad \eta p \leq \zeta\,.$$

For $|\eta| \leq |\zeta|$, we denote by $\mathrm{pf}(\eta,\zeta)$ the prefix of ζ of length $|\eta|$. If $|\eta| > |\zeta|$, we set $\mathrm{pf}(\eta,\zeta) := \zeta$.

(b) We define

$$\mathfrak{T}_*(\kappa^{<\alpha}) := \langle \kappa^{<\alpha}, \sqcap, \leq, (\prec_p)_{p\in\kappa} \rangle\,,$$

$$\text{and} \quad \mathfrak{T}_n(\kappa^{<\alpha}) := \langle \kappa^{<\alpha}, \sqcap, \leq, (\prec_p)_{p\in\kappa}, \mathrm{pf}, (\eta)_{\eta\in\kappa^{<n}} \rangle\,, \quad \text{for } n \leq \alpha\,.$$

We denote the substructure of $\mathfrak{T}_n(\kappa^{<\alpha})$ generated by a set $X \subseteq \kappa^{<\alpha}$ by $\langle\!\langle X \rangle\!\rangle_n$.

Remark. (a) Note that the substructure $\langle\!\langle X \rangle\!\rangle_n$ generated by a set $X \subseteq \kappa^{<\alpha}$ has universe

$$\langle\!\langle X \rangle\!\rangle_n = \kappa^{<n} \cup \{ \mathrm{pf}(\xi \sqcap \eta, \zeta) \mid \xi, \eta, \zeta \in X \}\,.$$

Thus, it consists of (i) all elements of $X \cup \kappa^{<n}$, (ii) all elements of the form $\eta \sqcap \zeta$, with $\eta, \zeta \in X$, and (iii) all prefixes of some element of X that have the same length as an element of the form (i) or (ii).

(b) Note that we have

$$|\eta| = |\zeta| \quad \text{iff} \quad \mathrm{pf}(\eta,\zeta) = \zeta \text{ and } \mathrm{pf}(\zeta,\eta) = \eta\,.$$

Hence, every embedding $h : \mathfrak{T}_n(\kappa^{<\alpha}) \to \mathfrak{T}_n(\kappa^{<\alpha})$ has the property that

$$|\eta| = |\zeta| \quad \text{implies} \quad h(|\eta|) = h(|\zeta|)\,, \quad \text{for all } \eta, \zeta \in \kappa^{<\alpha}.$$

Definition 2.2. (a) The set of *levels* of a tuple $\bar{\eta} \in (\kappa^{<\alpha})^d$ is

$$\mathrm{Lvl}(\bar{\eta}) := \{\, |\eta_i \sqcap \eta_j| \mid i, j < d \,\} = \{\, |\zeta| \mid \zeta \in \langle\!\langle \bar{\eta} \rangle\!\rangle_0 \,\}.$$

(b) Let $h : \mathfrak{T}_n(\kappa^{<\alpha}) \to \mathfrak{T}_n(\kappa^{<\alpha})$ be an embedding. The *level embedding function* associated with h is the function $f : \alpha \to \alpha$ such that

$$|h(\eta)| = f(|\eta|), \quad \text{for all } \eta \in \kappa^{<\alpha}.$$

Our first result is a generalisation of a strong version of the Pigeon Hole Principle. We omit the proof, which is quite involved.

Theorem 2.3 (Halpern, Läuchli)**.** *Let $m, d < \omega$ and let C be a finite set. For every function $c : (m^{<\omega})^d \to C$ there exist embeddings*

$$g_i : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega}), \quad \text{for } i < d,$$

such that all g_i have the same level embedding function and

$$c(g_0(\eta_0), \dots, g_{d-1}(\eta_{d-1})) = c(g_0(\zeta_0), \dots, g_{d-1}(\zeta_{d-1})),$$

for all tuples $\bar{\eta}, \bar{\zeta} \in (m^{<\omega})^d$ with $|\eta_0| = \dots = |\eta_{d-1}|$ and $|\zeta_0| = \dots = |\zeta_{d-1}|$.

In the remainder of this section we generalise the Theorem of Ramsey to trees. In the version for linear orders we required tuples to have the same colour if they have the same order type. When dealing with other kinds of structures we replace the order type of a tuple by its atomic type.

Definition 2.4. (a) Let $c : A^d \to C$ a function, for $d < \omega$, and let $\approx$ be an equivalence relation on A^d. A subset $X \subseteq A$ is *$\approx$-homogeneous* with respect to c if

$$\bar{\eta} \approx \bar{\zeta} \quad \text{implies} \quad c(\bar{\eta}) = c(\bar{\zeta}), \quad \text{for all } \bar{\eta}, \bar{\zeta} \in X^d.$$

(b) For tuples $\bar{\eta}, \bar{\zeta} \subseteq \kappa^{<\alpha}$, we define

$$\begin{aligned} \bar{\eta} \approx_* \bar{\zeta} \quad &:\text{iff} \quad \mathrm{atp}(\bar{\eta}/\mathfrak{T}_*(\kappa^{<\alpha})) = \mathrm{atp}(\bar{\zeta}/\mathfrak{T}_*(\kappa^{<\alpha})), \\ \bar{\eta} \approx_n \bar{\zeta} \quad &:\text{iff} \quad \mathrm{atp}(\bar{\eta}/\mathfrak{T}_n(\kappa^{<\alpha})) = \mathrm{atp}(\bar{\zeta}/\mathfrak{T}_n(\kappa^{<\alpha})). \end{aligned}$$

Our goal is to prove the following variant of the Theorem of Ramsey for trees.

Theorem 2.5 (Milliken). *Let $m, d < \omega$ and let C be a finite set. For every function $c : (m^{<\omega})^d \to C$ there exists an embedding $g : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega})$ such that* $\operatorname{rng} g$ *is $\approx_0$-homogeneous with respect to c.*

The proof of the Theorem of Ramsey was by induction on the length of tuples. We prove the Theorem of Milliken by a similar argument where the induction is on the number of levels of a tuple. The next lemma contains the inductive step of this argument. It is based on the following variant of the relation $\approx_n$.

Definition 2.6. Let $k, n < \omega$. For $\bar\eta, \bar\zeta \subseteq m^{<\omega}$, we set

$$\bar\eta \approx_{n,k} \bar\zeta \quad :\text{iff} \quad \bar\eta = \bar\zeta\,, \text{ or}$$
$$\bar\eta \approx_n \bar\zeta \text{ and } |\mathrm{Lvl}(\bar\eta) \setminus [n]|\,,\ |\mathrm{Lvl}(\bar\zeta) \setminus [n]| \leq k\,,$$

and we denote by $\approx_{\omega,k}$ the transitive closure of the union $\bigcup_{n<\omega} \approx_{n,k}$.

Remark. (a) Note that

$$\bar\eta \approx_{n,0} \bar\zeta \qquad \text{iff} \quad \bar\eta = \bar\zeta\,,$$

and the fact that $|\mathrm{Lvl}(\bar\eta)| \leq 2|\bar\eta|$ implies that

$$\bar\eta \approx_{n,2|\bar\eta|} \bar\zeta \quad \text{iff} \quad \bar\eta \approx_n \bar\zeta\,.$$

(b) A set X is $\approx_{\omega,k}$-homogeneous if, and only if, it is $\approx_{n,k}$-homogeneous, for every $n < \omega$.

Lemma 2.7. *Let $m, d < \omega$, let C be a finite set, and let $c : (m^{<\omega})^d \to C$ be a function such that $m^{<\omega}$ is $\approx_{\omega,k}$-homogeneous with respect to c. For every $n < \omega$, there exists an embedding*

$$g : \mathfrak{T}_{n+1}(m^{<\omega}) \to \mathfrak{T}_{n+1}(m^{<\omega})$$

such that $\operatorname{rng} g$ *is $\approx_{n,k+1}$-homogeneous with respect to c.*

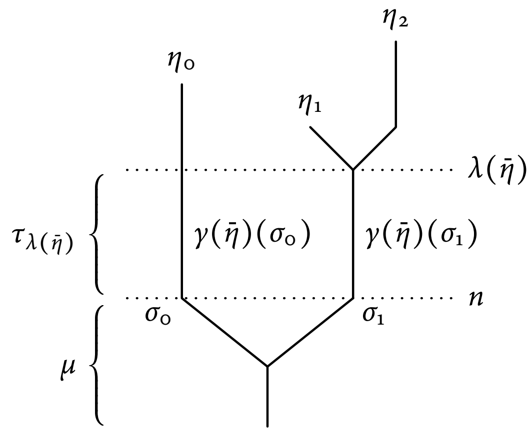

Figure 1.. The definition of μ, λ, τ_l, and γ.

Proof. Given $n < \omega$, set

$$\Gamma := \left\{ \bar\eta \in (m^{<\omega})^d \setminus (m^{<n})^d \;\middle|\; |\mathrm{Lvl}(\bar\eta) \setminus [n]| \leq k + 1 \right\}.$$

For $\bar\eta \in \Gamma$, let

$$\lambda(\bar\eta) := \min(\mathrm{Lvl}(\bar\eta) \setminus [n]).$$

Set $L := m^n$ and let

$$\mu : m^{<\omega} \setminus m^{<n} \to L : \eta \mapsto \eta \restriction n$$

be the function mapping each element to its prefix of length n. For $l \geq n$, let $\tau_l : m^{<\omega} \setminus m^{<l} \to m^{l-n}$ be the function mapping an element $\eta \in m^{<\omega}$ of length $|\eta| \geq l$ to the unique sequence $\sigma \in m^{<\omega}$ such that

$$|\sigma| = l - n \quad \text{and} \quad \mu(\eta)\sigma \leq \eta.$$

Let H be the set of all functions $h : L \to m^{<\omega}$ such that

$$|h(\rho)| = |h(\sigma)|, \quad \text{for all } \rho, \sigma \in L.$$

For $h, h' \in H$ and $\bar\eta \in \Gamma$, we set

$$h \sim_{\bar\eta} h' \quad :\text{iff} \quad h(\mu(\eta_i)) = h'(\mu(\eta_i)), \quad \text{for all } i < d \text{ with } |\eta_i| \geq n.$$

We define a function $\gamma : \Gamma \to H : \bar\eta \mapsto h_{\bar\eta}$ where

$$h_{\bar\eta}(\sigma) := \begin{cases} \tau_{\lambda(\bar\eta)}(\eta_i) & \text{if } \eta_i \in \mu^{-1}(\sigma), \\ \langle 0, \dots, 0 \rangle & \text{otherwise}. \end{cases}$$

Note that, in the first case of the definition of $h_{\bar\eta}(\sigma)$, the value does not depend on the choice of $i < d$ since

$$\eta_i, \eta_j \in \mu^{-1}(\sigma) \quad \text{implies} \quad \tau_{\lambda(\bar\eta)}(\eta_i) = \tau_{\lambda(\bar\eta)}(\eta_j).$$

Finally, we define a function $\beta : H \times \Gamma/{\approx_n} \to C$ by

$$\beta(h, [\bar\eta]_{\approx_n}) := c(\bar a[\bar\zeta]) \quad \text{where } \bar\zeta \in \gamma^{-1}\big[[h]_{\sim_{\bar\eta}}\big] \cap [\bar\eta]_{\approx_n}.$$

To prove that β is well-defined, we have to check that

$$\gamma^{-1}\big[[h]_{\sim_{\bar\eta}}\big] \cap [\bar\eta]_{\approx_n} \neq \varnothing$$

and that the value of β does not depend on the choice of $\bar\zeta$.

For non-emptiness, fix h and $[\bar\eta]_{\approx_n}$. For $i < d$ with $|\eta_i| \geq n$, let $\rho_i \in m^{<\omega}$ be the sequence such that

$$\eta_i = \mu(\eta_i)\, \tau_{\lambda(\bar\eta)}(\eta_i)\, \rho_i.$$

We set

$$\zeta_i := \mu(\eta_i)\, h(\mu(\eta_i))\, \rho_i.$$

For $i < d$ with $|\eta_i| < n$, we set $\zeta_i := \eta_i$. Then $\bar\zeta \approx_n \bar\eta$ and, since we have

$$\lambda(\bar\zeta) = n + |h(\mu(\eta_i))|, \quad \text{for any } i < d \text{ with } |\eta_i| \geq n,$$

it also follows that $\gamma(\bar{\zeta}) \sim_{\bar{\eta}} h$. Hence, $\bar{\zeta} \in \gamma^{-1}[[h]_{\sim_{\bar{\eta}}}] \cap [\bar{\eta}]_{\approx_n}$.

To show that the value of $\beta(h, [\bar{\eta}]_{\approx_n})$ does not depend on the choice of $\bar{\zeta}$, consider two tuples $\bar{\xi}, \bar{\zeta} \in \gamma^{-1}[[h]_{\sim_{\bar{\eta}}}] \cap [\bar{\eta}]_{\approx_n}$. First of all, note that $\bar{\xi} \approx_n \bar{\zeta}$ implies that $\mu(\xi_i) = \mu(\zeta_i)$, for all i with $|\xi_i| \geq n$, since

$$\sigma \prec_p \xi_i \quad \text{iff} \quad \sigma \prec_p \zeta_i\,, \quad \text{for all } \sigma \in m^{n-1} \text{ and all } p < m\,.$$

(For $n = 0$, we have $\mu(\xi_i) = \langle\rangle = \mu(\zeta_i)$, for all i.) Consequently, $\gamma(\bar{\xi}) \sim_{\bar{\eta}} h \sim_{\bar{\eta}} \gamma(\bar{\zeta})$ implies that

$$\tau_{\lambda(\bar{\xi})}(\xi_i) = h(\mu(\xi_i)) = h(\mu(\zeta_i)) = \tau_{\lambda(\bar{\zeta})}(\zeta_i)\,,$$

for all $i < d$ with $|\xi_i| \geq n$. In particular, $\lambda(\bar{\xi}) = \lambda(\bar{\zeta}) =: l$ and

$$\xi_i \upharpoonright l = \mu(\xi_i)\,\tau_l(\xi_i) = \mu(\zeta_i)\,\tau_l(\zeta_i) = \zeta_i \upharpoonright l\,.$$

As $\bar{\xi} \approx_n \bar{\zeta}$ it follows that $\bar{\xi} \approx_{l+1} \bar{\zeta}$. Since

$$|\mathrm{Lvl}(\bar{\xi}) \smallsetminus [l+1]| = |\mathrm{Lvl}(\bar{\zeta}) \smallsetminus [l+1]| \leq k\,,$$

we, therefore, have $\bar{\xi} \approx_{l+1,k} \bar{\zeta}$ and, by assumption on c, it follows that $c(\bar{\xi}) = c(\bar{\zeta})$, as desired.

To conclude the proof, consider the function $c_0 : H \to C^{\Gamma/\approx_n}$ mapping a tuple $h \in H$ to the function $[\bar{\eta}]_{\approx_n} \mapsto \beta(h, [\bar{\eta}]_{\approx_n})$, and let $c_1 : (m^{<\omega})^L \to C^{\Gamma/\approx_n}$ be an arbitrary extension of c_0.

Since $C^{\Gamma/\approx_n}$ is a finite set, we can use the Theorem of Halpern and Läuchli to obtain embeddings $g_\sigma : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega})$, for $\sigma \in L$, such that all g_σ have the same level embedding function and the restriction $c_1 \upharpoonright H \cap \prod_{\sigma \in L} \operatorname{rng} g_\sigma$ is constant. We can define the desired embedding $g : \mathfrak{T}_{n+1}(m^{<\omega}) \to \mathfrak{T}_{n+1}(m^{<\omega})$ by setting

$$g(\eta) := \begin{cases} \eta & \text{if } |\eta| \leq n\,, \\ \sigma g_\sigma(\xi) & \text{if } \eta = \sigma\xi \text{ for } \sigma \in L \text{ and } \xi \in m^{<\omega}\,. \end{cases}$$

It remains to prove that $\operatorname{rng} g$ is $\approx_{n,k+1}$-homogeneous with respect to c. Let $\bar{\eta}, \bar{\zeta} \in \Gamma \cap (\operatorname{rng} g)^d$ be tuples with $\bar{\eta} \approx_n \bar{\zeta}$. To show that $c(\bar{\eta}) = c(\bar{\zeta})$, set $h := \gamma(\bar\eta)$ and $h' := \gamma(\bar\zeta)$. For each $\sigma \in L$, fix some $\xi_\sigma \in \operatorname{rng} g_\sigma$ and set

$$h_0(\sigma) := \begin{cases} h(\sigma) & \text{if } \sigma \preceq \eta_i \text{ for some } i\,, \\ \xi_\sigma & \text{otherwise}\,. \end{cases}$$

Then $h_0 \in \prod_{\sigma \in L} \operatorname{rng} g_\sigma$ and $h_0 \sim_{\bar\eta} h$. Similarly, we can find some $h'_0 \in \prod_{\sigma \in L} \operatorname{rng} g_\sigma$ with $h'_0 \sim_{\bar\zeta} h'$. Since $c_0(h_0) = c_0(h'_0)$ and $[\bar\eta]_{\approx_n} = [\bar\zeta]_{\approx_n}$ it follows that

$$\begin{aligned} c(\bar\eta) = \beta(h, [\bar\eta]_{\approx_n}) &= \beta(h_0, [\bar\eta]_{\approx_n}) \\ &= c_0(h_0)([\bar\eta]_{\approx_n}) \\ &= c_0(h'_0)([\bar\eta]_{\approx_n}) \\ &= c_0(h'_0)([\bar\zeta]_{\approx_n}) \\ &= \beta(h'_0, [\bar\zeta]_{\approx_n}) = \beta(h', [\bar\zeta]_{\approx_n}) = c(\bar\zeta)\,. \end{aligned}$$

□

Lemma 2.8. *Let $m, d < \omega$, let C be a finite set, and let $c : (m^{<\omega})^d \to C$ be a function such that $m^{<\omega}$ is $\approx_{\omega,k}$-homogeneous with respect to c. There exists an embedding $g : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega})$ such that* $\operatorname{rng} g$ *is $\approx_{\omega,k+1}$-homogeneous with respect to c.*

Proof. To simplify notation, we write $c \circ g$ for the function mapping a tuple $\bar\eta \in (m^{<\omega})^d$ to the value $c(g(\eta_0), \dots, g(\eta_{d-1}))$. We construct a sequence of embeddings

$$g_n : \mathfrak{T}_n(m^{<\omega}) \to \mathfrak{T}_n(m^{<\omega})\,, \quad \text{for } n < \omega\,,$$

such that, for all $i < n < \omega$, the set $m^{<\omega}$ is $\approx_{i,k+1}$-homogeneous with respect to the function $c_n := c \circ g_0 \circ \dots\dots g_n$.

We start with $g_0 :=$ id. Then $c_0 = c$ trivially satisfies the above condition. For the inductive step, suppose that we have already found functions $g_0, \dots, g_n$ such that, for every $i < n$, $m^{<\omega}$ is $\approx_{i,k+1}$-homogeneous with respect to c_n. We can use Lemma 2.7 to find an embedding g_{n+1} :

$\mathfrak{T}_{n+1}(m^{<\omega}) \to \mathfrak{T}_{n+1}(m^{<\omega})$ such that $m^{<\omega}$ is $\approx_{n,k+1}$-homogeneous with respect to $c_n \circ g_{n+1} = c_{n+1}$. Furthermore, since $m^{<\omega}$ is $\approx_{i,k+1}$-homogeneous with respect to c_n, for all $i < n$, it follows that it is also $\approx_{i,k+1}$-homogeneous with respect to $c_n \circ g_{n+1}$.

Having constructed the sequence $g_0, g_1, \dots$ we obtain the desired embedding $g : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega})$ as follows. For $\eta \in m^n$, we set $g(\eta) := (g_0 \circ \dots \circ g_{n+1})(\eta)$. Clearly, g is an embedding. Hence, it remains to prove that $\operatorname{rng} g$ is $\approx_{\omega,k+1}$-homogeneous. Fix n and consider two tuples $\bar\eta, \bar\zeta \subseteq m^{<\omega}$ such that

$$\bar\eta \approx_n \bar\zeta \quad \text{and} \quad \left|\operatorname{Lvl}(\bar\eta) \smallsetminus [n]\right|, \left|\operatorname{Lvl}(\bar\zeta) \smallsetminus [n]\right| \leq k+1 .$$

Choose $n < l < \omega$ such that $\bar\eta, \bar\zeta \subseteq m^{<l}$. Then

$$g(\bar\eta) = (g_0 \circ \dots \circ g_l)(\bar\eta) \quad \text{and} \quad g(\bar\zeta) = (g_0 \circ \dots \circ g_l)(\bar\zeta) .$$

As $\operatorname{rng}(g_0 \circ \dots \circ g_l)$ is $\approx_{n,k+1}$-homogeneous with respect to c, it follows that $c(g(\bar\eta)) = c(g(\bar\zeta))$. □

Proof of Theorem 2.5. Note that, for every $n < \omega$, the set $m^{<\omega}$ is $\approx_{n,0}$-homogeneous with respect to c. Hence, repeating Lemma 2.8 we obtain embeddings

$$g_k : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega}), \quad \text{for } k \leq 2d ,$$

such that $\operatorname{rng}(g_0 \circ \dots \circ g_k)$ is $\approx_{\omega,k}$-homogeneous with respect to c. Setting $g := g_0 \circ \dots \circ g_{2d}$ it follows that $\operatorname{rng} g$ is $\approx_{0,2d}$-homogeneous with respect to c. Since $|\operatorname{Lvl}(\bar\eta)| \leq 2d$, for all $\bar\eta \in (m^{<\omega})^d$, this is the same as saying that $\operatorname{rng} g$ is $\approx_0$-homogeneous with respect to c. □

As for the Theorem of Ramsey, the Theorem of Milliken also has a finitary version. The proof follows exactly the same lines as that of Theorem 1.4.

Theorem 2.9. *Let $m, d, k < \omega$ and let C be a finite set. There exists a number $n < \omega$ such that, for every function $c : (m^{<n})^d \to C$, there exists an embedding $g : \mathfrak{T}_0(m^{<k}) \to \mathfrak{T}_0(m^{<n})$ such that* $\operatorname{rng} g$ *is $\approx_0$-homogeneous with respect to c.*

Proof. For a contradiction, suppose that there exists no number n as above. For $n < \omega$, let F_n be the set of all functions $c : (m^{<n})^d \to C$ such that there is no embedding $g : \mathfrak{T}_0(m^{<k}) \to \mathfrak{T}_0(m^{<n})$ such that $\operatorname{rng} g$ is $\approx_0$-homogeneous with respect to c. Each set F_n is finite and nonempty. Furthermore, $c \in F_{n+1}$ implies that $c \upharpoonright (m^{<n})^d \in F_n$. Hence, if we order the set $T := \bigcup_n F_n$ by inclusion, we obtain a tree $\langle T, \subseteq \rangle$. This tree is infinite and finitely branching. By the Lemma of Kőnig it therefore contains an infinite branch $(c_n)_{n<\omega}$ where $c_n \in F_n$. Set $c := \bigcup_n c_n$. Then c is a function $c : (m^{<\omega})^d \to C$. By Theorem 2.5, there exists an embedding $g : \mathfrak{T}_0(m^{<\omega}) \to \mathfrak{T}_0(m^{<\omega})$ such that $\operatorname{rng} g$ is $\approx_0$-homogeneous with respect to c. Fix a number $n < \omega$ such that $\operatorname{rng}(g \upharpoonright m^{<k}) \subseteq m^{<n}$. Then $g \upharpoonright m^{<k} : \mathfrak{T}_0(m^{<k}) \to \mathfrak{T}_0(m^{<n})$ is an embedding such that $\operatorname{rng} g$ is $\approx_0$-homogeneous with respect to c_n. A contradiction. □

Note that every $\approx_*$-homogeneous set is also $\approx_0$-homogeneous. Hence, we would obtain a stronger version of the Theorem of Milliken if we could replace the relation $\approx_0$ by $\approx_*$. For the finitary version this is possible.

Theorem 2.10. *Let $m, d, k < \omega$ and let C be a finite set. There exists a number $n < \omega$ such that, for every function $c : (m^{<n})^d \to C$, there exists an embedding $g : \mathfrak{T}_*(m^{<k}) \to \mathfrak{T}_*(m^{<n})$ such that* $\operatorname{rng} g$ *is $\approx_*$-homogeneous with respect to c.*

The proof consists in finding sets where the relations $\approx_*$ and $\approx_0$ coincide. To do so we introduce the following family of embeddings.

Definition 2.11. For $0 < k < \omega$, the *k-th skew embedding*

$$h_k : \mathfrak{T}_*(m^{<k}) \to \mathfrak{T}_*(m^{<l(k)})$$

is defined inductively as follows. We start with $h_1 : \langle\rangle \mapsto \langle\rangle$ and $l(1) = 1$. If h_k and $l(k)$ are already defined, we set

$$h_{k+1}(\langle\rangle) := \langle\rangle \quad \text{and} \quad h_{k+1}(p\eta) := \langle \underbrace{p, \dots, p}_{p+2+pl(k) \text{ times}} \rangle h_k(\eta),$$

for $\eta \in m^{<\omega}$ and $p < m$. Furthermore, $l(k+1) := ml(k) + m + 1$.

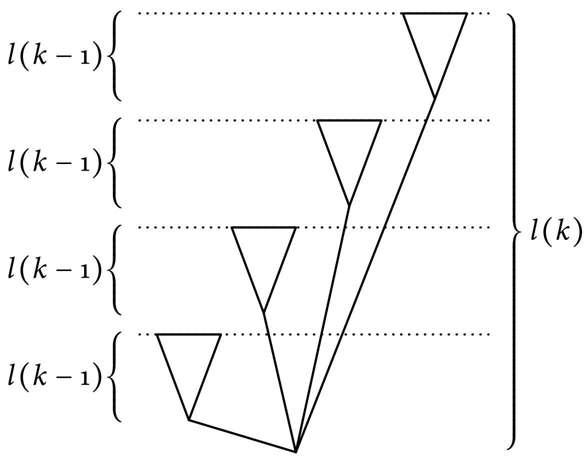

Figure 2.. The k-th skew embedding h_k.

Lemma 2.12. *The k-th skew embedding $h_k : \mathfrak{T}_*(m^{<k}) \to \mathfrak{T}_*(m^{<l(k)})$ is an embedding.*

Proof. By an easy induction on $|\eta|$, one can show that

$$\begin{aligned} \eta \leq \zeta \quad &\text{implies} \quad h_k(\eta) \leq h_k(\zeta)\,, \\ \text{and} \quad \eta \prec_p \zeta \quad &\text{implies} \quad h_k(\eta) \prec_p h_k(\zeta)\,. \end{aligned}$$

Similarly, an induction on $|\eta \sqcap \zeta|$ yields

$$h_k(\eta \sqcap \zeta) = h_k(\eta) \sqcap h_k(\zeta)\,. \qquad \square$$

A useful property of a skew embedding is that it upgrades $\approx_*$-equivalence to $\approx_0$-equivalence.

Lemma 2.13. *Let $\bar{\eta}, \bar{\zeta} \subseteq m^{<k}$. Then $\bar{\eta} \approx_* \bar{\zeta}$ implies $h_k(\bar{\eta}) \approx_0 h_k(\bar{\zeta})$.*

Proof. Let $\bar{\eta}, \bar{\zeta} \in (m^{<k})^d$ with $\bar{\eta} \approx_* \bar{\zeta}$. We start by proving the following

claims:

(a) $h_k(\bar{\eta}) \approx_* h_k(\bar{\zeta})$.

(b) $|h_k(\eta_i)| < |h_k(\eta_j)|$ iff $|h_k(\zeta_i)| < |h_k(\zeta_j)|$, for all $i, j < d$.

(c) $\mathrm{pf}(h_k(\eta_i), h_k(\eta_j)) \prec_p h_k(\eta_j)$
iff $\mathrm{pf}(h_k(\zeta_i), h_k(\zeta_j)) \prec_p h_k(\zeta_j)$, for all $i, j < d$.

(a) Since $h_k : \mathfrak{T}_*(m^{<k}) \to \mathfrak{T}_*(m^{<l(k)})$ is an embedding, it preserves atomic types. Consequently, we have $h_k(\bar{\eta}) \approx_* \bar{\eta} \approx_* \bar{\zeta} \approx_* h_k(\bar{\zeta})$.

(b) It follows by induction on $|\eta_i \sqcap \eta_j|$ that

$$|h_k(\eta_i)| < |h_k(\eta_j)| \quad \text{iff} \quad \eta_i <_{\mathrm{lex}} \eta_j .$$

Hence, $\bar{\eta} \approx_* \bar{\zeta}$ implies that

$$\begin{aligned} |h_k(\eta_i)| < |h_k(\eta_j)| \quad &\text{iff} \quad \eta_i <_{\mathrm{lex}} \eta_j \\ &\text{iff} \quad \zeta_i <_{\mathrm{lex}} \zeta_j \quad \text{iff} \quad |h_k(\zeta_i)| < |h_k(\zeta_j)| . \end{aligned}$$

(c) By definition of h_k, we have

$$\mathrm{pf}(h_k(\eta_i), h_k(\eta_j)) \prec_p h_k(\eta_j)$$

iff $|h_k(\eta_i)| < |h_k(\eta_j)|$ and $h_k(\eta_i \sqcap \eta_j) \prec_p h_k(\eta_j)$.

Therefore, (c) follows from (a) and (b).

To conclude the proof, suppose that $\bar{\eta} \approx_* \bar{\zeta}$. W.l.o.g. we may assume that, for all $i, j < d$, there is some $l < d$ such that $\eta_l = \eta_i \sqcap \eta_j$. Then it follows by (a), (b), and (c) that $h_k(\bar{\eta}) \approx_0 h_k(\bar{\zeta})$. □

Proof of Theorem 2.10. Let $h_k : m^{<k} \to m^{<l(k)}$ be the k-th skew embedding. By Theorem 2.9, there exists a number n such that, for every function $c : (m^{<n})^d \to C$, we can find an embedding $g : \mathfrak{T}_0(m^{<l(k)}) \to \mathfrak{T}_0(m^{<n})$ such that $\mathrm{rng}\, g$ is $\approx_0$-homogeneous with respect to c. We

claim that $g \circ h_k : \mathfrak{T}_*(m^{<k}) \to \mathfrak{T}_*(m^{<n})$ is the desired embedding. For $\bar\eta, \bar\zeta \in (m^{<k})^d$ it follows by Lemma 2.13 that

$$\begin{aligned} \bar\eta \approx_* \bar\zeta \quad &\Rightarrow \quad h_k(\bar\eta) \approx_\circ h_k(\bar\zeta) \\ &\Rightarrow \quad g(h_k(\bar\eta)) \approx_\circ g(h_k(\bar\zeta)) \\ &\Rightarrow \quad c(g(h_k(\bar\eta))) = c(g(h_k(\bar\zeta)))\,. \end{aligned}$$

Hence, $\mathrm{rng}(g \circ h_k)$ is $\approx_*$-homogeneous with respect to c. □

3. *Indiscernible sequences*

If we apply the Ramsey Theorem to sequences of elements in a structure coloured by their types we obtain subsequences where each tuple has the same type. Such sequences, called *indiscernible,* can be used to investigate the structure of the given model. Let us fix some notation.

Definition 3.1. Let $\langle I, \leq \rangle$ be a linear order and $(\bar a^i)_{i \in I}$ a sequence of tuples $\bar a^i \in A^\alpha$, for some ordinal α.

(a) For $\bar\imath \in I^n$, we set $\bar a[\bar\imath] := \bar a^{i_0} \dots \bar a^{i_{n-1}}$.

(b) The *order type* of a tuple $\bar\imath \in I^n$ is the atomic type of $\bar\imath$ in $\langle I, \leq \rangle$.

Definition 3.2. Suppose that X and Y are disjoint sets of variables and $\Delta \subseteq \mathrm{FO}[\Sigma, X \cup Y]$ a set of formulae. Let $\mathfrak{M}$ be a Σ-structure, $U \subseteq M$, and $(\bar a^i)_{i \in I}$ a sequence of tuples in M.

(a) The *Δ-type* of a tuple $\bar b \subseteq M$ over U is the set

$$\mathrm{tp}_\Delta(\bar b/U) := \big\{\, \varphi(\bar x; \bar c) \;\big|\; \mathfrak{M} \vDash \varphi(\bar b; \bar c),\ \bar c \subseteq U,\ \varphi(\bar x, \bar y) \in \Delta,\ \bar x \subseteq X,\ \bar y \subseteq Y \,\big\}$$

(b) We call $(\bar a^i)_{i \in I}$ a *Δ-indiscernible sequence* over U, or a *sequence of Δ-indiscernibles,* if

$$\mathrm{tp}_\Delta(\bar a[\bar\imath]/U) = \mathrm{tp}_\Delta(\bar a[\bar k]/U)\,, \quad \text{for all } \bar\imath, \bar k \in [I]^{<\omega}.$$

For $\Delta = \mathrm{FO}[\Sigma, X \cup Y]$ we drop the Δ and simply speak of *indiscernible sequences.*

(c) The sequence $(\bar{a}^i)_i$ is *totally Δ-indiscernible* over U if

$$\mathrm{tp}_\Delta(\bar{a}[\bar{\imath}]/U) = \mathrm{tp}_\Delta(\bar{a}[\bar{k}]/U)\,,$$

for all finite sequences $\bar{\imath}, \bar{k} \in I^{<\omega}$ of distinct elements with $|\bar{\imath}| = |\bar{k}|$.

Example. (a) If $\langle A, < \rangle$ is an open dense linear order then every strictly increasing sequence $(a^i)_{i \in I}$ in A is indiscernible. Such a sequence is obviously not totally indiscernible.

(b) Let $\mathfrak{K}$ be an algebraically closed field. Every sequence of algebraically independent elements is totally indiscernible. Similarly, if $\mathfrak{V}$ is a vector space then every sequence of linearly independent elements is totally indiscernible.

For finite sets Δ, we can use the Ramsey Theorem to show that every infinite sequence contains a Δ-indiscernible subsequence. For infinite Δ, we need to apply the Compactness Theorem to find Δ-indiscernible sequences.

Lemma 3.3. *Let $(\bar{a}^i)_{i \in I}$ be an infinite sequence. For every finite set Δ of formulae there exists an infinite subset $I_0 \subseteq I$ such that $(\bar{a}^i)_{i \in I_0}$ is Δ-indiscernible.*

Proof. Let n be the maximal number such that Δ contains a formula $\varphi(\bar{x}^0, \dots, \bar{x}^{n-1})$ with n tuples of variables. We define a colouring $c : [I]^n \to \mathscr{P}(\Delta)$ by

$$c(\bar{\imath}) := \{\, \varphi(\bar{x}^0, \dots, \bar{x}^{n-1}) \in \Delta \mid \mathbb{M} \vDash \varphi(\bar{a}[\bar{\imath}]) \,\}\,.$$

By the Ramsey Theorem there exists an infinite subset $I_0 \subseteq I$ that is homogeneous with respect to c. By definition of c it follows that $(\bar{a}^i)_{i \in I_0}$ is Δ-indiscernible. □

To find Δ-indiscernible sequences, for infinite sets Δ, we apply the Compactness Theorem. Before doing so, let us introduce the average type of a sequence.

Definition 3.4. The *average type* of a sequence $(\bar{a}^i)_i$ over U is the set

$$\mathrm{Av}((\bar{a}^i)_i/U) := \{ \varphi(\bar{x}^0, \dots, \bar{x}^{n-1}; \bar{c}) \mid \bar{c} \subseteq U \text{ and } \mathbb{M} \vDash \varphi(\bar{a}[\bar{\imath}]; \bar{c}) \text{ for all } \bar{\imath} \in [I]^n \} .$$

Lemma 3.5. *Let* $(\bar{a}^i)_{i \in I}$ *be a sequence. Then* $\mathrm{Av}((\bar{a}^i)_i/U)$ *is a partial type. If* $(\bar{a}^i)_i$ *is indiscernible over* U*, it is complete.*

Proposition 3.6. *Let* $\mathfrak{M}$ *be a* Σ*-structure and* $U \subseteq M$ *a set of parameters. For every infinite sequence* $(\bar{a}^i)_{i \in I}$ *and every linear order* J *there exists an elementary extension* $\mathfrak{N} \succeq \mathfrak{M}$ *containing an indiscernible sequence* $(\bar{b}^j)_{j \in J}$ *over* U *such that*

$$\mathrm{Av}((\bar{a}^i)_i/U) \subseteq \mathrm{Av}((\bar{b}^j)_j/U) .$$

Proof. For every $j \in J$, fix a tuple of new constant symbols $\bar{c}^j$ and set

$$\Phi := \{ \varphi(\bar{c}[\bar{\jmath}]; \bar{d}) \mid \varphi(\bar{x}; \bar{d}) \in \mathrm{Av}((\bar{a}^i)_i/U),\ \bar{\jmath} \in [J]^{<\omega},\ \bar{d} \subseteq U \}$$

$$\Psi := \{ \psi(\bar{c}[\bar{\imath}]; \bar{d}) \leftrightarrow \psi(\bar{c}[\bar{\jmath}]; \bar{d}) \mid \psi \text{ a formula, } \bar{\imath}, \bar{\jmath} \in [J]^{<\omega}, \text{ and } \bar{d} \subseteq U \} .$$

It is sufficient to prove that the set $\Gamma := \mathrm{Th}(\mathfrak{M}_M) \cup \Phi \cup \Psi$ is satisfiable. Consider a finite subset $\Gamma_0 \subseteq \Gamma$. Since $\mathrm{Th}(\mathfrak{M}_M)$ is closed under conjunctions, we may assume that $\Gamma_0 = \{\vartheta(\bar{d})\} \cup \Phi_0 \cup \Psi_0$ for finite sets $\Phi_0 \subseteq \Phi$ and $\Psi_0 \subseteq \Psi$. By Lemma 3.3, there is an infinite subset $I_0 \subseteq I$ such that we have

$$\mathfrak{M} \vDash \psi(\bar{a}[\bar{\imath}]; \bar{d}) \leftrightarrow \psi(\bar{a}[\bar{\jmath}]; \bar{d}) ,$$

for every formula $\psi(\bar{x}; \bar{d}) \leftrightarrow \psi(\bar{y}; \bar{d}) \in \Psi_0$ and all increasing $\bar{\imath}, \bar{\jmath} \subseteq I_0$. For every formula $\varphi(\bar{x}; \bar{d}) \in \Phi_0$, there are only finitely many indices $\bar{\imath} \subseteq I_0$ such that $\mathfrak{M} \nvDash \varphi(\bar{a}[\bar{\imath}]; \bar{d})$. Hence, we can find an infinite subset $I_1 \subseteq I_0$ containing no such tuple $\bar{\imath}$. Let $J_0 \subseteq J$ be the finite set of all indices $j \in J$ such that the constant $\bar{c}^j$ appears in $\Phi_0 \cup \Psi_0$, and fix an embedding $g : J_0 \to I_1$. We can satisfy Γ_0 by interpreting $\bar{c}^j$ by the tuple $\bar{a}^{g(j)}$. □

We can improve the preceding proposition as follows.

Theorem 3.7. *Let $\mathfrak{M}$ be a Σ-structure, $U \subseteq M$ a set of parameters, $\bar{s}$ a sequence of sorts, and λ a cardinal such that $\lambda \geq |S^{\bar{s}^n}(U)|$, for all $n < \omega$. Set $\mu := \beth_{\lambda^+}$.*

For every sequence $(\bar{a}^\alpha)_{\alpha<\mu}$ with $\bar{a}^\alpha \in M^{\bar{s}}$ and for every linear order I, there exists an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ containing an indiscernible sequence $(\bar{b}^i)_{i\in I}$ over U such that, for every $\bar{\imath} \in [I]^n$, there are indices $\bar{\alpha} \in [\mu]^n$ with

$$\operatorname{tp}(\bar{b}[\bar{\imath}]/U) = \operatorname{tp}(\bar{a}[\bar{\alpha}]/U) .$$

Proof. It is sufficient to prove the claim for $I = \omega$. Then the general statement will follow by compactness. We define a sequence of types $(\mathfrak{p}_n)_{n<\omega}$ with $\mathfrak{p}_n \in S^{\bar{s}^n}(U)$ satisfying the following conditions:

(1) $\mathfrak{p}_n(\bar{x}_0, \dots, \bar{x}_{n-1}) \vDash \mathfrak{p}_m(\bar{x}_{i_0}, \dots, \bar{x}_{i_{m-1}})$, for all $i_0 < \dots < i_{m-1} < n$.

(2) For every cardinal $\nu < \mu$, there is some set $I \subseteq \mu$ of size $|I| = \nu$ such that

$$\operatorname{tp}(\bar{a}[\bar{\imath}]/U) = \mathfrak{p}_n , \quad \text{for every tuple } \bar{\imath} \in [I]^n .$$

Any sequence $(\bar{b}^n)_{n<\omega}$ realising the limit $\mathfrak{p}_\omega := \bigcup_{n<\omega} \mathfrak{p}_n$ has the desired properties.

We start with $\mathfrak{p}_0 := \operatorname{Th}(\mathfrak{M}_U)$. If we have already defined $\mathfrak{p}_n$, we consider the set X of all $\bar{s}^{n+1}$-types over U satisfying condition (1). If there is some type $\mathfrak{q} \in X$ that also satisfies condition (2), we are done. Suppose there is no such type. Then we can choose, for every $\mathfrak{q} \in X$, a cardinal $\nu_\mathfrak{q} < \mu$ such that no subset $I \subseteq \mu$ of size $\nu_\mathfrak{q}$ satisfies the above condition. Since $|X| \leq \lambda < \lambda^+ = \operatorname{cf} \mu$ it follows that

$$\nu_* := \lambda \oplus \sup \{ \nu_\mathfrak{q} \mid \mathfrak{q} \in X \} < \mu .$$

By choice of ν_* there exists, for every $\mathfrak{q} \in X$ and all $I \subseteq \mu$ of size $|I| = \nu_*$, some increasing tuple $\bar{\imath} \in I^{n+1}$ such that $\operatorname{tp}(\bar{a}[\bar{\imath}]/U) \neq \mathfrak{q}$. Since $\nu_* < \mu = \beth_{\lambda^+}$ there is some ordinal $\alpha < \lambda^+$ with $\nu_* < \beth_\alpha$. Let $\rho := \beth_{\alpha+n+1}$. Then

$$\beth_n(\nu_*)^+ \leq \rho < \mu .$$

By choice of $\mathfrak{p}_n$ there is some set $I \subseteq \mu$ of size $|I| = \rho$ such that

$$\operatorname{tp}(\bar{a}[\bar{\imath}]/U) = \mathfrak{p}_n\,, \quad \text{for every } \bar{\imath} \in [I]^n.$$

Since $|S^{\bar{s}^n}(U)| \leq \lambda \leq \nu_*$ we can use the Theorem of Erdős and Rado to find a subset $I_0 \subseteq I$ of size $|I_0| = \nu_*^+$ such that the types

$$\operatorname{tp}(\bar{a}[\bar{\imath}]/U)\,, \quad \text{for } \bar{\imath} \in [I_0]^{n+1}\,,$$

are all equal. This contradicts the choice of ν_*. □

There is a close relationship between automorphisms and indiscernible sequences. The next observation follows immediately from the definitions of an indiscernible sequence and a strongly κ-homogeneous structure.

Lemma 3.8. *Let $\mathfrak{M}$ be strongly κ-homogeneous and let $(\bar{a}^i)_{i\in I}$ be a sequence of indiscernible over U. Suppose that $|U| \oplus |I| \oplus |\bar{a}^i| < \kappa$. For every partial automorphism $\pi \in \operatorname{pIso}(I, I)$ of the index set I (considered as a linear order), there exists an automorphism $h \in \operatorname{Aut}\mathfrak{M}$ such that*

$$h \upharpoonright U = \operatorname{id}_U \quad \textit{and} \quad h(\bar{a}^i) = \bar{a}^{\pi(i)}\,, \quad \textit{for all } i \in I\,.$$

In a sufficient saturated structure, we can extend every indiscernible sequence to a longer one.

Lemma 3.9. *Let $\mathfrak{M}$ be κ-saturated. If $(\bar{a}^i)_{i\in I}$ is indiscernible over U and $g : I \to J$ is an embedding with $|J| \oplus |U| \oplus |\bar{a}^i| < \kappa$ then there exists an indiscernible sequence $(\bar{b}^j)_{j\in J}$ such that $\bar{a}^i = \bar{b}^{g(i)}$, for $i \in I$.*

Proof. We can use Proposition 3.6 to find an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ containing an indiscernible sequence $(\bar{c}^j)_{j\in J}$ with $\operatorname{Av}((\bar{c}^j)_j/U) = \operatorname{Av}((\bar{a}^i)_i/U)$. This implies that

$$\operatorname{tp}(\textstyle\bigcup_i \bar{c}^{g(i)}/U) = \operatorname{tp}(\bigcup_i \bar{a}^i/U)\,.$$

W.l.o.g. we may assume that $\mathfrak{N}$ is strongly κ-homogeneous. Therefore, there exists an automorphism π of $\mathfrak{N}_U$ mapping $\bar{c}^{g(i)}$ to $\bar{a}^i$. Since $\mathfrak{M}$ is κ-saturated it contains a sequence $(\bar{b}^j)_{j\in J}$ such that

$$\operatorname{tp}(\bigcup_j \bar{b}^j/U \cup \bigcup_i \bar{a}^i) = \operatorname{tp}(\bigcup_j \pi(\bar{c}^j)/U \cup \bigcup_i \bar{a}^i)\,.$$

It follows that $(\bar{b}^j)_j$ is the desired sequence of indiscernibles. □

Corollary 3.10. *If $(\bar{a}^i)_{i\in I}$ is indiscernible over U and $g : I \to J$ an embedding, then there exists an elementary extension $\mathfrak{N}$ containing an indiscernible sequence $(\bar{b}^i)_{i\in J}$ such that $\bar{b}^{g(i)} = \bar{a}^i$, for $i \in I$.*

Let us record the following consequence of Theorem 3.7.

Lemma 3.11. *Let $(\bar{a}_i)_{i\in I}$ be an indiscernible sequences over U. For every set $C \subseteq \mathbb{M}$, there exists a set $C' \equiv_U C$ such that $(\bar{a}_i)_{i\in I}$ is indiscernible over $U \cup C'$.*

Proof. Let $\kappa := |T| \oplus |U \cup C|$ and $\lambda := \beth_{(2^\kappa)^+}$. By Corollary 3.10, there exists an indiscernible sequence $(\bar{b}_\alpha)_{\alpha<\kappa}$ over U with

$$\operatorname{Av}((\bar{b}_\alpha)_\alpha/U) = \operatorname{Av}((\bar{a}_i)_i/U)\,.$$

Furthermore, with the help of Theorem 3.7 we can find an indiscernible sequence $(\bar{c}_n)_{n<\omega}$ over $U\cup C$ such that, for every $n < \omega$, there are indices $\alpha_0 < \cdots < \alpha_{n-1}$ with

$$\bar{c}_0 \dots \bar{c}_{n-1} \equiv_{U\cup C} \bar{b}_{\alpha_0} \dots \bar{b}_{\alpha_{n-1}}\,.$$

By Lemma 3.9, we can extend $(\bar{c}_n)_{n<\omega}$ to an indiscernible sequence $(\bar{c}_i)_{i\in\omega+I}$ over $U \cup C$. Since

$$\operatorname{Av}((\bar{c}_i)_i/U) = \operatorname{Av}((\bar{a}_i)_i/U)\,,$$

there exists an automorphism $\pi \in \operatorname{Aut}\mathbb{M}_U$ such that $\pi(\bar{c}_{\omega+i}) = \bar{a}_i$, for all $i \in I$. Then $\pi[C] \equiv_U C$ and $(\bar{a}_i)_{i\in I}$ is indiscernible over $U \cup \pi[C]$. □

The following technical lemma can be used to simplify proofs of indiscernibility. It states that, if some formula is a witness for the failure of indiscernibility, we can detect this fact already by varying a single element of the sequence.

Lemma 3.12. *Let $\alpha = (\bar{a}_i)_{i \in I}$ be a sequence and $\varphi(\bar{x})$ a formula such that*

$$\mathbb{M} \vDash \varphi(\bar{a}[\bar{\imath}]) \wedge \neg\varphi(\bar{a}[\bar{\jmath}])\,, \quad \text{for some } \bar{\imath}, \bar{\jmath} \in [I]^n.$$

Then there are indices $\bar{u} < s < t < \bar{v}$ in I such that

$$\mathbb{M} \vDash \varphi(\bar{a}[\bar{u}s\bar{v}]) \leftrightarrow \neg\varphi(\bar{a}[\bar{u}t\bar{v}])\,.$$

Proof. We define a sequence $\bar{k}^0, \dots, \bar{k}^{2n} \in [I]^n$ by setting

$$k^l_m := \begin{cases} \min\{i_m, j_m\} & \text{if } l \leq n \text{ and } m < l\,, \\ i_m & \text{if } l \leq n \text{ and } m \geq l\,, \\ \min\{i_m, j_m\} & \text{if } l > n \text{ and } m < 2n - l\,, \\ j_m & \text{if } l > n \text{ and } m \geq 2n - l\,. \end{cases}$$

Then every $\bar{k}_l$ belongs to $[I]^n$, $\bar{k}_0 = \bar{\imath}$, $\bar{k}_{2n} = \bar{\jmath}$, and, for each $l < 2n$, the tuples $\bar{k}_l$ and $\bar{k}_{l+1}$ differ in at most one component. Let $l < 2n$ be the maximal index such that $\mathbb{M} \vDash \varphi(\bar{a}[\bar{k}_l])$. Then $\mathbb{M} \vDash \neg\varphi(\bar{a}[\bar{k}_{l+1}])$ and it follows by definition of $\bar{k}_l$ that $\bar{k}_l = \bar{u}s\bar{v}$ and $\bar{k}_{l+1} = \bar{u}t\bar{v}$ for indices $\bar{u} < s < \bar{v}$ and $\bar{u} < t < \bar{v}$. Interchanging $\bar{k}_l$ and $\bar{k}_{l+1}$ if necessary, we may assume that $s < t$. □

Recall that stable theories do not have the order property. This implies that in a model of a stable theory every indiscernible sequence is totally indiscernible.

Theorem 3.13. *A theory T is stable if, and only if, every infinite indiscernible sequence in a model of T is totally indiscernible.*

Proof. ($\Leftarrow$) Suppose that there is a formula $\varphi(\bar{x}, \bar{y})$ with the order property and let $(\bar{a}^n)_{n<\omega}$ and $(\bar{b}^n)_{n<\omega}$ be sequences such that

$$\mathbb{M} \models \varphi(\bar{a}^i, \bar{b}^k) \quad \text{iff} \quad i \leq k\,.$$

By Proposition 3.6, there exists an indiscernible sequence $(\bar{c}^n \bar{d}^n)_{n<\omega}$ with $\mathrm{Av}((\bar{a}^n \bar{b}^n)_n) \subseteq \mathrm{Av}((\bar{c}^n \bar{d}^n)_n)$. Setting $\psi(\bar{x}\bar{y}, \bar{x}'\bar{y}') := \varphi(\bar{x}, \bar{y}')$ it follows that

$$\mathbb{M} \models \psi(\bar{c}^i \bar{d}^i, \bar{c}^k \bar{d}^k) \quad \text{iff} \quad i \leq k\,.$$

Hence, $(\bar{c}^n \bar{d}^n)_n$ is not totally indiscernible.

($\Rightarrow$) Suppose that $(\bar{a}^i)_{i \in I}$ is an infinite indiscernible sequence over U that is not totally indiscernible. By Corollary 3.10, we may assume that the ordering I is dense. There are a formula φ and two tuples of indices $\bar{\imath}, \bar{k} \in I^n$ such that both $\bar{\imath}$ and $\bar{k}$ consist of distinct elements and we have

$$\mathbb{M} \models \varphi(\bar{a}[\bar{\imath}]) \wedge \neg\varphi(\bar{a}[\bar{k}])\,.$$

Set $\bar{l}^r := i_0 \dots i_{r-1} k_r \dots k_{n-1}$ and let r be the maximal number such that

$$\mathbb{M} \models \neg\varphi(\bar{a}[\bar{l}^r])\,.$$

Note that r is well-defined since $\bar{l}^0 = \bar{k}$ implies $\mathbb{M} \models \neg\varphi(\bar{a}[\bar{l}^0])$. Replacing $\bar{\imath}$ by $\bar{l}^{r+1}$ and $\bar{k}$ by $\bar{l}^r$, we may assume that $\bar{\imath}$ and $\bar{k}$ differ in exactly one component. Hence, suppose that

$$\bar{\imath} = s\bar{u}\bar{v}\bar{w} \quad \text{and} \quad \bar{k} = t\bar{u}\bar{v}\bar{w}\,, \quad \text{where } \bar{u} < s < \bar{v} < t < \bar{w}\,.$$

(Reversing the order of I, if necessary, we may assume that $s < t$.)

By indiscernibility, we know that the tuple $\bar{v}$ is not empty. We claim that we may assume that $\bar{v}$ is a singleton. If $\bar{v} = v_0 \dots v_{n-1}$ with $n > 1$ then, choosing some index $v_0 < v' < v_{n-1}$, we may replace either s or t by v', depending on whether or not the formula $\varphi(\bar{a}[v'\bar{u}\bar{v}\bar{w}])$ holds. Hence, the claim follows by induction. Thus, we have arrived at the situation that

$$\bar{\imath} = sv\bar{u}\bar{w} \quad \text{and} \quad \bar{k} = vt\bar{u}\bar{w}\,, \quad \text{where } \bar{u} < s < v < t < \bar{w}\,.$$

By indiscernibility, it follows that

$$\mathbb{M} \vDash \varphi(\bar{a}[st\bar{u}\bar{w}]) \wedge \neg\varphi(\bar{a}[ts\bar{u}\bar{w}]), \quad \text{for all} \quad \bar{u} < s < t < \bar{w}.$$

Fix an infinite increasing sequence of indices k_n, $n < \omega$, with

$$\bar{u} < k_0 < k_1 < \cdots < \bar{w},$$

set $\bar{b}^i := \bar{a}^{k_i}$, and define

$$\psi(\bar{x}, \bar{y}) := \bar{x} = \bar{y} \vee [\varphi(\bar{x}, \bar{y}, \bar{a}[\bar{u}\bar{w}]) \wedge \neg\varphi(\bar{y}, \bar{x}, \bar{a}[\bar{u}\bar{w}])].$$

Then we have

$$\mathbb{M} \vDash \psi(\bar{b}^i, \bar{b}^k) \quad \text{iff} \quad i \leq k.$$

Hence, T is unstable. □

When considering the automorphism group of a structure, an indiscernible sequence looks like a linear order while a totally indiscernible sequence looks like a set. We can generalise the definition of an indiscernible sequence to include automorphism groups of other structures.

Definition 3.14. Let L be an algebraic logic, $\mathfrak{I}$ a Γ-structure, $\mathfrak{M}$ a Σ-structure, and $U \subseteq M$.

(a) A *U-indiscernible system* over $\mathfrak{I}$ (w.r.t. L) is an injective function $\bar{a} : I \to M^\alpha$, for some ordinal α, such that, for every partial isomorphism $\bar{\imath} \mapsto \bar{k} \in \mathrm{pIso}_{\aleph_0}(\mathfrak{I}, \mathfrak{I})$, we have

$$\mathrm{tp}_L(\bar{a}[\bar{\imath}]/U) = \mathrm{tp}_L(\bar{a}[\bar{k}]/U).$$

(b) The *average type* of a U-indiscernible system $\bar{a}$ over $\mathfrak{I}$ is the function $\mathrm{Av}_L(\bar{a})$ with

$$\mathrm{Av}_L(\bar{a}/U) : \mathrm{atp}(\bar{\imath}/\mathfrak{I}) \mapsto \mathrm{tp}_L(\bar{a}[\bar{\imath}]/U), \quad \text{for } \bar{\imath} \in I^{<\omega}.$$

For $L = \mathrm{FO}$, we drop the index and just write $\mathrm{Av}(\bar{a}/U)$.

(c) Let $\mathfrak{I}$ and $\mathfrak{K}$ be two index structures and $\bar{a} : I \to M^\alpha$, $\bar{b} : K \to M^\alpha$ arbitrary families of α-tuples. We say that $\bar{a}$ is *inspired* by $\bar{b}$ over U if, for every finite set of formulae Δ and every finite tuple $\bar{\imath} \in I^{<\omega}$, there is a finite tuple $\bar{k} \in K^{<\omega}$ such that

$$\operatorname{atp}(\bar{\imath}/\mathfrak{I}) = \operatorname{atp}(\bar{k}/\mathfrak{K}) \quad \text{and} \quad \operatorname{tp}_\Delta(\bar{a}[\bar{\imath}]/U) = \operatorname{tp}_\Delta(\bar{b}[\bar{k}]/U)\,.$$

Remark. (a) Using the terminology of the previous definition we can restate Proposition 3.6 as: for every infinite sequence $(\bar{a}^i)_{i \in I}$, every linear order J, and every set U of parameters, there exists an indiscernible sequence $(\bar{b}^i)_{i \in J}$ over U inspired by $(\bar{a}^i)_{i \in I}$.

(b) Note that, for indiscernible systems $\bar{a}$ and $\bar{b}$ over U, $\bar{a}$ is inspired by $\bar{b}$ over U if, and only if, $\operatorname{Av}(\bar{a}/U) = \operatorname{Av}(\bar{b}/U)$.

In the same way as in Proposition 3.6 we can use the Compactness Theorem to show that we can extend every indiscernible system.

Lemma 3.15. *Let $\mathfrak{M}$ be a structure containing a U-indiscernible system $\bar{a}$ over $\mathfrak{I}$. If $\mathfrak{H}$ is a structure with* $\operatorname{Sub}_{\aleph_0}(\mathfrak{H}) \subseteq \operatorname{Sub}_{\aleph_0}(\mathfrak{I})$ *then there exists an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ containing a U-indiscernible system $\bar{b}$ over $\mathfrak{H}$ with* $\operatorname{Av}(\bar{b}/U) = \operatorname{Av}(\bar{a}/U)$.

In general, it is hard to prove the existence of indiscernible systems over structures that are not linear orders. For trees we can use the Theorem of Milliken to show that such systems always exist. Recall the trees $\mathfrak{T}_*(\kappa^{<\alpha})$ introduced in Section 2.

Definition 3.16. Let κ be a cardinal and α an ordinal. A family $(\bar{a}_\eta)_{\eta \in \kappa^{<\alpha}}$ is called *tree-indiscernible* over a set U if it is a U-indiscernible system over $\mathfrak{T}_*(\kappa^{<\alpha})$.

Theorem 3.17 (Džamonja, Shelah, B. Kim, H.-J. Kim). *Let $m < \omega$. For every family $\bar{a} = (\bar{a}_\eta)_{\eta \in m^{<\omega}}$ and every set U, there exists a family of tree-indiscernibles $(\bar{b}_\eta)_{\eta \in m^{<\omega}}$ over U inspired by $\bar{a}$.*

Proof. Fix variable symbols $\bar{x}_\eta$, for each $\eta \in m^{<\omega}$, and define

$$\begin{aligned} \Psi_{\bar\eta} &:= \big\{ \varphi(\bar{x}[\bar\zeta]) \bigm| \varphi \text{ a formula over } U,\ \bar\zeta \approx_* \bar\eta,\ \text{and} \\ &\qquad\qquad \mathbb{M} \vDash \varphi(\bar{a}[\bar\xi]) \text{ for all } \bar\xi \approx_* \bar\eta \big\}, \\ \Xi &:= \big\{ \varphi(\bar{x}[\bar\eta]) \leftrightarrow \varphi(\bar{x}[\bar\zeta]) \bigm| \varphi \text{ a formula over } U, \bar\eta \approx_* \bar\zeta \big\}, \end{aligned}$$

and $\quad \Phi := \Xi \cup \bigcup_{\bar\eta \subseteq m^{<\omega}} \Psi_{\bar\eta}\,.$

We claim that Φ is satisfiable. Let $\Phi_0 \subseteq \Phi$ be finite. There exists a finite set Δ of formulae such that every formula in Φ_0 is of the form

$$\varphi(\bar{x}[\bar\eta]) \leftrightarrow \varphi(\bar{x}[\bar\zeta]) \quad \text{or} \quad \varphi(\bar{x}[\bar\zeta])\,,$$

for some $\varphi(\bar{x}_0, \dots, \bar{x}_{n-1}) \in \Delta$. Let d be the number of variables appearing in Δ and let $c : (m^{<\omega})^d \to S(\Delta)$ be the function mapping each tuple $\bar\eta \in (m^{<\omega})^d$ to the type $\mathrm{tp}_\Delta(\bar{a}[\bar\eta])$.

Let $k < \omega$ be some number such that Φ_0 only contains variables $\bar{x}_\eta$ with $\eta \in m^{<k}$. We can use Theorem 2.10 to find an embedding $g : \mathfrak{T}_*(m^{<k}) \to \mathfrak{T}_*(m^{<\omega})$ such that $\mathrm{rng}\, g$ is $\approx_*$-homogeneous with respect to c. It follows that the family $(\bar{a}_{g(\eta)})_{\eta \in m^{<k}}$ satisfies Φ_0.

By the Compactness Theorem we conclude that Φ is satisfiable. Let $\bar{b} = (\bar{b}_\eta)_{\eta \in m^{<\omega}}$ be a family realising Φ. Then $\bar{b}$ is tree-indiscernible over U since it satisfies Ξ. Hence, it remains to show that $\bar{b}$ is inspired by $\bar{a}$.

For a contradiction, suppose otherwise. Then there exist a finite tuple $\bar\eta \subseteq m^{<\omega}$ and a finite set of formulae Δ over U such that

$$\mathrm{tp}_\Delta(\bar{b}[\bar\eta]) \neq \mathrm{tp}_\Delta(\bar{a}[\bar\zeta])\,, \quad \text{for all } \bar\zeta \approx_* \bar\eta\,.$$

W.l.o.g. we may assume that Δ is closed under negation. Set

$$\vartheta(\bar{x}) := \bigwedge \mathrm{tp}_\Delta(\bar{b}[\bar\eta])\,.$$

Then

$$\mathbb{M} \vDash \neg\vartheta(\bar{a}[\bar\zeta])\,, \quad \text{for all } \bar\zeta \approx_* \bar\eta\,.$$

Consequently, $\neg\vartheta(\bar{x}[\bar\eta]) \in \Psi_{\bar\eta}$. Since $\bar{b}$ satisfies $\Psi_{\bar\eta}$ it therefore follows that $\mathbb{M} \vDash \neg\vartheta(\bar{b}[\bar\eta])$. A contradiction. □

4. *The independence and strict order properties*

In this section we use indiscernible sequences to study concepts related to the order property. Recall that

$$[\![\varphi(\bar{a}, \bar{b}^i)]\!]_{i \in I} := \{\, i \in I \mid \mathbb{M} \vDash \varphi(\bar{a}, \bar{b}^i) \,\}\,.$$

Definition 4.1. Let T be a theory. A formula $\varphi(\bar{x}, \bar{y})$ has the *independence property* (with respect to T) if there exists a model $\mathfrak{M} \vDash T$ containing two sequences $(\bar{a}^w)_{w \in \wp(\omega)}$ and $(\bar{b}^n)_{n<\omega}$ such that

$$\mathfrak{M} \vDash \varphi(\bar{a}^w, \bar{b}^n) \quad \text{iff} \quad n \in w\,.$$

If some formula has the independence property with respect to T, we also say that T has the *independence property.*

Proposition 4.2. *Let T be a first-order theory and $\varphi(\bar{x}, \bar{y})$ a formula. The following statements are equivalent:*

(1) *φ has the independence property.*

(2) *For every finite number $m < \omega$, there exist sequences $(\bar{a}^w)_{w \in \wp[m]}$ and $(\bar{b}^n)_{n<m}$ such that*

$$\mathbb{M} \vDash \varphi(\bar{a}^w, \bar{b}^n) \quad \text{iff} \quad n \in w\,.$$

(3) *There exist a sequence $(\bar{a}^w)_{w \in \wp(\omega)}$ and an indiscernible sequence $(\bar{b}^n)_{n<\omega}$ such that*

$$\mathfrak{M} \vDash \varphi(\bar{a}^w, \bar{b}^n) \quad \text{iff} \quad n \in w\,.$$

(4) *There exist a tuple $\bar{a}$ and an indiscernible sequence $(\bar{b}^n)_{n<\omega}$ such that*

$$[\![\varphi(\bar{a}, \bar{b}^n)]\!]_{n<\omega} = \{\, 2n \mid n < \omega \,\}\,.$$

(5) *There exist a tuple $\bar{a}$ and an indiscernible sequence $(\bar{b}^i)_{i \in I}$ such that $[\![\varphi(\bar{a}, \bar{b}^i)]\!]_{i \in I}$ is not a finite union of segments.*

Proof. The implications (3) ⇒ (4) ⇒ (5) are trivial and (2) ⇒ (1) follows by compactness.

For (1) ⇒ (3), let $(\bar{a}^w)_{w \in \wp(\omega)}$ and $(\bar{b}^n)_{n<\omega}$ be sequences such that

$$\mathfrak{M} \vDash \varphi(\bar{a}^w, \bar{b}^n) \quad \text{iff} \quad n \in w\,.$$

By Proposition 3.6, there exists an indiscernible sequence $(\bar{d}^n)_{n<\omega}$ with the same average type as $(\bar{b}^n)_{n<\omega}$. By compactness, we can find a sequence $(\bar{c}^w)_{w \in \wp(\omega)}$ such that

$$\mathfrak{M} \vDash \varphi(\bar{c}^w, \bar{d}^n) \quad \text{iff} \quad n \in w\,.$$

It remains to prove (5) ⇒ (2). Fix $m < \omega$ and let $\bar{a}$ and $(\bar{b}^i)_{i \in I}$ be such that $[\![\varphi(\bar{a}, \bar{b}^i)]\!]_{i \in I}$ is not a finite union of segments. We can find a strictly increasing sequence $i_0 < \dots < i_{2m-1}$ of indices in I such that

$$\mathbb{M} \vDash \varphi(\bar{a}, \bar{b}^{i_k}) \quad \text{iff} \quad k \text{ is odd}\,.$$

Set $\bar{d}^k := \bar{b}^{i_k}$ and let

$$\chi_w(k) := \begin{cases} 0 & \text{if } k \notin w\,, \\ 1 & \text{if } k \in w\,, \end{cases}$$

be the characteristic function of w. Note that the sequence $(\bar{d}^k)_{k<2m}$ is also indiscernible. For each $w \subseteq [m]$, we can therefore find an automorphism π_w of $\mathbb{M}$ such that

$$\pi_w(\bar{d}^k) = \bar{d}^{2n+\chi_w(k)}\,, \quad \text{for } k < m\,.$$

Setting $\bar{c}^w := \pi_w^{-1}(\bar{a})$ it follows that

$$\begin{aligned} \mathbb{M} \vDash \varphi(\bar{c}^w, \bar{d}^k) \quad &\text{iff} \quad \mathbb{M} \vDash \varphi(\pi_w(\bar{c}^w), \pi_w(\bar{d}^k)) \\ &\text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a}, \bar{d}^{2n+\chi_w(k)}) \\ &\text{iff} \quad \chi_w(k) = 1 \\ &\text{iff} \quad k \in w\,. \end{aligned}$$

□

We can generalise Condition (4) above as follows.

Corollary 4.3. *Let $\varphi(\bar{x}; \bar{y}_0, \dots, \bar{y}_{n-1})$ be a formula. If there exist a tuple $\bar{c}$ and an indiscernible sequence $(\bar{a}_i)_{i \in I}$ such that the order I has no last element,*

$$\begin{aligned} &\mathbb{M} \vDash \varphi(\bar{c}; \bar{a}[\bar{\imath}])\,, && \textit{for arbitrarily large } \bar{\imath} \in [I]^n, \\ \textit{and} \quad &\mathbb{M} \vDash \neg\varphi(\bar{c}; \bar{a}[\bar{\imath}])\,, && \textit{for arbitrarily large } \bar{\imath} \in [I]^n, \end{aligned}$$

then φ has the independence property.

Proof. By assumption we can inductively choose tuples $\bar{k}_0 < \bar{k}_1 < \dots$ in $[I]^n$ such that

$$\mathbb{M} \vDash \varphi(\bar{c}; \bar{a}[\bar{k}_i]) \quad \text{iff} \quad i \text{ is even.}$$

Since the sequence $(\bar{a}[\bar{k}_i])_{i<\omega}$ is indiscernible, the claim follows by Proposition 4.2 (4). □

Lemma 4.4. *Let T be a first-order theory. If $\varphi(\bar{x}, \bar{y})$ has the independence property then so does $\varphi(\bar{y}, \bar{x})$.*

Proof. We apply the characterisation in Proposition 4.2 (2). Let $m < \omega$. Since $\varphi(\bar{x}, \bar{y})$ has the independence property there are tuples $\bar{a}^w$ and $\bar{b}^n$ for $w \subseteq \wp(2^m)$ and $n < 2^m$ such that

$$\mathbb{M} \vDash \varphi(\bar{a}^w, \bar{b}^n) \quad \text{iff} \quad n \in w\,.$$

We identify each number $k < 2^m$ with the function $k : [m] \to [2]$ such that $k = \sum_{i<m} k(i)2^i$. For $i < m$ and $s \subseteq [m]$, we define

$$\bar{c}^s := \bar{b}^{n_s} \quad \text{and} \quad \bar{d}^i := \bar{a}^{w_i}\,,$$

where

$$n_s := \textstyle\sum_{i \in s} 2^i \quad \text{and} \quad w_i := \{\, k < 2^m \mid k(i) = 1 \,\}\,.$$

It follows that

$$\begin{aligned} \mathbb{M} \vDash \varphi(\bar{d}^i, \bar{c}^s) \quad &\text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a}^{w_i}, \bar{b}^{n_s}) \\ &\text{iff} \quad n_s \in w_i \\ &\text{iff} \quad i \in s . \end{aligned}$$

□

Lemma 4.5. *Let T be a first-order theory and $\varphi(\bar{x}, \bar{y})$ a formula with the independence property. There exist formulae $\psi(x, \bar{y})$ and $\vartheta(\bar{x}, y)$ with, respectively, a single variable x and a single variable y that have the independence property.*

Proof. We construction ψ using Proposition 4.2 (3). Let $\bar{a}$ and $(\bar{b}^n)_{n<\omega}$ be tuples such that $[\![\varphi(\bar{a}, \bar{b}^n)]\!]_{n<\omega} = \{ 2n \mid n < \omega \}$. Suppose that $\bar{a} = a_0\bar{a}'$. We define a new sequence $\bar{c}^n := \bar{b}^n\bar{a}'$ and the formula $\psi(x, \bar{y}\bar{z}) := \varphi(x\bar{z}, \bar{y})$. It follows that $[\![\psi(a, \bar{c}^n)]\!]_{n<\omega} = \{ 2n \mid n < \omega \}$. Hence, ψ has the independence property.

To find $\vartheta(\bar{x}, y)$ it is sufficient to note that, according to Lemma 4.4, the formula $\varphi(\bar{y}, \bar{x})$ also has the independence property. Hence, we can apply the first part of the lemma. □

The independence property is closely related to the order property which characterises unstable theories.

Lemma 4.6. *Every formula with the independence property has the order property.*

Proof. Suppose that φ is a formula with the independence property and let $(\bar{a}^w)_{w\subseteq\wp(\omega)}$ and $(\bar{b}^n)_{n<\omega}$ be sequences such that

$$\mathbb{M} \vDash \varphi(\bar{a}^w, \bar{b}^n) \quad \text{iff} \quad n \in w .$$

Setting $w_n := \omega \setminus [n]$ and $\bar{c}^n := \bar{a}^{w_n}$ it follows that

$$\mathbb{M} \vDash \varphi(\bar{c}^n, \bar{b}^k) \quad \text{iff} \quad n \leq k .$$

Hence, φ has the order property. □

Lemma 4.7. *No o-minimal theory has the independence property.*

Proof. Let T be a theory with the independence property. Then there exist a model $\mathfrak{M}$ of T, a formula $\varphi(x, \bar{y})$, parameters $\bar{c} \subseteq M$, and an indiscernible sequence $(a_n)_{n<\omega}$ such that

$$\mathfrak{M} \vDash \varphi(a_n, \bar{c}) \quad \text{iff} \quad n \equiv 0 \pmod 2 .$$

Since $(a_n)_n$ is indiscernible we either have $a_0 < a_1 < \dots$ or $a_0 > a_1 > \dots$. In both cases it follows that the set $\varphi(x, \bar{c})^{\mathfrak{M}}$ is not a finite union of intervals. Hence, T is not o-minimal. □

Lemma 4.8. *Let $\varphi(\bar{x}, \bar{y})$ be a formula without the independence property. Suppose that there exists a tuple $\bar{c}$ and a sequence $(\bar{a}^i)_{i \in I}$ such that the sets $[\![\varphi(\bar{c}, \bar{a}^i)]\!]_i$ and $[\![\neg\varphi(\bar{c}, \bar{a}^i)]\!]_i$ are both infinite. Then there exists a formula $\chi(\bar{y}, \bar{y}'; \bar{d})$ with parameters $\bar{d}$ such that*

$$\mathbb{M} \vDash \chi(\bar{a}^i, \bar{a}^k; \bar{d}) \quad \text{iff} \quad i \leq k .$$

Proof. Let J be an open dense linear order with $I \subseteq J$ such that J contains infinitely many elements above I and below I. By Lemma 3.9, we can extend $(\bar{a}^i)_{i \in I}$ to an indiscernible sequence $(\bar{a}^i)_{i \in J}$. Replacing φ by $\neg\varphi$ if necessary, we may assume that $[\![\varphi(\bar{c}, \bar{a}^i)]\!]_i$ contains a final segment of J. By Proposition 4.2 (2), there exists a number m such that, for all indices $\bar{s} \in [I]^m$,

$$\mathbb{M} \vDash \neg\exists\bar{x} \bigwedge_{i<m-1} [\varphi(\bar{x}, \bar{a}^{s_i}) \leftrightarrow \neg\varphi(\bar{x}, \bar{a}^{s_{i+1}})] .$$

Consequently, there exists a number $0 < n \leq m$, a set $w \subseteq [n]$, and indices $\bar{s} \in [I]^n$ such that there is no $\bar{c}'$ with

$$\begin{aligned} &\downarrow s_0 \cup \{ s_i \mid i \notin w \} \subseteq [\![\neg\varphi(\bar{c}', \bar{a}^i)]\!]_i \\ \text{and} \quad &\uparrow s_{n-1} \cup \{ s_i \mid i \in w \} \subseteq [\![\varphi(\bar{c}', \bar{a}^i)]\!]_i . \end{aligned}$$

We choose n and w such that $\langle n, w \rangle$ is minimal with respect to the lexicographic order (treating $w \subseteq [n]$ as a word in $[2]^n$). By minimality of n, it follows that $0 \in w$ and $n - 1 \notin w$. Hence, there is some index $k < n$ with $[k] \subseteq w$ and $k \notin w$.

By compactness, there are finite sets $J_- \subseteq {\downarrow}s_0$ and $J_+ \subseteq {\uparrow}s_{n-1}$ such that there is no $\bar{c}'$ with

$$J_- \cup \{ s_i \mid i \notin w \} \subseteq [\![\neg\varphi(\bar{c}', \bar{a}^i)]\!]_i$$
$$\text{and} \quad J_+ \cup \{ s_i \mid i \in w \} \subseteq [\![\varphi(\bar{c}', \bar{a}^i)]\!]_i \,.$$

By indiscernibility, we may assume that

$$J_- \cup \{ s_i \mid i < k \} < I < J_+ \cup \{ s_i \mid i \geq k \} \,.$$

Let $w_+ := w \smallsetminus \{k - 1\}$ and $w_- := [n] \smallsetminus (w \cup \{k\})$. We define

$$\psi(\bar{x}) := \bigwedge_{i \in J_- \cup w_-} \neg\varphi(\bar{x}, \bar{a}^i) \wedge \bigwedge_{i \in J_+ \cup w_+} \varphi(\bar{x}, \bar{a}^i) \,.$$

Then

$$\mathbb{M} \vDash \neg\exists\bar{x}\big[\psi(\bar{x}) \wedge \varphi(\bar{x}, \bar{a}^{s_{k-1}}) \wedge \neg\varphi(\bar{x}, \bar{a}^{s_k})\big] \,.$$

Hence,

$$\mathbb{M} \vDash \forall\bar{x}\big[\psi(\bar{x}) \wedge \varphi(\bar{x}, \bar{a}^{s_{k-1}}) \rightarrow \varphi(\bar{x}, \bar{a}^{s_k})\big] \,.$$

Moreover, $(w \smallsetminus \{k - 1\}) \cup \{k\} <_{\text{lex}} w$ implies, by choice of w, that

$$\mathbb{M} \vDash \exists\bar{x}\big[\psi(\bar{x}) \wedge \neg\varphi(\bar{x}, \bar{a}^{s_{k-1}}) \wedge \varphi(\bar{x}, \bar{a}^{s_k})\big] \,.$$

Consequently, it follows by indiscernibility that, for all $i, l \in [s_{k-1}, s_k]$,

$$\mathbb{M} \vDash \forall\bar{x}[\psi(\bar{x}) \wedge \varphi(\bar{x}, \bar{a}^i) \rightarrow \varphi(\bar{x}, \bar{a}^l)] \quad \text{iff} \quad i \leq l \,.$$

In particular, this holds for all $i, l \in I$. □

Lemma 4.7 shows that there are unstable theories without the independence property. Such theories can be characterised as follows.

Definition 4.9. Let T be a theory. A formula $\varphi(\bar{x}, \bar{y})$ has the *strict order property* (with respect to T) if there exists a model $\mathfrak{M} \vDash T$ containing a sequence $(\bar{a}^n)_{n<\omega}$ such that

$$\mathfrak{M} \vDash \exists \bar{x}[\neg\varphi(\bar{x}, \bar{a}^i) \wedge \varphi(\bar{x}, \bar{a}^k)] \quad \text{iff} \quad i < k\,.$$

If some formula has the strict order property with respect to T then we also say that T has the *strict order property.*

Lemma 4.10. *A theory T has the strict order property if and only if there exists a formula $\varphi(\bar{x}, \bar{y})$ such that $\varphi^{\mathbb{M}}$ is a preorder with infinite chains.*

Proof. ($\Leftarrow$) Suppose that $\varphi(\bar{x}, \bar{y})$ defines a preorder with an infinite chain $(\bar{a}^i)_{i\in I}$. By compactness, there exists an infinite ascending $\varphi^{\mathbb{M}}$-chain $(\bar{b}^n)_{n<\omega}$. It follows that

$$\mathbb{M} \vDash \exists \bar{x}[\neg\varphi(\bar{x}, \bar{b}^i) \wedge \varphi(\bar{x}, \bar{b}^k)] \quad \text{iff} \quad i < k\,.$$

($\Rightarrow$) Suppose that there exists a formula $\psi(\bar{x}, \bar{y})$ with the strict order property and let $(\bar{a}^n)_{n<\omega}$ be a sequence with

$$\mathbb{M} \vDash \exists \bar{x}[\neg\psi(\bar{x}, \bar{a}^i) \wedge \psi(\bar{x}, \bar{a}^k)] \quad \text{iff} \quad i < k\,.$$

We set

$$\varphi(\bar{y}, \bar{y}') := \forall \bar{x}[\psi(\bar{x}, \bar{y}) \rightarrow \psi(\bar{x}, \bar{y}')]\,.$$

Clearly, $\varphi^{\mathbb{M}}$ is reflexive and transitive. Furthermore, we have

$$\mathbb{M} \vDash \varphi(\bar{a}^i, \bar{a}^k) \quad \text{iff} \quad i \geq k\,.$$

Hence, $(\bar{a}^n)_{n<\omega}$ is an infinite descending $\varphi^{\mathbb{M}}$-chain. □

Proposition 4.11. *A first-order theory T is unstable if, and only if, it has the independence property or the strict order property.*

Proof. ($\Leftarrow$) If there is a formula φ with the independence property then, according to Lemma 4.6, φ has also the order property and T is unstable.

Similarly, suppose that there exists a formula φ with the strict order property and let $(\bar{a}^n)_{n<\omega}$ be a sequence with

$$\mathbb{M} \vDash \exists \bar{x}[\neg\varphi(\bar{x}, \bar{a}^i) \wedge \varphi(\bar{x}, \bar{a}^k)] \quad \text{iff} \quad i < k\,.$$

Setting

$$\psi(\bar{x}, \bar{y}) := \bar{x} = \bar{y} \vee \exists \bar{z}[\neg\varphi(\bar{z}, \bar{x}) \wedge \varphi(\bar{z}, \bar{y})]$$

it follows that

$$\mathbb{M} \vDash \psi(\bar{a}^i, \bar{a}^k) \quad \text{iff} \quad i \leq k\,.$$

Hence, ψ has the order property and T is unstable.

($\Rightarrow$) Let $\varphi(\bar{x}, \bar{y})$ be a formula with the order property and suppose that $(\bar{a}^n)_{n<\omega}$ and $(\bar{b}^n)_{n<\omega}$ are indiscernible sequences such that

$$\mathbb{M} \vDash \psi(\bar{a}^i, \bar{b}^k) \quad \text{iff} \quad i \leq k\,.$$

By compactness, there are indiscernible sequences $(\bar{a}^i)_{i\in\mathbb{Z}}$ and $(\bar{b}^i)_{i\in\mathbb{Z}}$ such that

$$\mathbb{M} \vDash \psi(\bar{a}^i, \bar{b}^k) \quad \text{iff} \quad i \leq k\,.$$

If φ has the independence property we are done. Hence, suppose otherwise. Since $[\![\psi(\bar{a}^\circ, \bar{b}^i)]\!]_i$ and $[\![\neg\psi(\bar{a}^\circ, \bar{b}^i)]\!]_i$ are both infinite, we can use Lemma 4.8 to construct a formula $\chi(\bar{y}, \bar{y}; \bar{d})$ such that

$$\mathbb{M} \vDash \chi(\bar{b}^i, \bar{b}^k; \bar{d}) \quad \text{iff} \quad i \leq k\,.$$

It follows that

$$\mathbb{M} \vDash \exists \bar{x}[\neg\chi(\bar{x}, \bar{b}^i; \bar{d}) \wedge \chi(\bar{x}, \bar{b}^k; \bar{d})] \quad \text{iff} \quad i < k\,.$$

Consequently, the sequence $(\bar{b}^i \bar{d})_{i<\omega}$ witnesses that $\chi(\bar{x}, \bar{y}; \bar{z})$ has the strict order property. □

Proposition 4.12. *Let $\varphi(\bar{x}, \bar{y})$ be a formula over a set U. The following statements are equivalent:*

(1) *$\varphi(\bar{x}, \bar{y})$ has the order property.*

(2) *There exist an indiscernible sequence $(\bar{a}^i)_{i \in I}$ over U and a tuple $\bar{c}$ such that both the set $[\![\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I}$ and its complement are infinite.*

(3) *There exists an indiscernible sequence $(\bar{a}^i)_{i \in I}$ such that, for every number $m < \omega$, there exists a tuple $\bar{c}$ such that*

$$\big|[\![\varphi(\bar{a}^i; \bar{c})]\!]_{i \in I}\big| > m \quad \text{and} \quad \big|[\![\neg\varphi(\bar{a}^i; \bar{c})]\!]_{i \in I}\big| > m\,.$$

Proof. (1) ⇒ (3) By Proposition 3.6 and compactness, it is sufficient to find, for every $m < \omega$, a tuple $\bar{c}$ and a sequence $(\bar{a}^i)_{i < \omega}$ such that

$$\big|[\![\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I}\big| \geq m \quad \text{and} \quad \big|[\![\neg\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I}\big| \geq m\,.$$

Since φ has the order property there are sequences $(\bar{c}^n)_{n < \omega}$ and $(\bar{d}^n)_{n < \omega}$ such that

$$\mathbb{M} \vDash \varphi(\bar{c}^i, \bar{d}^k) \quad \text{iff} \quad i \leq k\,.$$

Given $m < \omega$ we consider the tuple $\bar{c} := \bar{d}^m$ and the sequence $\bar{a}^i := \bar{c}^i$, $i < \omega$. Then

$$\begin{aligned} &[\![\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I} = [\![\varphi(\bar{c}^i, \bar{d}^m)]\!]_{i \in I} = \{m, m+1, \dots\} \\ \text{and} \quad &[\![\neg\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I} = [\![\neg\varphi(\bar{c}^i, \bar{d}^m)]\!]_{i \in I} = \{0, \dots, m-1\} \end{aligned}$$

contain both at least m elements.

(2) ⇒ (1) Let $\bar{c}$ and $(\bar{a}^i)_{i \in I}$ be given. According to Proposition 4.2, if neither

$$I_0 := [\![\neg\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I} \quad \text{nor} \quad I_1 := [\![\varphi(\bar{a}^i, \bar{c})]\!]_{i \in I}$$

can be written as a finite union of segments then φ has the independence property. By Lemma 4.6, this implies that φ has the order property.

Hence, it remains to consider the case that both I_0 and I_1 are finite unions of segments. Since these sets are both infinite it follows that each contains at least one infinite segment. By taking a suitable subsequence of $(\bar{a}^i)_{i \in I}$ we may assume that both sets consist of a single infinite segment. Reversing the sequence $(\bar{a}^i)_{i \in I}$ if necessary, we may further assume that $I_0 < I_1$.

By compactness it is sufficient to find, for every $m < \omega$, sequences $(\bar{c}^i)_{i<m}$ and $(\bar{d}^i)_{i<m}$ such that

$$\mathbb{M} \vDash \varphi(\bar{c}^i, \bar{d}^k) \quad \text{iff} \quad i \leq k\,.$$

Given $m < \omega$ we pick indices $k_0 < \cdots < k_{m-1}$ in I_0 and $k_m < \cdots < k_{2m-1}$ in I_1. For $i < m$, let π_i be an automorphism with $\pi_i(\bar{a}^{k_j}) = \bar{a}^{k_{j-i}}$ and define

$$\bar{c}^i := \bar{a}^{k_{m-i}} \quad \text{and} \quad \bar{d}^i := \pi_i(\bar{c})\,.$$

For $i, l < m$, it then follows that

$$\begin{aligned} \mathbb{M} \vDash \varphi(\bar{c}^i, \bar{d}^l) \quad &\text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a}^{k_{m-i}}, \pi_l(\bar{c})) \\ &\text{iff} \quad \mathbb{M} \vDash \varphi(\pi_l(\bar{a}^{k_{m-i+l}}), \pi_l(\bar{c})) \\ &\text{iff} \quad \mathbb{M} \vDash \varphi(\bar{a}^{k_{m-i+l}}, \bar{c}) \\ &\text{iff} \quad m - i + l \geq m \\ &\text{iff} \quad i \leq l\,. \end{aligned}$$

(3) ⇒ (2) By Corollary 3.10, we may assume that the order I is dense. Set

$$\Phi := \mathrm{Av}((\bar{a}^i)_i/U) \cup \{\, \varphi(\bar{x}^n; \bar{y}) \leftrightarrow \neg\varphi(\bar{x}^{n+1}; \bar{y}) \mid n < \omega \,\}\,.$$

If Φ is satisfiable, there exists an indiscernible sequence $(\bar{b}^n)_{n<\omega}$ over U and a tuple $\bar{c}$ such that

$$[\![\varphi(\bar{b}^n; \bar{c})]\!]_{n<\omega} = \{\, 2n \mid n < \omega \,\}$$

and $[\![\neg\varphi(\bar{b}^n;\bar{c})]\!]_{n<\omega} = \{\, 2n+1 \mid n<\omega \,\}$.

In particular, both sets are infinite.

Hence, it remains to prove that Φ is satisfiable. Consider a finite subset $\Phi_0 \subseteq \Phi$. Let $n<\omega$ be the maximal number such that Φ_0 contains a formula of the form

$$\varphi(\bar{x}^n;\bar{y}) \leftrightarrow \neg\varphi(\bar{x}^{n+1};\bar{y})\,.$$

By (3), there exists a tuple $\bar{c}$ such that

$$\big|[\![\varphi(\bar{a}^i;\bar{c})]\!]_{i\in I}\big| > n \quad\text{and}\quad \big|[\![\neg\varphi(\bar{a}^i;\bar{c})]\!]_{i\in I}\big| > n\,.$$

If both sets are infinite, we are done. Hence, suppose that one of them is finite. Choose indices $k_0<\dots<k_{n-1}$ in the finite set. As the other set is dense and cofinite, it contains indices $l_0<\dots<l_{n-1}$ such that

$$k_0<l_0<k_1<l_1<\dots<k_{n-1}<l_{n-1}\,.$$

Let K be this set of indices. Then $(\bar{a}^i)_{i\in K}$ and $\bar{c}$ satisfy Φ_0. □

Corollary 4.13. *A first-order theory T is stable if, and only if, for every formula $\varphi(\bar{x})$ with parameters and all indiscernible sequences $(\bar{a}^i)_{i\in I}$ at least one of the sets $[\![\varphi(\bar{a}^i)]\!]_{i\in I}$ and $[\![\neg\varphi(\bar{a}^i)]\!]_{i\in I}$ is finite.*

Corollary 4.14. *Let T be a stable theory and $(\bar{a}^i)_{i\in I}$ an indiscernible sequence over U. For every set $C\subseteq\mathbb{M}$, the set*

$$\mathrm{Av}_1((\bar{a}^i)_i/C) := \{\, \varphi(\bar{x}) \mid \varphi \text{ a formula over } C \text{ such that } [\![\varphi(\bar{a}^i)]\!]_{i\in I} \text{ is cofinite} \,\}$$

forms a complete type over C.

Proof. By the preceding corollary, we have

$$\begin{aligned}\varphi(\bar{x}) \in \mathrm{Av}_1((\bar{a}^i)_i/C) \quad &\text{iff}\quad [\![\varphi(\bar{a}^i)]\!]_{i\in I} \text{ is cofinite}\\ &\text{iff}\quad [\![\neg\varphi(\bar{a}^i)]\!]_{i\in I} \text{ is finite}\\ &\text{iff}\quad [\![\neg\varphi(\bar{a}^i)]\!]_{i\in I} \text{ is not cofinite}\\ &\text{iff}\quad \neg\varphi(\bar{x}) \notin \mathrm{Av}_1((\bar{a}^i)_i/C)\,.\end{aligned}$$

Hence, it remains to prove that $\mathrm{Av}_1((\bar{a}_i)_i/C)$ is consistent with T. Let $\varphi_0, \dots, \varphi_n \in \mathrm{Av}_1((\bar{a}_i)_i/C)$. Then

$$[\![\varphi_0(\bar{a}_i)]\!]_{i \in I}, \dots, [\![\varphi_n(\bar{a}_i)]\!]_{i \in I} \text{ are cofinite.}$$

Hence, so is

$$[\![\varphi_0(\bar{a}_i) \wedge \dots \wedge \varphi_n(\bar{a}_i)]\!]_{i \in I} = [\![\varphi_0(\bar{a}_i)]\!]_{i \in I} \cap \dots \cap [\![\varphi_n(\bar{a}_i)]\!]_{i \in I} .$$

Fixing some index i in this set, it follows that

$$\mathbb{M} \vDash \varphi_0(\bar{a}_i) \wedge \dots \wedge \varphi_n(\bar{a}_i) .$$

Consequently, every finite subset of $\mathrm{Av}_1((\bar{a}_i)_i/C)$ is satisfiable. □

E6. *Functors and embeddings*

1. *Local functors*

In this section we consider functors preserving back-and-forth equivalence. Recall that $\mathrm{Sub}_\kappa(\mathfrak{M})$ denotes the class of all κ-generated substructures of $\mathfrak{M}$, and that a class $\mathcal{K}$ is κ-hereditary if $\mathfrak{M} \in \mathcal{K}$ implies $\mathrm{Sub}_\kappa(\mathfrak{M}) \subseteq \mathcal{K}$.

Definition 1.1. Let $\mathcal{K}$ be a class of Σ-structures. We denote the subcategory of $\mathfrak{Emb}(\Sigma)$ induced by $\mathcal{K}$ by $\mathfrak{Emb}(\mathcal{K})$.

Below we will show that functors preserving direct limits also preserve ∞-equivalence. We start by giving an alternative characterisation of such functors.

Definition 1.2. A functor $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ is *κ-local* if, for every embedding $f : \mathfrak{B} \to F(\mathfrak{A})$ where $\mathfrak{B} \in \mathcal{K}$ is κ-generated and $\mathfrak{A} \in \mathcal{C}$, there exists an embedding $g : \mathfrak{C} \to \mathfrak{A}$ where $\mathfrak{C} \in \mathcal{C}$ is κ-generated such that the map f factors through $F(g)$.

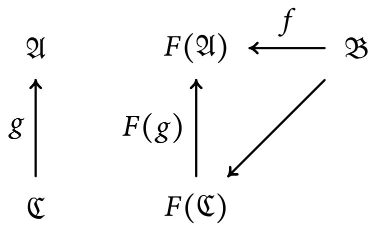

Example. The following operations are $\aleph_0$-local functors.

(a) The function mapping a ring $\mathfrak{R}$ to the polynomial ring $\mathfrak{R}[x]$.

(b) The function mapping an integral domain $\mathfrak{R}$ to its quotient field.

(c) The function mapping a set X to the free group generated by X.

(d) The function mapping a structure $\mathfrak{M}$ to the structure $\mathrm{HF}(\mathfrak{M})$ consisting of all hereditary finite sets with elements from $\mathfrak{M}$.

Lemma 1.3. *If $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{D})$ and $G : \mathfrak{Emb}(\mathcal{D}) \to \mathfrak{Emb}(\mathcal{K})$ are κ-local then so is $G \circ F$.*

Exercise 1.1. Prove the preceding lemma.

As a further, more involved example we show that quantifier-free interpretations are $\aleph_0$-local functors. While every interpretation is local in an intuitive sense we need the restriction to quantifier-free formulae to prove that the interpretation is a functor.

Lemma 1.4. *Every $\mathrm{QF}_{\infty\aleph_0}$-interpretation $\mathcal{I} : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ is an $\aleph_0$-local functor.*

Proof. First, we show that quantifier-free interpretations are functors. Suppose that

$$\mathcal{I} = \langle \alpha, (\delta_s)_{s\in S}, (\varepsilon_s)_{s\in S}, (\varphi_\xi)_{\xi\in\Sigma} \rangle$$

is quantifier-free and let $h : \mathfrak{A} \to \mathfrak{B}$ be an embedding. For $\bar{a} \in \delta_s^{\mathfrak{A}}$, we denote by $[\bar{a}]_s$ the element encoded by $\bar{a}$. We define $\mathcal{I}(h)$ by

$$\mathcal{I}(h)[\bar{a}]_s := [h(\bar{a})]_s \,.$$

Since embeddings preserve quantifier-free formulae it follows that this mapping is a well-defined embedding $\mathcal{I}(h) : \mathcal{I}(\mathfrak{A}) \to \mathcal{I}(\mathfrak{B})$. Obviously, we have $\mathcal{I}(f \circ g) = \mathcal{I}(f) \circ \mathcal{I}(g)$. Consequently, $\mathcal{I}$ is a functor.

To show that it is $\aleph_0$-local let $X \subseteq \mathcal{I}(\mathfrak{A})$ be finite. For each equivalence class $[\bar{a}]_s \in X$, fix a representative $\bar{a}$ and let A_0 be the set of these representatives. Then A_0 is finite and we have $X \subseteq \mathcal{I}(\langle\!\langle A_0 \rangle\!\rangle_{\mathfrak{A}})$. Note that $\mathcal{I}(\langle\!\langle A_0 \rangle\!\rangle_{\mathfrak{A}})$ is defined since $\mathcal{I}$ is quantifier-free. □

Local functors can be characterised in purely category-theoretical terms as those functors that preserve direct limits.

Theorem 1.5. *Let $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ be a functor where the classes $\mathcal{C}$ and $\mathcal{K}$ are κ-hereditary. The functor F is κ-local if and only if it preserves κ-filtered colimits.*

Proof. ($\Leftarrow$) Let $f : \mathfrak{B} \to F(\mathfrak{A})$ be an embedding where $\mathfrak{B} \in \mathcal{K}$ is κ-generated. According to Lemma ?? we can write $\mathfrak{A} = \varinjlim D$ where $D : \mathcal{I} \to \mathrm{Sub}_\kappa(\mathfrak{A})$ is the canonical κ-filtered diagram of all κ-generated substructures. The corresponding cocone μ from D to $\mathfrak{A}$ consists of all inclusion maps $\mu_\mathfrak{i} : D(\mathfrak{i}) \to \mathfrak{A}$. Since F preserves κ-direct limits we have $F(\mathfrak{A}) = \varinjlim(F \circ D)$ and the corresponding cone is $F[\mu]$.

To find the desired embedding $g : \mathfrak{C} \to \mathfrak{A}$ we fix a set $X \subseteq B$ of size $|X| < \kappa$ generating $\mathfrak{B}$. For each $x \in X$, we choose an index $\mathfrak{i}_x \in \mathcal{I}$ such that $f(x) \in \operatorname{rng} F(\mu_{\mathfrak{i}_x})$. Since I is κ-filtered there is some index $\mathfrak{k} \in I$ and morphisms $h_x : \mathfrak{i}_x \to \mathfrak{k}$, for all x. Hence, we have

$$f[X] \subseteq \operatorname{rng} F(\mu_\mathfrak{k}) ,$$

which, by Lemma B1.2.8, implies that

$$\begin{aligned} \operatorname{rng} f = f[\langle\!\langle X \rangle\!\rangle_\mathfrak{B}] &= \langle\!\langle f[X] \rangle\!\rangle_{F(\mathfrak{A})} \\ &\subseteq \langle\!\langle \operatorname{rng} F(\mu_\mathfrak{k}) \rangle\!\rangle_{F(\mathfrak{A})} = \operatorname{rng} F(\mu_\mathfrak{k}) . \end{aligned}$$

Since f and $F(\mu_\mathfrak{k})$ are injective and $\operatorname{rng} f \subseteq \operatorname{rng} F(\mu_\mathfrak{k})$ we can define a function $g : B \to F(D(\mathfrak{k}))$ by $g := F(\mu_\mathfrak{k})^{-1} \circ f$. Since f and $F(\mu_\mathfrak{k})$ preserve all quantifier-free formulae so does g. Hence, g is an embedding. Furthermore, we have $F(\mu_\mathfrak{k}) \circ g = f$.

($\Rightarrow$) Let $D : \mathcal{I} \to \mathfrak{Emb}(\mathcal{C})$ be a κ-filtered diagram with $\mathfrak{A} := \varinjlim D$, and suppose that μ is a limiting cocone from D to $\mathfrak{A}$. We claim that $\varinjlim(F \circ D) = F(\mathfrak{A})$. Let $\mathfrak{D} := \varinjlim(F \circ D)$ and let λ be a limiting cocone from $F \circ D$ to $\mathfrak{D}$. Since $F[\mu]$ is a cocone from $F \circ D$ to $F(\mathfrak{A})$ it follows that there exists an embedding $h : \mathfrak{D} \to F(\mathfrak{A})$ with $h * \lambda = F[\mu]$.

We only have to show that h is surjective. Fix $c \in F(\mathfrak{A})$. There exists some substructure $\mathfrak{B} \in \mathrm{Sub}_\kappa(F(\mathfrak{A}))$ with $c \in B$. Let $j : \mathfrak{B} \to F(\mathfrak{A})$ be the inclusion map. Since F is κ-local we can find a κ-generated structure $\mathfrak{C} \in \mathcal{C}$ and an embedding $g : \mathfrak{C} \to \mathfrak{A}$ such that $j = F(g) \circ j_0$, for some

$j_0 : \mathfrak{B} \to F(\mathfrak{C})$. In the same way as above we can show that there is some index $\mathfrak{k} \in \mathcal{I}$ and an embedding $g_0 : \mathfrak{C} \to D(\mathfrak{k})$ with $g = \mu_{\mathfrak{k}} \circ g_0$.

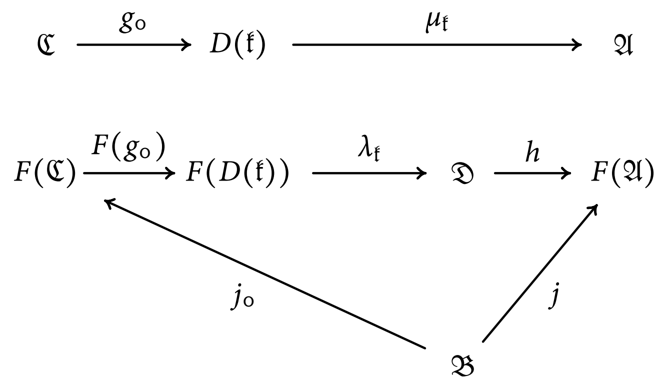

Since $a = j(a) = (h \circ \lambda_{\mathfrak{k}} \circ F(g_0) \circ j_0)(a)$ it follows that $a \in \operatorname{rng} h$. □

Let us show that local functors preserve back-and-forth equivalences.

Definition 1.6. Suppose that $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ is a functor where the classes $\mathcal{C}$ and $\mathcal{K}$ are κ-hereditary. Let $p = \bar{a} \mapsto \bar{b} \in \operatorname{pIso}(\mathfrak{A}, \mathfrak{B})$ be a partial isomorphism between $\mathfrak{A}, \mathfrak{B} \in \mathcal{C}$ and let $\pi : \mathfrak{A}_0 \to \mathfrak{B}_0$ be the unique isomorphism extending p, where $\mathfrak{A}_0 := \langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{A}}$ and $\mathfrak{B}_0 := \langle\!\langle \bar{b} \rangle\!\rangle_{\mathfrak{B}}$ are the structures induced by, respectively, the domain and range of p. Let $i : \mathfrak{A}_0 \to \mathfrak{A}$ and $j : \mathfrak{B}_0 \to \mathfrak{B}$ be the corresponding inclusion maps and suppose that $F(\pi) = \bar{a}' \mapsto \bar{b}'$. We define

$$p^F := F(i)(\bar{a}') \mapsto F(j)(\bar{b}') .$$

Proposition 1.7. *Let $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ be an $\aleph_0$-local functor where the classes $\mathcal{C}$ and $\mathcal{K}$ are $\aleph_0$-hereditary.*

$$p \in I^{\aleph_0}_{\omega\alpha}(\mathfrak{A}, \mathfrak{B}) \quad \textit{implies} \quad p^F \in I_\alpha(F(\mathfrak{A}), F(\mathfrak{B})) .$$

Proof. The claim follows by induction on α. Let $p := \bar{a} \mapsto \bar{b} \in I^{\aleph_0}_{\omega\alpha}(\mathfrak{A}, \mathfrak{B})$, set $\mathfrak{A}_0 := \langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{A}}$ and $\mathfrak{B}_0 := \langle\!\langle \bar{b} \rangle\!\rangle_{\mathfrak{B}}$, and let $\pi : \mathfrak{A}_0 \to \mathfrak{B}_0$ be the isomorphism extending p. Let $i : \mathfrak{A}_0 \to \mathfrak{A}$ and $j : \mathfrak{B}_0 \to \mathfrak{B}$ be the corresponding inclusion maps and suppose that $F(\pi) = \bar{a}' \mapsto \bar{b}'$.

For $\alpha = 0$, we have to check that $F(p)$ is a partial isomorphism. Since $F(i)$, $F(j)$, and $F(\pi)$ are embeddings it follows, for every quantifier-free formula $\varphi(\bar{x})$, that

$$\begin{aligned} F(\mathfrak{A}) \vDash \varphi(F(i)(\bar{a}')) \quad &\text{iff} \quad F(\mathfrak{A}_0) \vDash \varphi(\bar{a}') \\ &\text{iff} \quad F(\mathfrak{B}_0) \vDash \varphi(\bar{b}') \\ &\text{iff} \quad F(\mathfrak{B}) \vDash \varphi(F(j)(\bar{b}')) \,. \end{aligned}$$

If α is a limit ordinal then the claim follows immediately by inductive hypothesis. Hence, suppose that $\alpha = \beta + 1$. By symmetry, we only need to check the forth property. Fix $c \in F(\mathfrak{A})$. Since F is $\aleph_0$-local there exist a finitely generated structure $\mathfrak{C}$ and an embedding $g : \mathfrak{C} \to \mathfrak{A}$ such that the inclusion $h : \langle\!\langle c \rangle\!\rangle_{F(\mathfrak{A})} \to F(\mathfrak{A})$ factors through $F(g)$, i.e., $h = F(g) \circ h_0$. Choose a finite tuple $\bar{e}_0$ of generators of $\mathfrak{C}$ and set $\bar{e} := g(\bar{e}_0)$ and $\mathfrak{A}_1 := \langle\!\langle \bar{a}\bar{e} \rangle\!\rangle_{\mathfrak{A}}$. Since $p = \bar{a} \mapsto \bar{b} \in \mathrm{pIso}^{\aleph_0}_{\omega(\beta+1)}(\mathfrak{A}, \mathfrak{B})$ we can find some $\bar{f} \subseteq B$ with $q := \bar{a}\bar{e} \mapsto \bar{b}\bar{f} \in \mathrm{pIso}^{\aleph_0}_{\omega\beta}(\mathfrak{A}, \mathfrak{B})$. Set $\mathfrak{B}_1 := \langle\!\langle \bar{b}\bar{f} \rangle\!\rangle_{\mathfrak{B}}$ and let $\rho : \mathfrak{A}_1 \to \mathfrak{B}_1$ be the unique isomorphism extending q. We claim that q^F is an extension of p^F with $c \in \operatorname{dom} q^F$.

Let i_0, i_1, j_0, j_1, g_0 be the inclusion maps as depicted in the following diagram

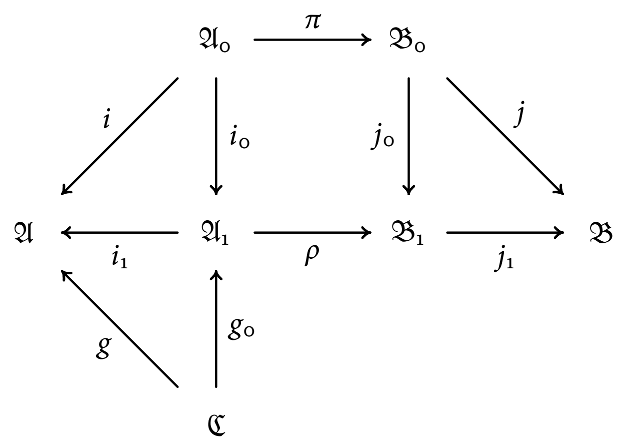

Applying F to this diagram we obtain

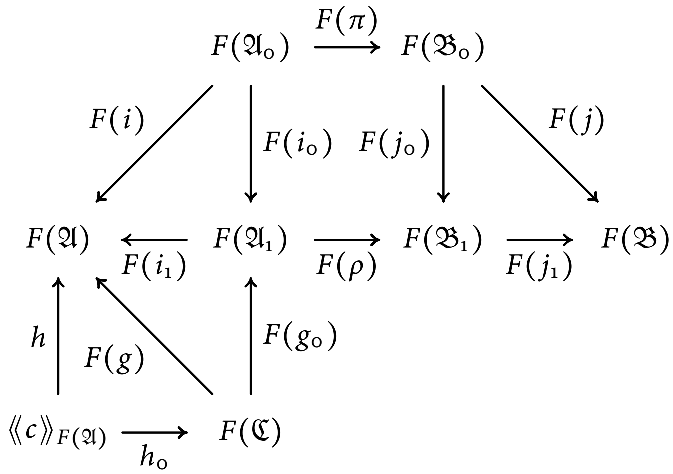

First, let us show that $c \in \operatorname{dom} q^F$. We have

$$c = h(c) = (F(i_1) \circ F(g_0) \circ h_0)(c)$$

which implies that $c \in \operatorname{rng} F(i_1) = \operatorname{dom} q^F$.

It remains to prove that $p^F \subseteq q^F$. Let $x \in \operatorname{dom} p^F$. Then $x = F(i)(a'_l)$, for some l. Setting $w := F(i_0)(a'_l)$ we have

$$F(i_1)(w) = (F(i_1) \circ F(i_0))(a'_l) = F(i)(a'_l) = x\,.$$

It follows that

$$\begin{aligned} q^F(x) &= (F(j_1) \circ F(\rho))(w) \\ &= (F(j_1) \circ F(\rho) \circ F(i_0))(a'_l) \\ &= (F(j_1) \circ F(j_0) \circ F(\pi))(a'_l) \\ &= (F(j) \circ F(\pi))(a'_l) = p^F(x)\,. \end{aligned}$$

□

Corollary 1.8. *Let $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ be an $\aleph_0$-local functor where the classes $\mathcal{C}$ and $\mathcal{K}$ are $\aleph_0$-hereditary. For all $\mathfrak{A}, \mathfrak{B}$, we have*

$$\mathfrak{A} \cong_{\omega\alpha} \mathfrak{B} \quad \textit{implies} \quad F(\mathfrak{A}) \cong_{\alpha} F(\mathfrak{B})\,.$$

In particular,

$$\mathfrak{A} \cong_{\infty} \mathfrak{B} \quad \textit{implies} \quad F(\mathfrak{A}) \cong_{\infty} F(\mathfrak{B})\,.$$

We conclude this section by showing that local functors are compatible with universal theories.

Definition 1.9. Let $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ be a functor and L a logic. The *L-theory* of F is the set

$$\mathrm{Th}_L(F) := \{ \varphi \in L \mid F(\mathfrak{A}) \vDash \varphi \text{ for all } \mathfrak{A} \in \mathcal{C} \}.$$

Lemma 1.10. *Let $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ be an $\aleph_0$-local functor where the classes $\mathcal{C}$ and $\mathcal{K}$ are $\aleph_0$-hereditary. If $\mathfrak{U} \in \mathcal{C}$ is $\aleph_0$-universal then*

$$\mathrm{Th}_{\forall_{\infty\aleph_0}}(F(\mathfrak{U})) = \mathrm{Th}_{\forall_{\infty\aleph_0}}(F).$$

Proof. $(\supseteq)$ follows immediately from the definitions.

$(\subseteq)$ We prove by induction on $\psi(\bar{x}) \in \forall_{\infty\aleph_0}$ that

$$F(\mathfrak{U}) \vDash \psi(\bar{c}), \quad \text{for all } \bar{c} \subseteq F(\mathfrak{U}),$$

implies that

$$F(\mathfrak{A}) \vDash \psi(\bar{a}), \quad \text{for all } \mathfrak{A} \in \mathcal{C} \text{ and every } \bar{a} \subseteq F(\mathfrak{A}).$$

First, suppose that ψ is quantifier-free. Let $\mathfrak{A} \in \mathcal{C}$ and $\bar{a} \subseteq F(\mathfrak{A})$. We have to show that $F(\mathfrak{A}) \vDash \psi(\bar{a})$. Since F is $\aleph_0$-local we can find a finitely generated substructure $\mathfrak{A}_0 \subseteq \mathfrak{A}$ with $\bar{a} \subseteq F(\mathfrak{A}_0)$. Since $\mathfrak{U}$ is $\aleph_0$-universal there exists an embedding $f : \mathfrak{A}_0 \to \mathfrak{U}$. We set $\bar{b} := F(f)(\bar{a})$. By assumption $F(\mathfrak{U}) \vDash \psi(\bar{b})$. Since ψ is quantifier-free and $F(f)$ is an embedding it follows that $F(\mathfrak{A}_0) \vDash \psi(\bar{a})$. Hence, $F(\mathfrak{A}) \vDash \psi(\bar{a})$.

For the inductive step, we have to distinguish three cases. Either

$$\psi(\bar{x}) = \bigwedge \Psi, \quad \text{or} \quad \psi(\bar{x}) = \bigvee \Psi, \quad \text{or} \quad \psi(\bar{x}) = \forall y \vartheta(\bar{x}, y).$$

In each of these cases the claim follows directly from the inductive hypothesis. □

2. *Word constructions*

Local functors can be characterised in terms of a certain family of comorphisms called *word constructions.* Instead of defining these operations as a single, complex construction we will introduce several simple operations which, when combined with first-order interpretations, yield the required expressive power.

We start with the main ingredient in a word construction, the so-called *term-algebra* operation.

Definition 2.1. Let Γ be a functional S-sorted signature and Σ a relational one that is S_0-sorted for some $S_0 \subseteq S$. The *Γ-term algebra* $\mathcal{T}[\Gamma, \mathfrak{A}]$ over a Σ-structure $\mathfrak{A}$ is the $T[\Gamma, S_0]$-sorted structure whose universe $T[\Gamma, A]$ consists of all Γ-terms over A. Every element $t(\bar{a}) \in T[\Gamma, A]$ has sort $t(\bar{s})$, where $\bar{s}$ are the sorts of $\bar{a}$. For each relation symbol $R \in \Sigma$, we have the relation

$$R^{\mathcal{T}[\Gamma,\mathfrak{A}]} = R^{\mathfrak{A}},$$

and, for each n-ary function symbol $f \in \Gamma$, we have an n-ary function defined by

$$f^{\mathcal{T}[\Gamma,\mathfrak{A}]}(t_0, \dots, t_{n-1}) := f t_0 \dots t_{n-1}.$$

Example. Let us give two simple examples showing the versatility of the term algebra operation in conjunction with a first-order interpretation.

(a) First, we interpret the product $\mathfrak{A} \times \mathfrak{A}$ in the structure $\mathcal{T}[\{f\}, \mathfrak{A}]$ where f is a binary function symbol. When we encode a pair $\langle a, b \rangle \in A \times A$ by the term $f(a, b)$, we can define the universe by the formula

$$\delta(x) := \text{“}x = f(a, b) \text{ for some } a, b \in A.\text{”}$$

Then we define each relation R by

$$\varphi_R(\bar{x}) := \text{“}x_i = f(a_i, b_i) \text{ for some } a_i, b_i \in A \text{ such that } \bar{a}, \bar{b} \in R.\text{”}$$

(b) Similarly, we can interpret the disjoint union $\mathfrak{A} \cup \mathfrak{A}$ in the structure $\mathcal{T}[\{f\}, \mathfrak{A}]$ where f is a unary function symbol. The universe is the set

$$A \cup \{ f(a) \mid a \in A \}$$

which is obviously definable in $\mathcal{T}[\{f\}, \mathfrak{A}]$. We can define the relations R by

$$\varphi_R(\bar{x}) := \text{“Either } \bar{x} = \bar{a} \text{ or } \bar{x} = f(\bar{a}), \text{ for some } \bar{a} \in R.\text{”}$$

Lemma 2.2. *Let Σ a relational signature and Γ a functional one. The Γ-term-algebra operation*

$$\mathcal{T}[\Gamma, -] : \mathfrak{Emb}(\Sigma) \to \mathfrak{Emb}(\Sigma \cup \Gamma)$$

is an $\aleph_0$-local functor.

Proof. First, let us show that it is a functor. Let $h : \mathfrak{A} \to \mathfrak{B}$ be an embedding of Σ-structures. We obtain an embedding

$$\mathcal{T}[\Gamma, h] : \mathcal{T}[\Gamma, \mathfrak{A}] \to \mathcal{T}[\Gamma, \mathfrak{B}]$$

by setting

$$\mathcal{T}[\Gamma, h](t(\bar{a})) := t(h(\bar{a})) .$$

To prove that $\mathcal{T}[\Gamma, -]$ is $\aleph_0$-local suppose that $X \subseteq T[\Gamma, A]$ is finite. Then we have $X \subseteq T[\Gamma, A_0] = \langle\!\langle A_0 \rangle\!\rangle_{\mathcal{T}[\Gamma, \mathfrak{A}]}$ where the set

$$A_0 := \bigcup \{ \bar{a} \mid t(\bar{a}) \in X \}$$

is finite. □

It follows from the results of the previous section that $\mathcal{T}[\Gamma, -]$ preserves ∞-equivalence. The next lemma gives a more precise statement.

Lemma 2.3. *Suppose that Σ is a relational signature, Γ a functional one, and κ an infinite cardinal. For each $\mathrm{FO}_{\kappa\aleph_0}$-formula $\varphi(x_0, \dots, x_{n-1})$ and all terms $t_i(\bar{x}^i) \in T^{<\omega}[\Gamma]$, for $i < n$, we can construct an $\mathrm{FO}_{\kappa\aleph_0}$-formula $\varphi_{t_0 \dots t_{n-1}}(\bar{x}^0, \dots, \bar{x}^{n-1})$ such that*

$$\mathcal{T}[\Gamma, \mathfrak{A}] \vDash \varphi(t_0(\bar{a}_0), \dots, t_{n-1}(\bar{a}_{n-1}))$$

iff
$$\mathfrak{A} \vDash \varphi_{t_0 \dots t_{n-1}}(\bar{a}_0, \dots, \bar{a}_{n-1}) .$$

Proof. W.l.o.g. we may assume that φ is term reduced. We construct $\varphi_{\bar{t}}$ inductively. First, suppose that φ is an atomic formula. If $\varphi = R\bar{x}$ with $R \in \Sigma$ then we can set

$$(R\bar{x})_{\bar{t}} := \begin{cases} Rx^0 \dots x^{n-1} & \text{if } t_i = x \text{ for all } i\,, \\ \text{false} & \text{otherwise}\,. \end{cases}$$

For $\varphi = x = y$ we set

$$(x = y)_{st} := \begin{cases} \bigwedge_i x_i = y_i & \text{if } s = t\,, \\ \text{false} & \text{otherwise}\,. \end{cases}$$

Finally, if $\varphi = f\bar{x} = y$ then we define

$$(f\bar{x} = y)_{\bar{s}t} := \begin{cases} \bigwedge_{i,j} x^i_j = y^i_j & \text{if } f\bar{s} = t\,, \\ \text{false} & \text{otherwise}\,, \end{cases}$$

where $s_i = s_i(\bar{x}^i)$ and $t = t(\bar{y}^0, \dots, \bar{y}^{n-1})$. Boolean operations are unchanged:

$$(\neg\varphi)_{\bar{t}} := \neg\varphi_{\bar{t}} \quad \text{and} \quad \left(\bigwedge \Phi\right)_{\bar{t}} := \bigwedge \{\, \varphi_{\bar{t}} \mid \varphi \in \Phi \,\}\,.$$

For a quantifier over a variable y of sort $s \in T[\Gamma, S_0]$, we have

$$(\exists y \varphi(\bar{x}, y))_{\bar{t}} := \exists \bar{y} \varphi_{\bar{t}s}(\bar{x}^0, \dots, \bar{x}^{n-1}, \bar{y})\,. \qquad \square$$

The term-algebra operation creates structures with many sorts. To reduce the number of sorts we employ a second operation that merges several sorts into a single one. Recall that with every morphism $\langle S, \Sigma\rangle \to \langle T, \Gamma\rangle$ of $\mathfrak{Sig}$ we have associated a reduct mapping $\mathfrak{Str}[\Gamma] \to \mathfrak{Str}[\Sigma]$. For relational signatures we can also define a mapping $\mathfrak{Str}[\Sigma] \to \mathfrak{Str}[\Gamma]$ in the other direction.

Definition 2.4. Let $\alpha = \langle \chi, \mu\rangle : \langle S, \Sigma\rangle \to \langle T, \Gamma\rangle$ be a morphism of $\mathfrak{Sig}$ where the signatures Σ and Γ are relational. The *inverse α-reduct* of a Σ-structure $\mathfrak{A}$ is the Γ-structure $\mathfrak{A}^\alpha$ where the domain of sort $t \in T$ is

$$A_t^\alpha := \dot{\bigcup} \{ A_s \mid s \in \chi^{-1}(t) \},$$

and, for each relation symbol $R \in \Gamma$, we have

$$R^{\mathfrak{A}^\alpha} := \dot{\bigcup} \{ Q^{\mathfrak{A}} \mid Q \in \mu^{-1}(R) \}.$$

Remark. We have defined inverse reducts only for relational signatures in order to avoid the complications arising from the fact that we require functions to be total. For instance, if $\mathfrak{V} = \langle V, K, +, \cdot\rangle$ is a $\{v, s\}$-sorted vector space and α maps both sorts to the same value, then we get problems defining $\mathfrak{V}^\alpha$ since the operation $+$ is a function $V \times V \to V$ and not a function $(V \cup K) \times (V \cup K) \to V \cup K$.

Lemma 2.5. *Let α be a morphism of $\mathfrak{Sig}$. The operation $\mathfrak{A} \mapsto \mathfrak{A}^\alpha$ is an $\aleph_0$-local functor.*

Proof. Clearly the operation is $\aleph_0$-local: for every finite subset $X \subseteq \mathfrak{A}^\alpha$ we have $X \subseteq (\langle\!\langle X\rangle\!\rangle_{\mathfrak{A}})^\alpha$. It remains to show that it is a functor. Let $h : \mathfrak{A} \to \mathfrak{B}$ be an embedding. We define

$$h^\alpha : \mathfrak{A}^\alpha \to \mathfrak{B}^\alpha \quad \text{by setting} \quad h^\alpha(a) := h(a).$$

To show that this function is an embedding suppose that $\bar{a} \in R^{\mathfrak{A}^\alpha}$. Then there is some relation $Q \in \alpha^{-1}(R)$ with $\bar{a} \in Q^{\mathfrak{A}}$. Hence, $h(\bar{a}) \in Q^{\mathfrak{B}} \subseteq R^{\mathfrak{B}^\alpha}$. $\square$

It follows that inverse reducts preserve $\mathrm{FO}_{\infty\aleph_0}$-equivalence. The next lemma states that they also preserve $\mathrm{FO}_{\kappa\aleph_0}$-equivalence for sufficiently large cardinals κ.

Lemma 2.6. *Let $\alpha = \langle \chi, \mu \rangle : \langle S, \Sigma \rangle \to \langle T, \Gamma \rangle$ be a morphism of $\mathfrak{Sig}$ where the signatures Σ and Γ are relational, and let κ be an infinite cardinal such that*

$$|\chi^{-1}(t)| < \kappa \quad \text{and} \quad |\mu^{-1}(R)| < \kappa\,, \quad \text{for all} \quad t \in T \text{ and } R \in \Gamma\,.$$

For every formula $\varphi(\bar{x}) \in \mathrm{FO}_{\kappa\aleph_0}[\Sigma]$ where x_i is of sort t_i and for all sorts $s_i \in \chi^{-1}(t_i)$, there exists a formula $\varphi^{\alpha}_{\bar{s}}(\bar{x}) \in \mathrm{FO}_{\kappa\aleph_0}[\Gamma]$ such that

$$\mathfrak{A}^{\alpha} \vDash \varphi(\bar{a}) \quad \text{iff} \quad \mathfrak{A} \vDash \varphi^{\alpha}_{\bar{s}}(\bar{a})\,,$$

for every Σ-structure $\mathfrak{A}$ and all $a_i \in A_{s_i}$.

Proof. We construct $\varphi^{\alpha}_{\bar{s}}$ by induction on φ. For atomic formulae we set

$$(x_0 = x_1)^{\alpha}_{\bar{s}} := x_0 = x_1 \quad \text{and} \quad (R\bar{x})^{\alpha}_{\bar{s}} := \bigvee \{\, Q \in \mu^{-1}(R) \mid Q\bar{x} \,\}$$

(where we consider x_i now to be of sort s_i). Boolean operations remain unchanged:

$$(\neg\varphi)^{\alpha}_{\bar{s}} := \neg\varphi^{\alpha}_{\bar{s}} \quad \text{and} \quad (\bigwedge \Phi)^{\alpha}_{\bar{s}} := \bigwedge \{\, \varphi^{\alpha}_{\bar{s}} \mid \varphi \in \Phi \,\}\,.$$

A quantifier with a variable y of sort $t \in T$ is replaced by a disjunction over all sorts $r \in \chi^{-1}(t)$

$$(\exists y\varphi)^{\alpha}_{\bar{s}} := \bigvee \{\, \exists y \varphi^{\alpha}_{\bar{s}r} \mid r \in \chi^{-1}(t) \,\}\,. \qquad \square$$

We obtain an alternative characterisation of $\aleph_0$-local functors by combining these two operations with quantifier-free interpretations.

Definition 2.7. (a) Let Σ be a signature and let Σ_{rel} be the signature obtained from Σ by replacing every function symbol f of type $\bar{s} \to t$ by a relation symbol R_f of type $\bar{s}t$. The *relational variant* of a Σ-structure $\mathfrak{M}$ is the Σ_{rel}-structure $\mathcal{R}(\mathfrak{M})$ obtained from $\mathfrak{M}$ by replacing every function f by its graph.

(b) A κ-*word construction* is an operation of the form

$$F = \mathcal{I} \circ \mathcal{S} \circ \mathcal{R} \circ \mathcal{T} \circ \mathcal{R} ,$$

where $\mathcal{I}$ is a $\mathrm{QF}_{\kappa\aleph_0}$-interpretation, $\mathcal{R}$ is the operation defined in (a), $\mathcal{S}$ is an inverse reduct, and $\mathcal{T}$ is a Γ-term-algebra operation where $|\Gamma| < \kappa$.

Remark. Note that $\mathcal{R}$ is a quantifier-free first-order interpretation.

Theorem 2.8. *Let $\mathcal{C}$ be an $\aleph_0$-hereditary class of Σ-structures and $\mathcal{K}$ a class of Γ-structures. Suppose that κ is a cardinal such that*

$$\kappa > 2^{|\Sigma| \oplus \aleph_0} \quad \textit{and} \quad \kappa > |F(\mathfrak{C})| , \quad \textit{for all finitely generated } \mathfrak{C} \in \mathcal{C} .$$

A mapping $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ is an $\aleph_0$-local functor if and only if it is an κ-word construction.

Proof. $(\Leftarrow)$ We have already seen that all operations a word construction is built up from are $\aleph_0$-local functors. Since $\aleph_0$-local functors are closed under composition the claim follows.

$(\Rightarrow)$ We have to express F as composition

$$F = \mathcal{I} \circ \mathcal{S} \circ \mathcal{R} \circ \mathcal{T} \circ \mathcal{R} .$$

To define $\mathcal{T}$ we use Theorem 1.5 which tells us that F preserves direct limits. Let $\mathcal{D} : I \to \mathrm{Sub}_{\aleph_0}(\mathfrak{A})$ be the canonical diagram with limit $\varinjlim \mathcal{D} = \mathfrak{A}$. We are looking for an operation mapping $\mathfrak{A}$ to $\varinjlim (F \circ \mathcal{D})$.

Fix an enumeration $(\mathfrak{C}_\alpha)_{\alpha<\lambda}$ of $\bigcup_{\mathfrak{A} \in \mathcal{K}} \mathrm{Sub}_{\aleph_0}(\mathfrak{A})$. Note that each structure $\mathfrak{C}_\alpha$ has at most $|\Sigma| \oplus \aleph_0$ elements. Hence, there are at most $2^{|\Sigma| \oplus \aleph_0}$ of them and we have $\lambda \leq 2^{|\Sigma| \oplus \aleph_0} < \kappa$.

For each $\alpha < \lambda$, we choose a finite tuple $\bar{c}_\alpha \subseteq C_\alpha$ generating $\mathfrak{C}_\alpha$. Set

$$\Xi := \{ f^\alpha_b \mid \alpha < \lambda,\ b \in F(\mathfrak{C}_\alpha) \} ,$$

where f_b^α is a new function symbol of arity $|\bar{c}_\alpha|$. Note that $|\Xi| < \kappa$ since $\lambda < \kappa$ and $|F(\mathfrak{C}_\alpha)| < \kappa$, for all α. For $\mathcal{T}$ we choose the Ξ-term-algebra operation $\mathfrak{A} \mapsto \mathcal{T}[\Xi, \mathfrak{A}]$. The inverse reduct $\mathcal{S}$ maps each element to the correct sort.

The main work is done by the interpretation $\mathcal{I}$. It creates the structures $F(\mathfrak{C}_\alpha)$ and pastes them together. The domain formula $\delta(x)$ states that x is a term of the form $f_b^\alpha(\bar{a})$, for some $\alpha < \lambda$ and $b \in F(\mathfrak{C}_\alpha)$, such that the substructure generated by $\bar{a}$ is isomorphic to $\mathfrak{C}_\alpha$. Each relation $R \in \Gamma$ can be defined by a formula $\varphi_R(\bar{x})$ stating that $x_i = f_{b_i}^{\alpha_i}(\bar{a})$ and the tuple $\bar{b}$ is in the relation $R^{F(\mathfrak{C}_\alpha)}$. The functions in Γ are defined in the same way. Two elements $f_b^\alpha(\bar{a})$ and $f_{b'}^{\alpha'}(\bar{a}')$ are defined to be equal iff we have $i(b) = i'(b')$ where $i : \mathfrak{C}_\alpha \to \langle\langle \bar{c}_\alpha \bar{c}_{\alpha'} \rangle\rangle_{\mathfrak{A}}$ and $i' : \mathfrak{C}_{\alpha'} \to \langle\langle \bar{c}_\alpha \bar{c}_{\alpha'} \rangle\rangle_{\mathfrak{A}}$ are the canonical inclusion maps. Since $\lambda < \kappa$ and every $\mathfrak{C}_\alpha$ has less than κ elements, it follows that each of the above statements can be expressed in $\mathrm{FO}_{\kappa\aleph_0}$. □

Corollary 2.9. *Let $F : \mathfrak{Emb}(\mathcal{C}) \to \mathfrak{Emb}(\mathcal{K})$ be $\aleph_0$-local and let Σ be the signature of $\mathcal{C}$. If κ is a cardinal such that*

$$\kappa > 2^{|\Sigma| \oplus \aleph_0} \quad \textit{and} \quad \kappa > |F(\mathfrak{C})|\,, \quad \textit{for all finitely generated } \mathfrak{C} \in \mathcal{C}\,,$$

then $\mathfrak{A} \cong_{\mathrm{FO}_{\kappa\aleph_0}} \mathfrak{B}$ implies $F(\mathfrak{A}) \cong_{\mathrm{FO}_{\kappa\aleph_0}} F(\mathfrak{B})$.

Remark. We have characterised $\aleph_0$-local functors in terms of word operations and we have shown that they preserve $\mathrm{FO}_{\infty\aleph_0}$-equivalence. These results can be generalised to κ-local functors for arbitrary cardinals κ. To do so we have to allow term algebras with operations of infinite arity less than κ. It follows that these operations preserve equivalence for the logic $\mathrm{FO}_{\infty\kappa}$ which extends $\mathrm{FO}_{\infty\aleph_0}$ by quantifiers $\exists\{ x_i \mid i < \alpha \}$ and $\forall\{ x_i \mid i < \alpha \}$ over sets of $\alpha < \kappa$ variables. We can give a back-and-forth characterisation of this logic if we replace the usual back-and-forth property by the requirement that, for every tuple $\bar{c}$ with $|\bar{c}| < \kappa$, we can find a corresponding tuple $\bar{d}$ in the other structure.

As an application of word constructions we consider varieties. With each variety $\mathcal{V}$ we can associated a so-called *replica functor* that maps a given structure to its closest approximation in $\mathcal{V}$.

Definition 2.10. Let $\Sigma \subseteq \Sigma_+$ be signatures, $P \in \Sigma_+ \setminus \Sigma$ a unary predicate, and $\mathcal{V}$ a quasivariety of Σ_+-structures.

The *replica functor* $R_{\mathcal{V}} : \mathfrak{Hom}(\Sigma) \to \mathfrak{Hom}(\mathcal{V})$ of $\mathcal{V}$ maps an arbitrary Σ-structure $\mathfrak{A}$ to the free model of the $\mathcal{V}$-presentation $\langle A; \Phi_{\mathfrak{A}} \rangle$ where

$$\Phi_{\mathfrak{A}} := \{ Pa \mid a \in A \} \cup \{ \varphi(\bar{a}) \mid \varphi \text{ atomic}, \bar{a} \subseteq A, \mathfrak{A} \vDash \varphi(\bar{a}) \}.$$

Remark. Note that replica functors differ from the functors considered so far since, in general, they do not preserve embeddings. Hence, they are functors $\mathfrak{Hom}(\Sigma) \to \mathfrak{Hom}(\Sigma_+)$, and not $\mathfrak{Emb}(\Sigma) \to \mathfrak{Emb}(\Sigma_+)$.

Lemma 2.11. *The replica functor $R_{\mathcal{V}} : \mathfrak{Hom}(\Sigma) \to \mathfrak{Hom}(\mathcal{V})$ is a functor.*

Proof. Let $h : \mathfrak{A} \to \mathfrak{B}$ be a homomorphism. By definition, the structure $R_{\mathcal{V}}(\mathfrak{A})$ is the free model of $\langle A; \Phi_{\mathfrak{A}} \rangle$. Let $\bar{a}$ be an enumeration of A and set $\bar{b} := h(\bar{a})$. Since homomorphisms preserve atomic formulae it follows that

$$\langle R_{\mathcal{V}}(\mathfrak{B}), \bar{b} \rangle \vDash \Phi_{\mathfrak{A}},$$

that is, $R_{\mathcal{V}}(\mathfrak{B})$ is a model of $\langle A; \Phi_{\mathfrak{A}} \rangle$. Since $R_{\mathcal{V}}(\mathfrak{A})$ is the free model of this presentation there exists a unique homomorphism $g : R_{\mathcal{V}}(\mathfrak{A}) \to R_{\mathcal{V}}(\mathfrak{B})$ with $g \upharpoonright A = h$. It is straightforward to check that we obtain a functor if we define $R_{\mathcal{V}}(h) := g$. □

Proposition 2.12. *Each replica functor is a word construction.*

Proof. Since the structure $R_{\mathcal{V}}(\mathfrak{A})$ is generated by the set A there exists a homomorphism $\mathfrak{T}[\Sigma_+, A] \to R_{\mathcal{V}}(\mathfrak{A})$ such that $h \upharpoonright A = \mathrm{id}_A$. We define a quantifier-free interpretation $\mathcal{I}$ such that

$$R_{\mathcal{V}} = \mathcal{I} \circ \mathcal{S} \circ \mathcal{R} \circ \mathcal{T} \circ \mathcal{R},$$

where $\mathcal{T}(\mathfrak{A}) := \mathcal{T}[\Sigma_+, \mathfrak{A}]$ and $\mathcal{S}$ is the inverse reduct that maps every sort $t \in T[\Sigma_+, S_0]$ of $\mathcal{T}[\Sigma_+, \mathfrak{A}]$ to the sort s such that $t \in T_s[\Sigma_+, S_0]$.

According to Lemma D2.4.2, we have

$$R_{\mathcal{V}}(\mathfrak{A}) \vDash \psi(\bar{a}) \quad \text{iff} \quad \mathrm{Th}(\mathcal{V}) \vDash \bigwedge \Phi_{\mathfrak{A}} \to \psi(\bar{a}),$$

for every atomic formula $\psi(\bar{x}) \in \mathrm{FO}^{<\omega}[\Sigma_+]$ and all $\bar{a} \subseteq A$.

Note that, by the interpolation theorem, we have

$$\mathrm{Th}(\mathcal{V}) \vDash \bigwedge \Phi_{\mathfrak{A}} \to \psi(\bar{a}) \quad \text{iff} \quad \mathrm{Th}(\mathcal{V}) \vDash \bigwedge \Phi_{\langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{A}}} \to \psi(\bar{a}).$$

For each atomic formula $\psi(\bar{x})$, we define

$$D_\psi := \{ \langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{A}} \mid \mathrm{Th}(\mathcal{V}) \vDash \bigwedge \Phi_{\mathfrak{A}} \to \psi(\bar{a}) \}.$$

Let $\eta_\psi(\bar{x})$ be the $\mathrm{FO}_{\infty\aleph_0}$-formula expressing that

$$\langle\!\langle \bar{x} \rangle\!\rangle_{\mathfrak{A}} \cong \mathfrak{C}, \quad \text{for some } \mathfrak{C} \in D_\psi.$$

It follows that

$$R_{\mathcal{V}}(\mathfrak{A}) \vDash \psi(\bar{a}) \quad \text{iff} \quad \langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{A}} \in D_\psi \quad \text{iff} \quad \mathfrak{A} \vDash \eta_\psi.$$

Consequently, we can define the desired interpretation

$$\mathcal{I} = \langle \alpha, (\delta_s)_{s \in S}, (\varepsilon_s)_{s \in S}, (\varphi_\xi)_{\xi \in \Sigma_+} \rangle$$

by setting

$$\begin{aligned} \alpha &:= \text{true}, \\ \delta_s(x) &:= \text{true}, \\ \varepsilon_s(x, y) &:= \text{“}x = s(\bar{a}) \text{ and } y = t(\bar{b}) \text{ and } \mathfrak{A} \vDash \eta_{s(\bar{x}) = t(\bar{y})}(\bar{a}, \bar{b})\text{”}, \\ \varphi_\xi(\bar{x}) &:= \text{“}x_i = t_i(\bar{a}_i) \text{ and } \mathfrak{A} \vDash \eta_{R\bar{t}}(\bar{a}_0, \dots, \bar{a}_{n-1})\text{”}. \end{aligned}$$

□

3. Ehrenfeucht-Mostowski models

If a functor F is $\aleph_0$-local then with every element c of $F(\mathfrak{A})$ we can associate some finitely generated substructure $\mathfrak{A}_0 \subseteq \mathfrak{A}$ such that c is contained in $F(\mathfrak{A}_0)$. We can think of the generators of $\mathfrak{A}_0$ as a code for c. In general, c can have several such codes and the connection between c and its codes is rather loose. In order to obtain a tighter relationship and a canonical way to encode elements of $F(\mathfrak{A})$, we add a function $s : A \to F(\mathfrak{A})$ assigning to every element a of $\mathfrak{A}$ some element of $F(\mathfrak{A})$ encoded by a.

Definition 3.1. Let $\mathcal{K}$ be a class of Γ-structures and Σ a signature. A functor $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ is *strongly local* if there exists a family of injective functions $s_{\mathfrak{I}} : I \to F(\mathfrak{I})$, for $\mathfrak{I} \in \mathcal{K}$, such that

- $F(\mathfrak{I})$ is generated by $\operatorname{rng} s_{\mathfrak{I}}$ and
- $F(h) \circ s_{\mathfrak{I}} = s_{\mathfrak{K}} \circ h$, for every embedding $h : \mathfrak{I} \to \mathfrak{K}$.

We call $s_{\mathfrak{I}}$ the *spine* of $F(\mathfrak{I})$.

Remark. Translated into category-theoretical terms the second of the above conditions on $s_{\mathfrak{I}}$ simply means that $(s_{\mathfrak{I}})_{\mathfrak{I}}$ is a natural transformation between the functors U and $V \circ F$, where

$$U : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Set} \quad \text{and} \quad V : \mathfrak{Emb}(\Sigma) \to \mathfrak{Set}$$

are the forgetful functors mapping a structure to its universe.

Every strongly local functor is $\aleph_0$-local. For the proof we need a technical lemma.

Lemma 3.2. *Let $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ be a strongly local functor and $h : \mathfrak{I} \to \mathfrak{K}$ an embedding in $\mathcal{K}$. Then*

$$F(h) : F(\mathfrak{I}) \cong \langle\!\langle s_{\mathfrak{K}}[\operatorname{rng} h] \rangle\!\rangle_{F(\mathfrak{K})} .$$

Proof. It is sufficient to show that $\operatorname{rng} F(h) = \langle\!\langle s_{\mathfrak{K}}[\operatorname{rng} h] \rangle\!\rangle_{F(\mathfrak{K})}$. Note that $F(h) \circ s_{\mathfrak{I}} = s_{\mathfrak{K}} \circ h$ implies

$$F(h)[\operatorname{rng} s_{\mathfrak{I}}] = s_{\mathfrak{K}}[\operatorname{rng} h] .$$

Therefore, $\langle\!\langle \operatorname{rng} s_{\mathfrak{I}} \rangle\!\rangle_{F(\mathfrak{I})} = F(I)$ implies

$$\begin{aligned}\operatorname{rng} F(h) &= F(h)\big[\langle\!\langle \operatorname{rng} s_{\mathfrak{I}} \rangle\!\rangle_{F(\mathfrak{I})}\big] \\ &= \big\langle\!\big\langle F(h)[\operatorname{rng} s_{\mathfrak{I}}] \big\rangle\!\big\rangle_{F(\mathfrak{K})} = \langle\!\langle s_{\mathfrak{K}}[\operatorname{rng} h] \rangle\!\rangle_{F(\mathfrak{K})} . \qquad \square\end{aligned}$$

Proposition 3.3. *Let $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ be a strongly local functor where $\mathcal{K}$ is $\aleph_0$-hereditary. Then F is $\aleph_0$-local.*

Proof. Fix $\mathfrak{I} \in \mathcal{K}$ and suppose that $X \subseteq F(\mathfrak{I})$ is finite. Then there is a finite subset $Z \subseteq \operatorname{rng} s_{\mathfrak{I}}$ such that $X \subseteq \langle\!\langle Z \rangle\!\rangle_{F(\mathfrak{I})}$. Set

$$\mathfrak{I}_0 := \langle\!\langle s_{\mathfrak{I}}^{-1}[Z] \rangle\!\rangle_{\mathfrak{I}} .$$

Note that $\mathfrak{I}_0 \in \mathcal{K}$ since $\mathcal{K}$ is $\aleph_0$-hereditary. By Lemma 3.2, it follows that

$$X \subseteq \langle\!\langle Z \rangle\!\rangle_{F(\mathfrak{I})} = \langle\!\langle \operatorname{rng} s_{\mathfrak{I}_0} \rangle\!\rangle_{F(\mathfrak{I})} \cong F(\mathfrak{I}_0) . \qquad \square$$

By Corollary 2.9 it follows that strongly local functors preserve $\mathrm{FO}_{\kappa\aleph_0}$-equivalence.

Corollary 3.4. *Let $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ be a strongly local functor where $\mathcal{K}$ is an $\aleph_0$-hereditary class of Γ-structures. For every cardinal $\kappa \geq 2^{|\Gamma| \oplus \aleph_0}$ and all $\mathfrak{I}, \mathfrak{K} \in \mathcal{K}$,*

$$\mathfrak{I} \equiv_{\mathrm{FO}_{\kappa\aleph_0}} \mathfrak{K} \quad \text{implies} \quad F(\mathfrak{I}) \equiv_{\mathrm{FO}_{\kappa\aleph_0}} F(\mathfrak{K}) .$$

Strongly local functors also preserve QF-equivalence.

Lemma 3.5. *Suppose that $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ is a strongly local functor where the class $\mathcal{K}$ is $\aleph_0$-hereditary.*

Let $\mathfrak{I}, \mathfrak{K} \in \mathcal{K}$ be structures and $\bar{a} \subseteq I$ and $\bar{b} \subseteq K$ finite tuples. Then

$$\langle \mathfrak{I}, \bar{a} \rangle \equiv_0 \langle \mathfrak{K}, \bar{b} \rangle \quad \text{implies} \quad \big\langle F(\mathfrak{I}), s_{\mathfrak{I}}(\bar{a}) \big\rangle \equiv_0 \big\langle F(\mathfrak{K}), s_{\mathfrak{K}}(\bar{b}) \big\rangle .$$

Proof. Set $\mathfrak{L} := \langle\!\langle \bar{a} \rangle\!\rangle_{\mathfrak{I}}$ and let $s_{\mathfrak{L}}$ be the spine of $\mathfrak{L}$. Since $\mathcal{K}$ is $\aleph_0$-hereditary we have $\mathfrak{L} \in \mathcal{K}$. Since $\langle \mathfrak{I}, \bar{a} \rangle \equiv_0 \langle \mathfrak{K}, \bar{b} \rangle$, there are embeddings $f : \mathfrak{L} \to \mathfrak{I}$ and $g : \mathfrak{L} \to \mathfrak{K}$ with $f(\bar{a}) = \bar{a}$ and $g(\bar{a}) = \bar{b}$. Note that

$$(F(f) \circ s_{\mathfrak{L}})(\bar{a}) = (s_{\mathfrak{I}} \circ f)(\bar{a}) = s_{\mathfrak{I}}(\bar{a}) ,$$

and $\quad (F(g) \circ s_{\mathfrak{L}})(\bar{a}) = (s_{\mathfrak{K}} \circ g)(\bar{a}) = s_{\mathfrak{K}}(\bar{b})\,.$

Since embeddings preserve every quantifier-free formula φ, it follows that

$$\begin{aligned} F(\mathfrak{J}) \vDash \varphi(s_{\mathfrak{J}}(\bar{a})) \quad &\text{iff} \quad F(\mathfrak{L}) \vDash \varphi(s_{\mathfrak{L}}(\bar{a})) \\ &\text{iff} \quad F(\mathfrak{K}) \vDash \varphi(s_{\mathfrak{K}}(\bar{b}))\,. \end{aligned}$$

□

Corollary 3.6. *Let $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ be a strongly local functor where the class $\mathcal{K}$ is $\aleph_0$-hereditary. For every $\mathfrak{J} \in \mathcal{K}$, the spine $s_{\mathfrak{J}}$ of $F(\mathfrak{J})$ is a* QF*-indiscernible system over $\mathfrak{J}$.*

Next we study the first-order theory of structures in the range of a strongly local functor.

Proposition 3.7. *Let $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ be a strongly local functor and $\mathfrak{U} \in \mathcal{K}$ an $\aleph_0$-universal structure. If $\mathrm{Th}(F(\mathfrak{U}))$ is a Skolem theory then $\mathrm{Th}(F)$ is complete. In particular,*

$$F(\mathfrak{J}) \equiv F(\mathfrak{K})\,, \quad \text{for all } \mathfrak{J}, \mathfrak{K} \in \mathcal{K}\,.$$

Furthermore, each spine $s_{\mathfrak{J}}$ is an indiscernible system over $\mathfrak{J}$.

Proof. A Skolem theory is $\forall$-axiomatisable and admits quantifier elimination. Let $\Phi \subseteq \forall$ be an axiom system for $\mathrm{Th}(F(\mathfrak{U}))$. By Lemma 1.10. we have $\Phi \subseteq \mathrm{Th}(F)$. Hence,

$$\mathrm{Th}(F(\mathfrak{U})) = \Phi^{\vDash} \subseteq \mathrm{Th}(F) \subseteq \mathrm{Th}(F(\mathfrak{U}))$$

implies that $F(\mathfrak{J}) \equiv F(\mathfrak{U})$, for all $\mathfrak{J} \in \mathcal{K}$.

To show that every spine $s_{\mathfrak{J}}$ is indiscernible, fix $\mathfrak{J} \in \mathcal{K}$ and let $\bar{c}, \bar{d} \subseteq I$ be tuples with $\mathrm{atp}(\bar{c}) = \mathrm{atp}(\bar{d})$. For every formula $\varphi(\bar{x})$, there exists a quantifier-free formula $\psi(\bar{x})$ with $F(\mathfrak{J}) \vDash \varphi \leftrightarrow \psi$. By Lemma 3.5, it follows that

$$\begin{aligned} F(\mathfrak{J}) \vDash \varphi(s_{\mathfrak{J}}[\bar{c}]) \quad &\text{iff} \quad F(\mathfrak{J}) \vDash \psi(s_{\mathfrak{J}}[\bar{c}]) \\ &\text{iff} \quad F(\mathfrak{J}) \vDash \psi(s_{\mathfrak{J}}[\bar{d}]) \\ &\text{iff} \quad F(\mathfrak{J}) \vDash \varphi(s_{\mathfrak{J}}[\bar{d}])\,. \end{aligned}$$

□

Existence and uniqueness of strongly local functors is proved in the following proposition.

Proposition 3.8. *Let* $\mathfrak{A}$ *be a* Σ*-structure,* $\mathfrak{U}$ *a* Γ*-structure, and set*

$$\mathcal{K} := \{ \mathfrak{I} \mid \mathrm{Sub}_{\aleph_0}(\mathfrak{I}) \subseteq \mathrm{Sub}_{\aleph_0}(\mathfrak{U}) \} .$$

Suppose that $\mathfrak{A}$ *is generated by a* QF*-indiscernible system* $a : U \to A$ *over* $\mathfrak{U}$*. Up to natural isomorphism there exists a unique strongly local functor* $F : \mathfrak{Emb}(\mathcal{K}) \to \mathfrak{Emb}(\Sigma)$ *such that*

$$F(\mathfrak{U}) \cong \mathfrak{A} \quad \textit{and} \quad \mathrm{Av}_{\mathrm{QF}}(s_{\mathfrak{U}}) = \mathrm{Av}_{\mathrm{QF}}(a) .$$

Proof. For each $\mathfrak{I} \in \mathcal{K}$, we define a set $\Phi(\mathfrak{I}) \subseteq \mathrm{QF}^{\circ}[\Sigma_I]$ by

$$\Phi(\mathfrak{I}) := \big\{ \varphi(\bar{c}) \bigm| \bar{c} \subseteq I \text{ and } \varphi(\bar{x}) \in \mathrm{Av}_{\mathrm{QF}}(a)(\mathrm{atp}(\bar{c}/\mathfrak{I})) \big\} .$$

We claim that $\Phi(\mathfrak{I})$ is =-closed. Since every type $\mathfrak{q}$ contains the equation $t(\bar{x}) = t(\bar{x})$, we have

$$t(\bar{c}) = t(\bar{c}) \in \Phi(\mathfrak{I}) , \quad \text{for every term } t(\bar{c}) \in T[\Sigma_I, \varnothing] .$$

Furthermore, if $\Phi(\mathfrak{I})$ contains the formulae $\varphi(t(\bar{c}), \bar{c})$ and $t(\bar{c}) = t'(\bar{c})$ then

$$\varphi(t(\bar{x}), \bar{x}),\ t(\bar{x}) = t'(\bar{x}) \in \mathrm{Av}_{\mathrm{QF}}(a)(\mathrm{atp}(\bar{c}/\mathfrak{I}))$$

implies

$$\varphi(t'(\bar{x}), \bar{x}) \in \mathrm{Av}_{\mathrm{QF}}(a)(\mathrm{atp}(\bar{c}/\mathfrak{I})) .$$

Consequently, $\varphi(t'(\bar{c}), \bar{c}) \in \Phi(\mathfrak{I})$. Hence, we can use Lemma C2.4.4 to construct a Herbrand model $\mathfrak{H}(\mathfrak{I})$ of $\Phi(\mathfrak{I})$ such that

$$\Phi(\mathfrak{I}) = \big\{ \varphi \in \mathrm{QF}^{\circ}[\Sigma_I] \bigm| \mathfrak{H}(\mathfrak{I}) \vDash \varphi \big\} .$$

We define the desired strongly local functor by setting

$$F(\mathfrak{I}) := \mathfrak{H}(\mathfrak{I})|_{\Sigma} \quad \text{and} \quad s_{\mathfrak{I}}(c) := c^{\mathfrak{H}(\mathfrak{I})}, \quad \text{for } c \in I .$$

First, note that the mapping $s_{\mathfrak{J}}$ is injective since we have $x_0 \neq x_1 \in \operatorname{tp}(a[vv'])$, for all elements $v \neq v'$ of $\mathfrak{U}$. Furthermore, if $h : \mathfrak{J} \to \mathfrak{K}$ is an embedding, $\bar{c} \subseteq I$, and $\varphi(\bar{x})$ quantifier-free, then

$$\begin{aligned} F(\mathfrak{J}) \vDash \varphi(s_{\mathfrak{J}}(\bar{c})) \quad &\text{iff} \quad \varphi(\bar{x}) \in \operatorname{Av}_{\mathrm{QF}}(a)(\operatorname{atp}(\bar{c}/\mathfrak{J})) \\ &\text{iff} \quad \varphi(\bar{x}) \in \operatorname{Av}_{\mathrm{QF}}(a)(\operatorname{atp}(h(\bar{c})/\mathfrak{K})) \\ &\text{iff} \quad F(\mathfrak{K}) \vDash \varphi(s_{\mathfrak{K}}(h(\bar{c}))) . \end{aligned}$$

By the Diagram Lemma it follows that the function

$$F(h) : t^{F(\mathfrak{J})}(s_{\mathfrak{J}}(\bar{c})) \mapsto t^{F(\mathfrak{K})}(s_{\mathfrak{K}}(h(\bar{c})))$$

is an embedding $F(h) : F(\mathfrak{J}) \to F(\mathfrak{K})$. Consequently, F is a functor. By construction, it further follows that it is strongly local, that $F(\mathfrak{U}) \cong \mathfrak{A}$, and that $\operatorname{Av}_{\mathrm{QF}}(s_{\mathfrak{U}}) = \operatorname{Av}_{\mathrm{QF}}(a)$. Hence, it remains to check uniqueness.

Suppose that G is another strongly local functor such that $G(\mathfrak{U}) \cong \mathfrak{A}$ and $\operatorname{Av}_{\mathrm{QF}}(s'_{\mathfrak{U}}) = \operatorname{Av}_{\mathrm{QF}}(a)$, where $s'_{\mathfrak{U}}$ is the spine of $G(\mathfrak{U})$. For every $\mathfrak{J} \in \mathcal{K}$, each finite tuple $\bar{c} \subseteq I$, and all quantifier-free formulae $\varphi(\bar{x})$, it follows that

$$\begin{aligned} G(\mathfrak{J}) \vDash \varphi(s'_{\mathfrak{J}}(\bar{c})) \quad &\text{iff} \quad G(\mathfrak{U}) \vDash \varphi((s'_{\mathfrak{U}} \circ g)(\bar{c})) \\ &\text{iff} \quad \varphi(\bar{x}) \in \operatorname{Av}_{\mathrm{QF}}(s'_{\mathfrak{U}})(\operatorname{atp}(g(\bar{c})/\mathfrak{U})) \\ &\text{iff} \quad \varphi(\bar{x}) \in \operatorname{Av}_{\mathrm{QF}}(a)(\operatorname{atp}(g(\bar{c})/\mathfrak{U})) \\ &\text{iff} \quad \varphi(\bar{x}) \in \operatorname{Av}_{\mathrm{QF}}(a)(\operatorname{atp}(\bar{c}/\mathfrak{J})) \\ &\text{iff} \quad F(\mathfrak{J}) \vDash \varphi(s_{\mathfrak{J}}(\bar{c})) , \end{aligned}$$

where $g : \langle\!\langle \bar{c} \rangle\!\rangle_{\mathfrak{J}} \to \mathfrak{U}$ is an arbitrary embedding and $s'_{\mathfrak{J}}$ and $s'_{\mathfrak{U}}$ are the spines of $G(\mathfrak{J})$ and $G(\mathfrak{U})$, respectively. Since $F(\mathfrak{J})$ and $G(\mathfrak{J})$ are generated by, respectively, $\operatorname{rng} s_{\mathfrak{J}}$ and $\operatorname{rng} s'_{\mathfrak{J}}$ it follows that we obtain an isomorphism $\pi : F(\mathfrak{J}) \to G(\mathfrak{J})$ by setting

$$\pi\big(t^{F(\mathfrak{J})}(s_{\mathfrak{J}}(\bar{c}))\big) := t^{G(\mathfrak{J})}(s'_{\mathfrak{J}}(\bar{c})) ,$$

for all terms $t(\bar{x})$ and all $\bar{c} \subseteq I$. $\square$

Of particular importance are strongly local functors $F : \mathfrak{Emb}(\mathcal{L}) \to \mathfrak{Emb}(\Sigma)$ where $\mathcal{L}$ is the class of all linear orders. This is mainly due to the fact that we always can find enough indiscernible sequences, whereas arbitrary indiscernible systems do not need to exist. Note that $\mathcal{L}$ is hereditary and every infinite linear order is $\aleph_0$-universal.

Definition 3.9. Let $\mathfrak{Lin} := \mathfrak{Emb}(\mathcal{L})$ where $\mathcal{L}$ is the class of all linear orders.

(a) A strongly local functor $F : \mathfrak{Lin} \to \mathfrak{Emb}(\Sigma)$ is called an *Ehrenfeucht-Mostowski* functor. We say that F is an Ehrenfeucht-Mostowski functor *for a theory T* if F is an Ehrenfeucht-Mostowski functor such that $F(I) \vDash T$, for every linear order I.

(b) Let T be a first-order theory. An *Ehrenfeucht-Mostowski model* of T is a model of the form $F(I)$ where F is some Ehrenfeucht-Mostowski functor for T and I is a linear order.

(c) Let $F : \mathfrak{Lin} \to \mathfrak{Emb}(\Sigma)$ be an Ehrenfeucht-Mostowski functor. The *average type* of F is the set

$$\operatorname{Av}(F) := \big\{ \varphi(\bar{x}) \in \mathrm{FO}^{<\omega}[\Sigma] \bigm| F(\mathfrak{I}) \vDash \varphi(s_{\mathfrak{I}}(\bar{c})) \text{ for all } \mathfrak{I} \in \mathcal{K} \text{ and } \bar{c} \in [I]^{<\omega} \big\} .$$

Note that, by Proposition 3.7 and Lemma 3.5, the average type of an Ehrenfeucht-Mostowski function is complete.

Lemma 3.10. *If $F : \mathfrak{Lin} \to \mathfrak{Emb}(\Sigma)$ is an Ehrenfeucht-Mostowski functor, then* $\operatorname{Av}(F)$ *is a complete type.*

Theorem 3.11 (Ehrenfeucht-Mostowski). *Let $\mathfrak{M}$ be a model of a Skolem theory T. For every sequence $(a^i)_{i \in I}$ of distinct elements in $\mathfrak{M}$ there exists an Ehrenfeucht-Mostowski functor F for T such that*

$$\operatorname{Av}((a^i)_i/\varnothing) \subseteq \operatorname{Av}(F) .$$

Proof. By Proposition E5.3.6, there exists an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ containing an indiscernible sequence $(c_n)_{n<\omega}$ with

$$\operatorname{Av}((a^i)_{i \in I}/\varnothing) \subseteq \operatorname{Av}((c_n)_{n<\omega}/\varnothing) .$$

Let $s : \omega \to N$ be the function mapping $n < \omega$ to c_n and set

$$\mathfrak{U} := \langle\!\langle \operatorname{rng} s \rangle\!\rangle_{\mathfrak{N}} .$$

Note that the function s is injective, since $x_0 \neq x_1 \in \operatorname{Av}((a^i)_i/\varnothing)$. Furthermore, we have $\mathfrak{U} \preceq \mathfrak{N}$ since T is a Skolem theory. Hence, we can use Proposition 3.8 to find an Ehrenfeucht-Mostowski functor F with $F(\omega) = \mathfrak{U}$ and $s_\omega = s$. It follows that $\operatorname{Av}((a^i)_i/\varnothing) \subseteq \operatorname{Av}((c_n)_n/\varnothing) = \operatorname{Av}(F)$. □

Corollary 3.12. *If a first-order theory T has infinite models then there exists an Ehrenfeucht-Mostowski functor for T.*

Proof. Let T^+ be a Skolemisation of T. It is sufficient to find an Ehrenfeucht-Mostowski functor F for T^+ since we can obtain the desired Ehrenfeucht-Mostowski functor for T by composing F with a suitable reduct functor.

Let $\mathfrak{M}^+$ be an infinite model of T^+ that contains an indiscernible sequence $(a^n)_{n<\omega}$ of distinct elements. By Theorem 3.11, there exists an Ehrenfeucht-Mostowski functor F with $\operatorname{Av}((a^n)_n) \subseteq \operatorname{Av}(F)$. We claim that F is the desired Ehrenfeucht-Mostowski functor for T^+. As $(a^n)_n$ is indiscernible, its average type $\operatorname{Av}((a^n)_n)$ is complete and, therefore, equal to $\operatorname{Av}(F)$. Consequently, $F(\omega) \vDash T^+$. Since T^+ is a Skolem theory, it follows by Lemma 3.7 that $F(I) \vDash T^+$, for every I. □

We use Ehrenfeucht-Mostowski functors to construct models of a theory with certain properties. As a first simple application, we build models with many automorphisms.

Lemma 3.13. *Let T be a complete first-order theory with infinite models. For every cardinal $\kappa \geq |T|$, there exists a model $\mathfrak{M}$ of T of size $|M| = \kappa$ with 2^κ automorphisms.*

Proof. According to Corollary 3.12, there is an Ehrenfeucht-Mostowski functor $F : \mathfrak{Lin} \to \operatorname{Mod}(T)$ for T. We will construct a linear order I of size $|I| = \kappa$ with 2^κ automorphisms. It follows that $F(I)$ is the desired model of T.

Let $I := \mathbb{Z} \cdot \kappa$ be the product of the order $\mathbb{Z}$ of the integers and the well-order κ. For every set $X \subseteq \kappa$, we can define an automorphism $\pi_X : I \to I$ by

$$\pi_X\langle k, \alpha\rangle := \begin{cases} \langle k+1, \alpha\rangle & \text{if } \alpha \in X\,, \\ \langle k, \alpha\rangle & \text{if } \alpha \notin X\,. \end{cases}$$

Since $\pi_X \neq \pi_Y$, for $X \neq Y$, it follows that I has at least 2^κ automorphisms. □

One important application of Ehrenfeucht-Mostowski models rests on the fact that such models realise few types.

Theorem 3.14. *Let T be a Skolem theory over the signature Σ and let $\mathfrak{M}$ be an Ehrenfeucht-Mostowski model of T.*

(a) *For every finite sequence of sorts $\bar{s}$, $\mathfrak{M}$ realises at most $|\Sigma| \oplus \aleph_0$ types in $S^{\bar{s}}(T)$.*

(b) *Let s be a sort and $U \subseteq M$. If the spine of $\mathfrak{M}$ is well-ordered then $\mathfrak{M}$ realises at most $|\Sigma| \oplus |U| \oplus \aleph_0$ types in $S^s(U)$.*

Proof. (a) Suppose that $\mathfrak{M} = F(I)$ for some Ehrenfeucht-Mostowski functor F. Fix a finite tuple $\bar{s}$ of sorts and let $\bar{a} \in M^{\bar{s}}$ be a tuple of elements of the corresponding sorts. For each index l there exists a term $t_l(\bar{x})$ and an increasing tuple $\bar{\imath}^l \subseteq I$ such that $a_l = t_l^{\mathfrak{M}}(s_I[\bar{\imath}^l])$. By adding redundant variables we may assume that all the tuples $\bar{\imath}^l$ are equal. We denote this tuple by $\bar{\imath}$. If $\bar{k} \subseteq I$ is another increasing tuple of the same length then it follows from indiscernibility of the spine s_I that

$$\begin{aligned} & \mathfrak{M} \vDash \varphi(t_0(s_I[\bar{\imath}]), \dots, t_{n-1}(s_I[\bar{\imath}])) \\ \text{iff}\quad & \mathfrak{M} \vDash \varphi(t_0(s_I[\bar{k}]), \dots, t_{n-1}(s_I[\bar{k}]))\,, \end{aligned}$$

for every formula φ. Setting $b_l := t_l(s_I[\bar{k}])$ we obtain $\operatorname{tp}(\bar{a}) = \operatorname{tp}(\bar{b})$. Hence, the type of $\bar{a}$ is uniquely determined by the terms t_l. Since

$$|T^{<\omega}_{s_l}[\Sigma]| = |\Sigma| \oplus \aleph_0$$

it follows that $\mathfrak{M}$ realises at most $|\Sigma| \oplus \aleph_0$ types from $S^{\bar{s}}(T)$.

(b) Suppose that $\mathfrak{M} = F(\alpha)$, for some ordinal α, and let $U \subseteq M$. Each element $c \in U$ can be written as $c = t_c^{\mathfrak{M}}(s_\alpha(\bar{\imath}_c))$, for some term t_c and indices $\bar{\imath}_c \subseteq \alpha$. The set $W := \bigcup_{c \in U} \bar{\imath}_c$ has size $|W| \leq |U| \oplus \aleph_0$. Let $u(\bar{x}) \in T^{<\omega}[\Sigma]$ be a term and $\bar{k} \subseteq \alpha$. By indiscernibility of s_α the type of $u^{\mathfrak{M}}(\bar{k})$ is determined by the relative position of $\bar{k}$ with respect to the elements of W. Since α is well-ordered, there are at most $|W| \oplus \aleph_0$ ways in which $\bar{k}$ can lie relative to W. Consequently, the elements $u^{\mathfrak{M}}(\bar{k})$ with $\bar{k} \subseteq \alpha$ realise at most $|W| \oplus \aleph_0$ complete types over U. Therefore, at most

$$|T^{<\omega}_s[\Sigma]| \oplus |W| \oplus \aleph_0 \leq |\Sigma| \oplus |U| \oplus \aleph_0$$

complete s-types over U are realised in $\mathfrak{M}$. □

Corollary 3.15. *Let T be a complete first-order theory with infinite models. For every cardinal $\kappa \geq |T|$, T has an Ehrenfeucht-Mostowski model $\mathfrak{M}$ of size $|M| = \kappa$ such that, for every set $U \subseteq M$ and every finite tuple $\bar{s}$ of sorts, $\mathfrak{M}$ realises at most $|U| \oplus |T|$ types from $S^{\bar{s}}(U)$.*

Proof. According to Corollary 3.12, there is an Ehrenfeucht-Mostowski functor $F : \mathfrak{Lin} \to \operatorname{Mod}(T)$ for T. Let $\mathfrak{M} := F(\kappa)$. Then $|M| = \kappa$ and, by Theorem 3.14 (b), $\mathfrak{M}$ realises at most $|U| \oplus |T|$ types in $S^s(U)$, for every $U \subseteq M$. For a finite tuple $\bar{s} = s_0 \dots s_{n-1}$ it follows by induction that $\mathfrak{M}$ realises at most $(|U| \oplus |T|)^n = |U| \oplus |T|$ types in $S^{\bar{s}}(U)$. □

Theorem 3.16. *Let Σ be a signature. If a theory T over Σ is κ-categorical for some $\kappa \geq |\Sigma| \oplus \aleph_0$, then T is λ-stable, for every cardinal $|\Sigma| \oplus \aleph_0 \leq \lambda < \kappa$.*

Proof. Let $\mathfrak{M}$ be the Ehrenfeucht-Mostowski model from Corollary 3.15. For a contradiction, suppose that there is some set U of size $|U| = \lambda$ with $|S^s(U)| > \lambda$. Let $\mathfrak{N}$ be a model of T containing U that realises λ^+ of these types. By the Theorem of Löwenheim and Skolem we can choose $\mathfrak{N}$ to be of size $|N| = \lambda^+ \leq \kappa$. Hence, $\mathfrak{N}$ has an elementary extension $\mathfrak{N}_+$ of size $|N_+| = \kappa$. As T is κ-categorical this implies $\mathfrak{N}_+ \cong \mathfrak{M}$ and there exists an elementary embedding $h : \mathfrak{N} \to \mathfrak{M}$. Hence, $\mathfrak{M}$ contains a subset $h[U]$ of size λ such that more than λ types over $h[U]$ are realised in $\mathfrak{M}$. This contradicts our choice of $\mathfrak{M}$. □

Corollary 3.17. *Let T be a theory over a countable signature. If T is κ-categorical for some uncountable cardinal κ then T is $\aleph_0$-stable.*

The next proposition generalises Lemma E4.1.6.

Proposition 3.18. *Let T be a countable, complete theory. If there is some finite sequence $\bar{s}$ of sorts such that $S^{\bar{s}}(T)$ is uncountable then, for each infinite cardinal κ, T has at least $2^{\aleph_0}$ pairwise non-isomorphic models of cardinality κ.*

Proof. Let κ be an infinite cardinal and fix $\bar{s}$ such that $S^{\bar{s}}(T)$ is uncountable. By Corollary B5.7.5, it follows that $|S^{\bar{s}}(T)| = 2^{\aleph_0}$. Note that this also implies that T has infinite models. Let $\bar{c}$ be a tuple of new constant symbols of sorts $\bar{s}$. For each $\mathfrak{p}(\bar{x}) \in S^{\bar{s}}(T)$ we form the theory $T_\mathfrak{p} := T \cup \mathfrak{p}(\bar{c})$. Let $T_\mathfrak{p}^+$ be a Skolemisation of $T_\mathfrak{p}$. We can use Theorem 3.11 to find an Ehrenfeucht-Mostowski model $\mathfrak{A}_\mathfrak{p}$ of $T_\mathfrak{p}^+$ with a spine $s_\mathfrak{p} : \kappa \to A_\mathfrak{p}$. It follows that

$$\kappa \leq |A_\mathfrak{p}| \leq \kappa \oplus |T_\mathfrak{p}^+| = \kappa \oplus \aleph_0 = \kappa\,.$$

By Theorem 3.14 $\mathfrak{A}_\mathfrak{p}$ realises only countably many $\bar{s}$-types. Therefore, so does $\mathfrak{B}_\mathfrak{p} := \mathfrak{A}_\mathfrak{p}|_\Sigma$. Furthermore, the tuple $\bar{c}^{\mathfrak{A}_\mathfrak{p}}$ realises the type $\mathfrak{p}$ in $\mathfrak{B}_\mathfrak{p}$.

We claim that there are $2^{\aleph_0}$ pairwise non-isomorphic models among the $\mathfrak{B}_\mathfrak{p}$. Suppose otherwise. Then there exists a set $I \subseteq S^{\bar{s}}(T)$ of size $|I| < 2^{\aleph_0}$ such that every $\mathfrak{B}_\mathfrak{p}$ is isomorphic to some $\mathfrak{B}_\mathfrak{q}$ with $\mathfrak{q} \in I$. Since every type in $S^{\bar{s}}(T)$ is realised in some $\mathfrak{B}_\mathfrak{p}$, but each $\mathfrak{B}_\mathfrak{p}$ realises only countably many types, it follows that

$$|S^{\bar{s}}(T)| \leq |I| \otimes \aleph_0 < 2^{\aleph_0}.$$

Contradiction. □

Definable linear orders in an Ehrenfeucht-Mostowski model $F(I)$ are closely related to the order induced by I. We start with a technical lemma.

Lemma 3.19. *Let $\langle A, < \rangle$ be an infinite dense linear order and suppose that $\sqsubset$ is a linear order on $[A]^n$ with the following property. For all tuples $\bar{a}, \bar{a}', \bar{b}, \bar{b}' \in [A]^n$ such that $\bar{a}\bar{b}$ and $\bar{a}'\bar{b}'$ have the same order type with respect to $<$, we have*

$$\bar{a} \sqsubset \bar{b} \quad \text{iff} \quad \bar{a}' \sqsubset \bar{b}' .$$

Then there exist a linear order $\triangleleft$ on $[n]$ and a map $\sigma : [n] \to \{-1, 1\}$ such that,

$$\bar{a} \sqsubset \bar{b}$$

iff there is some $l \in [n]$ with $a_l <^{\sigma(l)} b_l$ and $a_i = b_i$, for $i \triangleleft l$,

where $<^1 := <$ and $<^{-1} := >$.

Proof. We start by defining linear orders $\prec_i$ on A, for $i < n$, by

$$a \prec_i b \quad :\text{iff} \quad \bar{c}[i/a] \sqsubset \bar{c}[i/b], \quad \text{for some } \bar{c} \in [A]^n \text{ with } c_{i-1} < a < c_{i+1} \text{ and } c_{i-1} < b < c_{i+1} .$$

(Recall that, according to Definition B3.1.12, $\bar{c}[i/a]$ denotes the tuple obtained from $\bar{c}$ by replacing c_i by a.) Note that, by our assumption on $\sqsubset$, if $a \prec_i b$ holds then we have $\bar{c}[i/a] \sqsubset \bar{c}[i/b]$ for *all* tuples $\bar{c}$ satisfying the above conditions. Furthermore, since we can always find such a tuple and $\sqsubset$ is linear it follows that $a \prec_i b$ or $b \prec_i a$. Finally, if $a \prec_i b$ holds for some $a < b$ then it holds for all $a < b$. Therefore, we have $\prec_i = <$ or $\prec_i = <^{-1}$. Let $\sigma : [n] \to \{1, -1\}$ be the function with $\prec_i = <^{\sigma(i)}$.

We define the ordering $\triangleleft$ on $[n]$ by

$$i \triangleleft j \quad \text{iff} \quad i \neq j \text{ and there are } a \prec_i a',\ b \prec_j b',\ \text{and } \bar{c} \text{ such that } \bar{c}[i/a, j/b'] \sqsubset \bar{c}[i/a', j/b] \text{ and these tuples are increasing.}$$

By assumption on $\sqsubset$ it follows that the definition of $i \lhd j$ does not depend on the choice of a, a', b, b' and $\bar{c}$. If there are some elements satisfying the definition above then we have $\bar{c}[i/a, j/b'] \sqsubset \bar{c}[i/a', j/b]$ for *all* elements as above. Consequently, $i \lhd j$ implies $j \ntriangleleft i$. Furthermore, since $\sqsubset$ is linear we have $i \lhd j$ or $j \lhd i$, for all i, j. In order to show that $\lhd$ is a linear order it therefore remains to prove that it is transitive.

Suppose that $i \lhd j \lhd k$. We have to show that $i \lhd k$. If $i = k$ we would have $i \lhd j$ and $j \lhd i$, which is impossible. Hence, $i \neq k$. Choose elements $a <_i a'$, $b <_k b'$, and $\bar{c}$ such that the tuples $\bar{c}[i/a, k/b']$ and $\bar{c}[i/a', k/b]$ are increasing. We claim that $\bar{c}[i/a, k/b'] \sqsubset \bar{c}[i/a', k/b]$. Since A is dense we can find some element $d <_j c_j$ such that $\bar{c}[i/a', j/d, k/b]$ is increasing. Then $i \lhd j$ implies that

$$\bar{c}[i/a, k/b'] = \bar{c}[i/a, j/c_j, k/b'] \sqsubset \bar{c}[i/a', j/d, k/b'] .$$

Similarly, $j \lhd k$ implies

$$\bar{c}[i/a', j/d, k/b'] \sqsubset \bar{c}[i/a', j/c_j, k/b] = \bar{c}[i/a', k/b] .$$

Therefore, we have

$$\bar{c}[i/a, k/b'] \sqsubset \bar{c}[i/a', k/b] ,$$

as desired.

It remains to prove that the ordering $\sqsubset$ coincides with the ordering $\sqsubset^{\sigma}_{\lhd}$ induced by $\lhd$ and σ as in the claim above. Since both relations are linear orders it is sufficient to prove that $\bar{a} \sqsubset^{\sigma}_{\lhd} \bar{b}$ implies $\bar{a} \sqsubset \bar{b}$.

For $\bar{a}, \bar{b} \in [A]^n$, let $d(\bar{a}, \bar{b})$ be the number of indices i with $a_i \neq b_i$. We prove the claim by induction on $d := d(\bar{a}, \bar{b})$. If $d = 0$ then $\bar{a} \not\sqsubset^{\sigma}_{\lhd} \bar{b}$ and there is nothing to prove.

Suppose that $d = 1$ and let l be the unique index with $a_l \neq b_l$. Then we have

$$\bar{a} \sqsubset \bar{b} \quad \text{iff} \quad a_l <_l b_l \quad \text{iff} \quad a_l <^{\sigma(l)} b_l \quad \text{iff} \quad \bar{a} \sqsubset^{\sigma}_{\lhd} \bar{b} .$$

Suppose that $d = 2$. Let l and j be the indices where $\bar{a}$ and $\bar{b}$ differ and suppose that $l \lhd j$. By definition of $\sqsubseteq^{\sigma}_{\lhd}$ we have $a_l \prec_l b_l$. Hence, if $b_j \prec_j a_j$ then $l \lhd j$ implies that

$$\bar{a} = \bar{a}[l/a_l, j/a_j] \sqsubset \bar{a}[l/b_l, j/b_j] = \bar{b}\,,$$

and we are done. Suppose therefore that $a_j \prec_j b_j$. Let $k_0 := \min\{l, j\}$ and $k_1 := \max\{l, j\}$ (with respect to the natural ordering on $[n]$). If $a_{k_0} < b_{k_0}$ then $\bar{a}[k_1/b_{k_1}] \in [A]^n$ and, by inductive hypothesis, we have

$$\bar{a} \sqsubset \bar{a}[k_1/b_{k_1}] = \bar{b}[k_0/a_{k_0}] \sqsubset \bar{b}\,.$$

Similarly, $b_{k_0} < a_{k_0}$ implies that

$$\bar{a} \sqsubset \bar{a}[k_0/b_{k_0}] = \bar{b}[k_1/a_{k_1}] \sqsubset \bar{b}\,.$$

Finally, suppose that $d > 2$. Let l be the $\lhd$-minimal index with $a_l \neq b_l$ and let k be the $<$-maximal one. First, consider the case that $k \neq l$. If $a_k \prec_k b_k$ then we have

$$\bar{a} \sqsubset^{\sigma}_{\lhd} \bar{a}[k/b_k] \sqsubset^{\sigma}_{\lhd} \bar{b}\,,$$

and the claim follows by inductive hypothesis. Therefore, suppose that $b_k \prec_k a_k$. Since A is dense we can find some element c with $a_l \prec_l c \prec_l b_l$ and $a_{l-1}, b_{l-1} < c < a_{l+1}, b_{l+1}$. Then

$$\bar{a} \sqsubset^{\sigma}_{\lhd} \bar{a}[l/c, k/b_k] \sqsubset^{\sigma}_{\lhd} \bar{b}\,,$$

and the claim follows by inductive hypothesis.

It remains to consider the case that $k = l$. Let k' be the $<$-minimal index with $a_{k'} \neq b_{k'}$. Then $k' \neq l$ and we can use a dual argument to show that $\bar{a} \sqsubset \bar{b}$. □

Theorem 3.20. *Let $F : \mathfrak{Lin} \to \mathfrak{Emb}(\Sigma)$ be an Ehrenfeucht-Mostowski functor and $t(x^0, \dots, x^{n-1})$ a term over Σ. Suppose that $\chi(x, y)$ is a quantifier-free formula such that $\mathrm{Av}(F)$ implies that χ linearly orders all elements of the form $t(s_I[\bar{\imath}])$ with $\bar{\imath} \in [I]^n$.*

Then there exist a linear order $\triangleleft$ *on* $[n]$ *and a map* $\sigma : [n] \to \{-1, 1\}$ *such that, for every linear order I and all tuples* $\bar{\imath}, \bar{\jmath} \in I^n$,

$$F(I) \vDash \chi\big(t(s_I[\bar{\imath}]), t(s_I[\bar{\jmath}])\big)$$

iff *there is some* $l \in [n]$ *with* $i_l <^{\sigma(l)} j_l$ *and* $i_s = j_s$, *for* $s \triangleleft l$,

where $<^1 := <$ *and* $<^{-1} := >$.

Proof. Note that we can embed every model $F(I)$ into a model $F(J)$ where J is a dense order. Since χ is quantifier-free it is therefore sufficient to consider the case of a dense order I. Define

$$\bar{\imath} \sqsubset \bar{\jmath} \quad :\text{iff} \quad F(I) \vDash \chi\big(t(s_I[\bar{\imath}]), t(s_I[\bar{\jmath}])\big) .$$

According to Lemma 3.19 the order $\sqsubset$ has the desired form. □

E7. *Abstract elementary classes*

1. *Abstract elementary classes*

For every algebraic logic L, we can form the category $\mathfrak{Emb}_L(\Sigma)$ of L-embeddings. This is a subcategory of the category $\mathfrak{Emb}(\Sigma)$ of all embeddings. It has the same objects but fewer morphisms. In this section we investigate to which extend these two categories determine L.

Definition 1.1. Suppose that $\mathcal{K}$ is a class of Σ-structures that is closed under isomorphisms and let $\mathcal{E}$ be a class of embeddings between structures in $\mathcal{K}$.

(a) The pair $\langle \mathcal{K}, \mathcal{E} \rangle$ forms an *abstract elementary class* if it satisfies the following conditions.

- (i) $\mathcal{E}$ is closed under composition and it contains all isomorphisms between structures in $\mathcal{K}$.
- (ii) $f, f \circ g \in \mathcal{E}$ implies $g \in \mathcal{E}$, for all embeddings f and g.
- (iii) The subcategory of $\mathfrak{Emb}(\mathcal{K})$ induced by $\mathcal{E}$ has direct limits and, for every directed diagram $D : I \to \mathcal{E}$, the direct limits of D in $\mathcal{E}$ and in $\mathfrak{Emb}(\Sigma)$ coincide.
- (iv) There exists a cardinal $\mathrm{ln}(\mathcal{K}) \geq |\Sigma| \oplus \aleph_0$ such that, for every structure $\mathfrak{M} \in \mathcal{K}$ and every set $X \subseteq M$, we can find a substructure $\mathfrak{C} \in \mathcal{K}$ of size $|C| \leq |X| \oplus \mathrm{ln}(\mathcal{K})$ such that $\langle\!\langle X \rangle\!\rangle_{\mathfrak{M}} \subseteq \mathfrak{C} \subseteq \mathfrak{M}$ and the inclusion map $\mathfrak{C} \to \mathfrak{M}$ belongs to $\mathcal{E}$.

The cardinal $\mathrm{ln}(\mathcal{K})$ is called the *Löwenheim number* of $\mathcal{K}$.

(b) Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class. The elements of $\mathcal{E}$ are called *$\mathcal{K}$-embeddings*. Usually, we drop the class $\mathcal{E}$ from our notation and just write $\mathcal{K}$ for $\langle \mathcal{K}, \mathcal{E} \rangle$.

(c) Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class and let $\mathfrak{A} \subseteq \mathfrak{B}$ be structures in $\mathcal{K}$. We define

$$\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{B} \quad :\text{iff} \quad \text{the inclusion map } i : \mathfrak{A} \to \mathfrak{B} \text{ belongs to } \mathcal{E}\,.$$

If $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{B}$ then we call $\mathfrak{A}$ a $\mathcal{K}$-*substructure* of $\mathfrak{B}$.

(d) The pair $\langle \mathcal{K}, \mathcal{E} \rangle$ forms an *algebraic class* if

(i) $\mathcal{E} = \operatorname{Emb}(\mathcal{K})$ is the set of all embeddings and

(ii) $\mathcal{K}$ is closed under isomorphisms, substructures, and direct limits of embeddings.

Example. (a) Every algebraic class $\langle \mathcal{K}, \mathcal{E} \rangle$ of Σ-structures is an abstract elementary class with Löwenheim number $\ln(\mathcal{K}) = |\Sigma| \oplus \aleph_0$.

(b) Let $L := \mathrm{FO}_{\kappa\aleph_0}$, let $T \subseteq L^0[\Sigma]$ be a theory, and let $\mathcal{E}$ be the class of all $L^{<\omega}$-embeddings between models of T. Then $\langle \operatorname{Mod}(T), \mathcal{E} \rangle$ is an abstract elementary class and the relation $\leq_{\mathcal{K}}$ coincides with the $L^{<\omega}$-substructure relation $\leq_{L^{<\omega}}$. The same holds for many other algebraic logics L.

Exercise 1.1. In (b) of the above example we have taken for $\mathcal{E}$ all embeddings that preserve every formula with finitely many free variables. What goes wrong if we take only those embeddings that also preserve formulae with infinitely many free variables?

Exercise 1.2. Let $\langle \mathcal{K}_i, \mathcal{E}_i \rangle$, $i \in I$, be a family of abstract elementary classes over the signature Σ. Show that the intersection $\langle \bigcap_i \mathcal{K}_i, \bigcap_i \mathcal{E}_i \rangle$ is an abstract elementary class with Löwenheim number $\sup_i \ln(\mathcal{K}_i)$.

Remark. (a) We have defined the $\mathcal{K}$-substructure relation $\leq_{\mathcal{K}}$ in terms of the class $\mathcal{E}$ of $\mathcal{K}$-embeddings. Conversely, $\leq_{\mathcal{K}}$ determines $\mathcal{E}$ since an embedding $h : \mathfrak{A} \to \mathfrak{B}$ belongs to $\mathcal{E}$ if and only if $\operatorname{rng} h \leq_{\mathcal{K}} \mathfrak{B}$.

(b) Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class and let $\mathcal{K}_0 \subseteq \mathcal{K}$ be the subclass of all structures of size at most $\ln(\mathcal{K})$. Every structure $\mathfrak{M} \in \mathcal{K}$ can be written as a direct limit $D : I \to \mathcal{E}$ of its $\mathcal{K}$-substructures in $\mathcal{K}_0$. Hence, $\mathcal{K}$ is the class of all direct limits of structures in $\mathcal{K}_0$. In particular, $\mathcal{K}_0$ and the restriction of $\mathcal{E}$ to $\mathcal{K}_0$ completely determine $\langle \mathcal{K}, \mathcal{E} \rangle$.

We have seen that many algebraic logics give rise to an abstract elementary class. Conversely, we can show that every such class arises from an algebraic logic in this way. To do so, we need the notion of a Galois type.

Definition 1.2. Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class. Let $\mathfrak{M} \in \mathcal{K}$ be a structure and $U \subseteq M$ a set of parameters.

We define the *Galois type* of a tuple $\bar{a} \subseteq M$ over U by

$$\mathrm{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{M}, U) := [\bar{a}, \mathfrak{M}, U]_{\approx}$$

where the equivalence relation $\approx$ is the transitive closure of the following relation $\sim$ on triples $\langle \bar{a}, \mathfrak{M}, U \rangle$ with $U, \bar{a} \subseteq M$. We set

$$\langle \bar{a}, \mathfrak{A}, U \rangle \sim \langle \bar{b}, \mathfrak{B}, V \rangle$$

iff $U = V$ and, for some $\mathfrak{M} \in \mathcal{K}$, there are $\mathcal{K}$-embeddings $f : \mathfrak{A}_0 \to \mathfrak{M}$ and $g : \mathfrak{B}_0 \to \mathfrak{M}$ where $\mathfrak{A}_0 \leq_{\mathcal{K}} \mathfrak{A}$ and $\mathfrak{B}_0 \leq_{\mathcal{K}} \mathfrak{B}$ are $\mathcal{K}$-substructures with $U \cup \bar{a} \subseteq A_0$ and $U \cup \bar{b} \subseteq B_0$ such that

$$f \restriction U = g \restriction U \quad \text{and} \quad f(\bar{a}) = g(\bar{b}) .$$

We write $S^{\bar{s}}_{\mathrm{Aut}}(U)$ for the set of all Galois types of $\bar{s}$-tuples over U.

Remark. (a) Let T be a first-order theory and $\mathrm{Mod}(T)$ the corresponding abstract elementary class. Then the Galois type of a tuple coincides with its first-order type.

(b) If an abstract elementary class $\mathcal{K}$ stems from an algebraic logic L then no L-formula can distinguish between tuples of the same Galois type. Hence, tuples with the same Galois type also have the same L-type. In general the converse fails.

(c) Below we will not consider Galois types over arbitrary parameters U. The set U will always be either empty or the universe of some $\mathcal{K}$-substructure $\mathfrak{U}$.

Proposition 1.3. *Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class of Σ-structures. There exists an algebraic logic L, a fragment $\Delta \subseteq L^{<\omega}[\Sigma]$, and a formula $\chi \in \Delta$ such that*

$$\mathcal{K} = \operatorname{Mod}_L(\chi) \quad \textit{and} \quad \mathcal{E} \textit{ is the class of all } \Delta\textit{-embeddings.}$$

Proof. For a set X of variables, we denote by Φ_X the set of all Galois types of X-tuples over the empty set. We start by defining the functor L. For a signature Γ and a set X of variables, we set

$$L[\Gamma, X] := \mathscr{P}(\Phi_X) \times \mathfrak{Sig}(\Sigma, \Gamma)\,,$$

and, for a morphism $\lambda \in \mathfrak{Sig}(\Gamma, \Gamma')$, we set

$$L[\lambda] : \langle \Psi, \mu \rangle \mapsto \langle \Psi, \lambda \circ \mu \rangle\,.$$

For a formula $\langle \Psi, \mu \rangle \in L[\Gamma, X]$, a Γ-structure $\mathfrak{A}$, and a tuple $\bar{a} \in A^X$, we define the satisfaction relation by

$$\mathfrak{A} \vDash \langle \Psi, \mu \rangle(\bar{a}) \quad :\text{iff} \quad \operatorname{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{A}|_\mu, \varnothing) \in \Psi\,.$$

Finally, we set

$$\Delta := \{\, \langle \Psi, \mu \rangle \in L^{<\omega}[\Sigma] \mid \mu = \mathrm{id} \,\} \quad \text{and} \quad \chi := \langle \Phi_\varnothing, \mathrm{id} \rangle\,. \qquad \square$$

This proposition provides a syntax for each abstract elementary class. But because of the high degree of generality in the definition of an algebraic logic, this result is of little practical use. A more concrete way of equipping an abstract elementary class with a kind of syntax is given by the notion of a Skolem expansion.

Definition 1.4. Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class of Σ-structures.

(a) An *expansion* of $\mathcal{K}$ is an abstract elementary class $\langle \mathcal{K}_+, \mathcal{E}_+ \rangle$ of Σ_+-structures, for some $\Sigma_+ \supseteq \Sigma$, such that

$$\operatorname{pr}_\Sigma(\mathcal{K}_+) = \mathcal{K}\,, \quad \operatorname{pr}_\Sigma(\mathcal{E}_+) = \mathcal{E}\,, \quad \text{and} \quad \ln(\mathcal{K}_+) = \ln(\mathcal{K})\,,$$

where $\mathrm{pr}_\Sigma : \mathfrak{Emb}(\Sigma_+) \to \mathfrak{Emb}(\Sigma)$ is the reduct functor.

(b) An expansion $\langle \mathcal{K}_+, \mathcal{E}_+ \rangle$ of $\langle \mathcal{K}, \mathcal{E} \rangle$ is a *Skolem expansion* if $\langle \mathcal{K}_+, \mathcal{E}_+ \rangle$ is an algebraic class.

Algebraic classes and, hence, Skolem expansions are very nicely behaved abstract elementary classes. For instance, the membership of a structure in such a class only depends on its finitely generated substructures.

Lemma 1.5. *Let $\mathcal{K}$ be an algebraic class and $\mathfrak{M}$ a structure. Then*

$$\mathfrak{M} \in \mathcal{K} \quad \text{iff} \quad \mathrm{Sub}_{\aleph_0}(\mathfrak{M}) \subseteq \mathcal{K}\,.$$

Proof. $(\Rightarrow)$ Suppose that $\mathcal{K}$ is algebraic, $\mathfrak{M} \in \mathcal{K}$, and $\mathfrak{A} \subseteq \mathfrak{M}$. Since $\mathcal{K}$ is algebraic, we have $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$. This implies that $\mathfrak{A} \in \mathcal{K}$.

$(\Leftarrow)$ Each structure $\mathfrak{M}$ can be written as direct limit $\mathfrak{M} = \varinjlim D$ where $D : I \to \mathrm{Sub}_{\aleph_0}(\mathfrak{M})$ is the diagram of the finitely generated substructures of $\mathfrak{M}$. By assumption we have $D(i) \in \mathcal{K}$, for every $i \in I$. Since $\mathcal{K}$ is algebraic it is closed under direct limits of embeddings. Consequently, we have $\mathfrak{M} = \varinjlim D \in \mathcal{K}$. ☐

As a corollary it follows that every algebraic class is $\forall_{\infty\aleph_0}$-axiomatisable.

Proposition 1.6. *Let Σ be a signature and set $\kappa := |\Sigma| \oplus \aleph_0$. Every algebraic class $\mathcal{K}$ of Σ-structures is $\forall_{(2^\kappa)^+\aleph_0}$-axiomatisable.*

Proof. Let

$$\mathcal{C}_n := \{ \langle \mathfrak{A}, \bar{a} \rangle \mid \mathfrak{A} \in \mathcal{K} \text{ is generated by } \bar{a} \in A^n \}$$

be the class of all structures in $\mathcal{K}$ that are generated by a set of size n. Note that every structure in $\mathcal{C}_n$ has size at most $\kappa = |\Sigma| \oplus \aleph_0$. Consequently, $\mathcal{C}_n$ contains, up to isomorphism, at most 2^κ structures. For every $\langle \mathfrak{A}, \bar{a} \rangle \in \mathcal{C}_n$, we can write down a quantifier-free formula $\varphi_{\mathfrak{A},\bar{a}}(\bar{x}) \in \mathrm{QF}^n_{\kappa^+\aleph_0}[\Sigma]$ such that

$$\mathfrak{B} \vDash \varphi_{\mathfrak{A},\bar{a}}(\bar{b}) \quad \text{iff} \quad \langle \langle\!\langle \bar{b} \rangle\!\rangle_{\mathfrak{B}}, \bar{b} \rangle \cong \langle \mathfrak{A}, \bar{a} \rangle\,.$$

By Lemma 1.5, it follows that the $\forall^{\circ}_{(2^\kappa)^+\aleph_0}[\Sigma]$-formula

$$\bigwedge_{n<\omega} \forall x_0 \cdots \forall x_{n-1} \bigvee_{\langle \mathfrak{A}, \bar{a} \rangle \in \mathcal{C}_n} \varphi_{\mathfrak{A},\bar{a}}(\bar{x})$$

axiomatises $\mathcal{K}$. ☐

If we can show that every abstract elementary class has a Skolem expansion, it follows that each such class is a projective $\forall_{\infty\aleph_0}$-class.

Theorem 1.7. *Let $\mathcal{K}$ be an abstract elementary class of Σ-structures. There exists a Skolem expansion $\mathcal{K}_+$ of $\mathcal{K}$ over a signature $\Sigma_+ \supseteq \Sigma$ of size $|\Sigma_+| = \ln(\mathcal{K})$.*

Proof. Let $\lambda := \ln(\mathcal{K})$ and set $\Sigma_+ := \Sigma \cup \{ f^n_\alpha \mid n < \omega,\ \alpha < \lambda \}$ where the f^n_α are new n-ary function symbols. We call a Σ_+-expansion $\mathfrak{M}_+$ of a structure $\mathfrak{M} \in \mathcal{K}$ *admissible* if

$$\mathfrak{A}|_\Sigma \leq_{\mathcal{K}} \mathfrak{M}, \quad \text{for every } \mathfrak{A} \subseteq \mathfrak{M}_+ .$$

We claim that the desired Skolem expansion $\langle \mathcal{K}_+, \mathcal{E}_+ \rangle$ is given by

$$\begin{aligned} \mathcal{K}_+ &:= \{ \mathfrak{M}_+ \mid \mathfrak{M}_+ \text{ an admissible expansion of some } \mathfrak{M} \in \mathcal{K} \}, \\ \mathcal{E}_+ &:= \mathrm{Emb}(\mathcal{K}_+) . \end{aligned}$$

Clearly, we have $\ln(\mathcal{K}_+) = |\Sigma_+| = \ln(\mathcal{K})$. Hence, it remains to prove the following claims.

Claim. (a) *For every pair $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{B}$ in $\mathcal{K}$, there exist admissible expansions $\mathfrak{A}_+$ and $\mathfrak{B}_+$ such that $\mathfrak{A}_+ \subseteq \mathfrak{B}_+$. In particular, we have $\mathrm{pr}_\Sigma(\mathcal{K}_+) = \mathcal{K}$.*

(b) $\mathrm{pr}_\Sigma(\mathcal{E}_+) = \mathcal{E}$.

(c) *$\mathcal{K}_+$ is closed under direct limits.*

(a) By induction on $n < \omega$, we can fix, for every subset $X \subseteq B$ of size n, a $\mathcal{K}$-substructure $\mathfrak{B}_X \leq_{\mathcal{K}} \mathfrak{B}$ of size at most λ containing $X \cup \bigcup_{Y \subset X} B_Y$.

Furthermore, if $X \subseteq A$ then we choose $\mathfrak{B}_X$ such that $B_X \subseteq A$. By construction, we have $\mathfrak{B}_X \subseteq \mathfrak{B}_Y$, for $X \subseteq Y$. Since $\mathfrak{B}_X, \mathfrak{B}_Y \leq_{\mathcal{K}} \mathfrak{B}$ this implies that $\mathfrak{B}_X \leq_{\mathcal{K}} \mathfrak{B}_Y$.

For every $\bar{a} \in B^n$, $n < \omega$, fix an enumeration $(c^{\bar{a}}_\alpha)_{\alpha<\lambda}$ (possibly with repetitions) of $B_{\bar{a}}$. To obtain the desired expansion $\mathfrak{B}_+$ we set $f^n_\alpha(\bar{a}) := c^{\bar{a}}_\alpha$, for $\bar{a} \in B^n$. Note that our construction ensures that A induces a substructure of $\mathfrak{B}_+$ since $\mathfrak{B}_X \subseteq \mathfrak{A}$, for $X \subseteq A$, implies that $\langle\!\langle X \rangle\!\rangle_{\mathfrak{B}_+} \subseteq A$. Therefore, we can set $\mathfrak{A}_+ := \mathfrak{B}_+|_A$.

To see that $\mathfrak{A}_+$ and $\mathfrak{B}_+$ are admissible, note that, by construction, we have $\mathfrak{B}_X \subseteq \langle\!\langle X \rangle\!\rangle_{\mathfrak{B}_+}|_\Sigma$, for every finite $X \subseteq B$. If $\mathfrak{C} \subseteq \mathfrak{B}_+$ is an arbitrary substructure then

$$\mathfrak{C}|_\Sigma = \varinjlim_{X \subseteq C \text{ finite}} \langle\!\langle X \rangle\!\rangle_{\mathfrak{C}}|_\Sigma = \varinjlim_{X \subseteq C \text{ finite}} \langle\!\langle X \rangle\!\rangle_{\mathfrak{B}_+}|_\Sigma = \varinjlim_{X \subseteq C \text{ finite}} \mathfrak{B}_X .$$

We have already seen that the $\mathfrak{B}_X$ form a directed system of $\mathcal{K}$-embeddings such that $\mathfrak{B}_X \leq_{\mathcal{K}} \mathfrak{B}$. Hence, the limit also satisfies $\mathfrak{C}|_\Sigma \leq_{\mathcal{K}} \mathfrak{B}$, as desired. Furthermore, if $\mathfrak{C} \subseteq \mathfrak{A}_+ \subseteq \mathfrak{B}_+$ then $\mathfrak{C}|_\Sigma, \mathfrak{A} \leq_{\mathcal{K}} \mathfrak{B}$ implies that $\mathfrak{C}|_\Sigma \leq_{\mathcal{K}} \mathfrak{A}$. Thus, $\mathfrak{A}_+$ and $\mathfrak{B}_+$ are admissible.

(b) ($\subseteq$) Let $h : \mathfrak{A}_+ \to \mathfrak{B}_+$ be a $\mathcal{K}_+$-embedding and set $C := \operatorname{rng} h$. Then C induces a substructure $\mathfrak{C}_+ \subseteq \mathfrak{B}_+$ and h induces an isomorphism $h' : \mathfrak{A}_+ \cong \mathfrak{C}_+$. The structure $\mathfrak{B}_+$ is an admissible expansion of some structure $\mathfrak{B} \in \mathcal{K}$. Hence, $\mathfrak{C}_+|_\Sigma \leq_{\mathcal{K}} \mathfrak{B}$ and the inclusion map $i : \mathfrak{C}_+|_\Sigma \to \mathfrak{B}$ belongs to $\mathcal{E}$. Since $\mathcal{E}$ contains all isomorphisms and it is closed under composition, it follows that $\operatorname{pr}_\Sigma(h) = i \circ \operatorname{pr}_\Sigma(h') \in \mathcal{E}$.

($\supseteq$) Let $h : \mathfrak{C} \to \mathfrak{B}$ be a $\mathcal{K}$-embedding. Setting $\mathfrak{A} := \operatorname{rng} h$ we can use (a) to find admissible expansions $\mathfrak{A}_+ \subseteq \mathfrak{B}_+$ of $\mathfrak{A}$ and $\mathfrak{B}$. Let $\mathfrak{C}_+$ be the expansion of $\mathfrak{C}$ that corresponds to $\mathfrak{A}_+$ via the isomorphism $h : \mathfrak{C} \cong \mathfrak{A}$. Then h induces an embedding $h_+ : \mathfrak{C}_+ \to \mathfrak{B}_+$. Since $\mathcal{K}_+$ is closed under isomorphisms we have $\mathfrak{C}_+ \in \mathcal{K}_+$. Hence, $h_+ \in \mathcal{E}_+$.

(c) Let $D : I \to \mathcal{K}_+$ be a directed diagram with limit $\mathfrak{M}_+ := \varinjlim D$. We have to show that $\mathfrak{M}_+ \in \mathcal{K}_+$. Let $p : \mathcal{K}_+ \to \mathcal{K}$ be the canonical projection functor and set $\mathfrak{M} := \mathfrak{M}_+|_\Sigma$. Then $p \circ D : I \to \mathcal{K}$ is a directed diagram with limit $\varinjlim (p \circ D) = \mathfrak{M}_+|_\Sigma = \mathfrak{M}$. By (b), it follows that $p \circ D$ is in fact a diagram $I \to \mathcal{E}$. Hence, the limit $\mathfrak{M}$ is in $\mathcal{K}$. We claim that $\mathfrak{M}_+$ is an admissible expansion of $\mathfrak{M}$. Let $\mathfrak{A} \subseteq \mathfrak{M}_+$ be a substructure. For every finite set $X \subseteq M$, there exists some i with $X \subseteq D(i)$. Since $D(i)$ is an admissible expansion it follows that

$$\langle\!\langle X \rangle\!\rangle_{D(i)}|_\Sigma \preceq_{\mathcal{K}} D(i)|_\Sigma \preceq_{\mathcal{K}} \varinjlim (p \circ D) = \mathfrak{M} .$$

The substructure $\mathfrak{A}$ is the direct limit of its finitely generated substructures $\mathfrak{X}$. We have just seen that $\mathfrak{X}|_\Sigma \preceq_{\mathcal{K}} \mathfrak{M}$, for all such $\mathfrak{X}$. By the definition of a direct limit, it follows that $\mathfrak{A}|_\Sigma = \varinjlim \mathfrak{X}|_\Sigma \preceq_{\mathcal{K}} \mathfrak{M}$. □

The existence of Skolem expansions enables us to apply the theory of Ehrenfeucht-Mostowski functors to abstract elementary classes. We will make extensive use of these functors in Section 4 below. As an example we use them in the remainder of this section to compute the Hanf number of a class.

Lemma 1.8. *Let $\mathcal{K}$ be an algebraic class of Σ-structures and set $\kappa := |\Sigma| \oplus \aleph_0$ and $\lambda := \beth_{(2^\kappa)^+}$. If $\mathcal{K}$ contains a structure of size at least λ then there exists an Ehrenfeucht-Mostowski functor $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$.*

Proof. Fix a structure $\mathfrak{M} \in \mathcal{K}$ of size $|M| \geq \lambda$ and let $(a_i)_{i<\lambda}$ be a sequence of distinct elements of M. Since $|S^{<\omega}(\varnothing)| \leq 2^\kappa$ we can apply Theorem E5.3.7 to $(a_i)_i$ to obtain an elementary extension $\mathfrak{M}_+ \succeq_{\mathrm{FO}} \mathfrak{M}$ that contains an indiscernible sequence $(b_i)_{i<\omega}$ such that, for all $n < \omega$ and every $\bar{\imath} \in [\omega]^n$, there is some $\bar{k} \in [\lambda]^n$ with

$$\mathrm{tp}(b[\bar{\imath}]) = \mathrm{tp}(a[\bar{k}]) .$$

Note that this implies in particular that $\langle\!\langle b[\bar{\imath}] \rangle\!\rangle_{\mathfrak{M}_+} \cong \langle\!\langle a[\bar{k}] \rangle\!\rangle_{\mathfrak{M}} \in \mathcal{K}$. By Proposition E6.3.8, there exists a unique strongly local functor $F : \mathfrak{Lin} \to \mathfrak{Emb}(\Sigma)$ such that $F(\omega) \cong \langle\!\langle (b_i)_i \rangle\!\rangle_{\mathfrak{M}_+}$. We claim that the range of F is contained in $\mathcal{K}$.

Let I be a linear order and consider a finitely generated substructure $\mathfrak{A} \subseteq F(I)$. Then there is a finite subset $I_0 \subseteq I$ such that $\mathfrak{A} \subseteq F(I_0)$. Consequently, for some $n < \omega$, $\mathfrak{A}$ is isomorphic to a substructure of

$$F(n) \cong \langle\!\langle b_0 \dots b_{n-1} \rangle\!\rangle_{\mathfrak{M}_+} \subseteq \mathfrak{M}_+ \in \mathcal{K} .$$

Since $\mathcal{K}$ is closed under substructures and isomorphisms, it follows that $\mathfrak{A} \in \mathcal{K}$. Hence, we have $\mathrm{Sub}_{\aleph_0}(F(I)) \subseteq \mathcal{K}$ which, by Lemma 1.5, implies that $F(I) \in \mathcal{K}$. Thus, $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ is the desired Ehrenfeucht-Mostowski functor. □

Using Skolem expansions we can extend this result to arbitrary abstract elementary classes.

Remark. Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class, $\mathcal{K}_+$ a Skolem expansion of $\mathcal{K}$, and $F_+ : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K}_+)$ an Ehrenfeucht-Mostowski functor. Composing F_+ with the reduct functor $\mathrm{pr}_\Sigma : \mathfrak{Emb}(\Sigma_+) \to \mathfrak{Emb}(\Sigma)$ we obtain a functor $F := \mathrm{pr}_\Sigma \circ F_+ : \mathfrak{Lin} \to \mathfrak{Emb}(\Sigma)$. By definition of a Skolem expansion, F is actually a functor $\mathfrak{Lin} \to \mathcal{E}$, i.e., it maps every embedding $I \to J$ of linear orders to a $\mathcal{K}$-embedding $F(I) \to F(J)$.

Definition 1.9. Let $\mathcal{K}$ be an abstract elementary class of Σ-structures and $\mathcal{K}_+$ a Skolem expansion of $\mathcal{K}$. An *Ehrenfeucht-Mostowski functor for* $\mathcal{K}$ is a functor $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ of the form $F = \mathrm{pr}_\Sigma \circ F_+$, where $F_+ : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K}_+)$ is an ordinary Ehrenfeucht-Mostowski functor.

Corollary 1.10. *Let $\mathcal{K}$ be an abstract elementary class and set $\kappa := 2^{\ln(\mathcal{K})}$. If $\mathcal{K}$ contains a structure of size at least $\beth_{\kappa^+}$, then there exists an Ehrenfeucht-Mostowski functor for $\mathcal{K}$.*

As promised we apply these results to compute the Hanf number of an abstract elementary class.

Definition 1.11. Let $\mathcal{K}$ be an arbitrary class of Σ-structures. The *Hanf number* of $\mathcal{K}$ is

$$\mathrm{hn}(\mathcal{K}) := \sup \{ |M|^+ \mid \mathfrak{M} \in \mathcal{K} \} .$$

If this supremum does not exist then we set $\mathrm{hn}(\mathcal{K}) := \infty$. In this case the class $\mathcal{K}$ is called *unbounded*.

Proposition 1.12. *Let $\mathcal{K}$ be an abstract elementary class of Σ-structures and set $\kappa := 2^{\ln(\mathcal{K})}$. We either have*

$$\mathrm{hn}(\mathcal{K}) \leq \beth_{\kappa^+} \quad \textit{or} \quad \mathrm{hn}(\mathcal{K}) = \infty .$$

Proof. Suppose that $\mathrm{hn}(\mathcal{K}) > \beth_{\kappa^+}$. By Corollary 1.10, there exists an Ehrenfeucht-Mostowski functor $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ for $\mathcal{K}$. For every cardinal λ, we have $F(\lambda) \in \mathcal{K}$. This implies that

$$\mathrm{hn}(\mathcal{K}) > |F(\lambda)| = \lambda \oplus \mathrm{ln}(\mathcal{K}) .$$

Consequently, $\mathrm{hn}(\mathcal{K}) = \infty$. □

With this proposition we are finally able to provide the missing part of the proof of Theorem C5.2.7. (Except that we do not obtain a strict inequality $\mathrm{hn}_1(\mathrm{FO}_{\kappa^+\aleph_0}) < \beth_{(2^\kappa)^+}$.)

Corollary 1.13. $\mathrm{hn}_1(\mathrm{FO}_{\kappa^+\aleph_0}) \leq \beth_{(2^\kappa)^+}$.

2. *Amalgamation and saturation*

In this section we consider saturated structures in abstract elementary classes. As we have already seen in the first-order case, an important ingredient in the construction of such structures is the amalgamation property.

Definition 2.1. Let $\langle \mathcal{K}, \mathcal{E} \rangle$ be an abstract elementary class.

(a) For a cardinal κ, we set

$$\mathcal{K}_\kappa := \{ \mathfrak{M} \in \mathcal{K} \mid |M| = \kappa \} \quad \text{and} \quad \mathcal{K}_{<\kappa} := \{ \mathfrak{M} \in \mathcal{K} \mid |M| < \kappa \} .$$

We define $\mathcal{K}_{>\kappa}$, $\mathcal{K}_{\leq\kappa}$, and $\mathcal{K}_{\geq\kappa}$ analogously.

(b) $\mathcal{K}$ has the *amalgamation property* if, for all $\mathcal{K}$-embeddings $f : \mathfrak{A} \to \mathfrak{B}$ and $g : \mathfrak{A} \to \mathfrak{C}$, there exist $\mathcal{K}$-embeddings $h : \mathfrak{B} \to \mathfrak{D}$ and $k : \mathfrak{C} \to \mathfrak{D}$ with $h \circ f = k \circ g$.

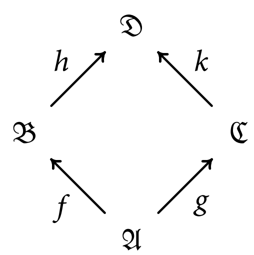

(c) $\mathcal{K}$ has the *joint embedding property* if, for all $\mathfrak{A}_0, \mathfrak{A}_1 \in \mathcal{K}$, there are $\mathcal{K}$-embeddings $\mathfrak{A}_0 \to \mathfrak{B}$ and $\mathfrak{A}_1 \to \mathfrak{B}$, for some $\mathfrak{B} \in \mathcal{K}$.

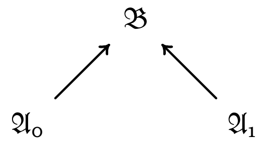

(d) An *amalgamation class* is an abstract elementary class with the amalgamation property. A *Jónsson class* is an abstract elementary class with the amalgamation property and the joint embedding property.

Example. Let T be an $\forall\exists$-theory and $\mathcal{K}$ the class of all existentially closed models of T. Then $\langle \mathcal{K}, \mathrm{Emb}(\mathcal{K}) \rangle$ forms an abstract elementary class with the amalgamation property.

In the same way that the class of all algebraically closed fields can be decomposed into the classes of algebraically closed fields of characteristic p, for the various p, we can write each amalgamation class as a union of Jónsson classes.

Lemma 2.2. *Every amalgamation class $\mathcal{K}$ is a disjoint union of at most $2^{\ln(\mathcal{K})}$ Jónsson classes.*

Proof. We define an equivalence relation on $\mathcal{K}$ by

$$\mathfrak{A} \sim \mathfrak{B} \quad :\text{iff} \quad \text{there are } \mathcal{K}\text{-embeddings } \mathfrak{A} \to \mathfrak{C} \text{ and } \mathfrak{B} \to \mathfrak{C}, \text{ for some } \mathfrak{C} \in \mathcal{K}.$$

Clearly, $\sim$ is reflexive and symmetric. For transitivity, let us assume that $\mathfrak{A}_0 \sim \mathfrak{A}_1 \sim \mathfrak{A}_2$. Then there are structures $\mathfrak{B}_0, \mathfrak{B}_1 \in \mathcal{K}$ and $\mathcal{K}$-embeddings $f_i : \mathfrak{A}_i \to \mathfrak{B}_0$, for $i \in \{0, 1\}$, and $g_i : \mathfrak{A}_i \to \mathfrak{B}_1$, for $i \in \{1, 2\}$.

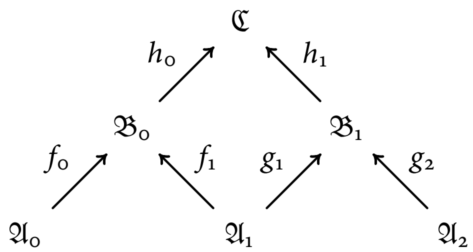

By the amalgamation property, we can find some structure $\mathfrak{C} \in \mathcal{K}$ and $\mathcal{K}$-embeddings $h_i : \mathfrak{B}_i \to \mathfrak{C}$, for $i < 2$, such that $h_0 \circ f_1 = h_1 \circ g_1$. Consequently, there are $\mathcal{K}$-embeddings $h_0 \circ f_0 : \mathfrak{A}_0 \to \mathfrak{C}$ and $h_1 \circ g_2 : \mathfrak{A}_2 \to \mathfrak{C}$. This implies that $\mathfrak{A}_0 \sim \mathfrak{A}_2$.

By definition, every $\sim$-class is a Jónsson class. Furthermore, $\mathfrak{A} \nsim \mathfrak{B}$ implies that there is no $\mathcal{K}$-embedding $\mathfrak{A} \to \mathfrak{B}$. Hence, $\mathcal{K}$ is the disjoint union of all $\sim$-classes. Finally, every $\sim$-class contains a structure of size at most $\ln(\mathcal{K})$. Consequently, there are at most $2^{\ln(\mathcal{K})}$ such classes. $\square$

For amalgamation classes, the definition of a Galois type can be simplified quite a bit.

Lemma 2.3. *Let $\mathcal{K}$ be an amalgamation class, $\mathfrak{A}, \mathfrak{B}, \mathfrak{U} \in \mathcal{K}$ structures with $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{A}, \mathfrak{B}$, and let $\bar{a} \subseteq A$ and $\bar{b} \subseteq B$. Then we have*

$$\mathrm{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{A}, U) = \mathrm{tp}_{\mathrm{Aut}}(\bar{b}/\mathfrak{B}, U)$$

if and only if there exists a structure $\mathfrak{M} \in \mathcal{K}$ of size $|M| \leq |A| \oplus |B| \oplus \ln(\mathcal{K})$ and $\mathcal{K}$-embeddings $g : \mathfrak{A} \to \mathfrak{M}$ and $h : \mathfrak{B} \to \mathfrak{M}$ such that

$$g \restriction U = h \restriction U \quad \text{and} \quad g(\bar{a}) = h(\bar{b})\,.$$

Proof. ($\Leftarrow$) is trivial. For ($\Rightarrow$), suppose that the Galois types are equal. Recall the relation $\sim$ from Definition 1.2. There exists a finite sequence $\langle \mathfrak{C}_0, \bar{c}_0 \rangle, \dots, \langle \mathfrak{C}_n, \bar{c}_n \rangle$ of structures such that

$$\langle \mathfrak{C}_0, \bar{c}_0 \rangle = \langle \mathfrak{A}, \bar{a} \rangle\,, \quad \langle \mathfrak{C}_n, \bar{c}_n \rangle = \langle \mathfrak{B}, \bar{b} \rangle\,,$$

$$\text{and} \quad \langle \bar{c}_i, \mathfrak{C}_i, U \rangle \sim \langle \bar{c}_{i+1}, \mathfrak{C}_{i+1}, U \rangle\,, \quad \text{for all } i < n\,.$$

We prove the claim by induction on n. For $n = 0$, we have $\mathfrak{A} = \mathfrak{B}$ and $\bar{a} = \bar{b}$, and there is nothing to do. Hence, suppose that $n > 0$. By inductive hypothesis, there exist a structure $\mathfrak{N}_0 \in \mathcal{K}$ and $\mathcal{K}$-embeddings $g_0 : \mathfrak{A} \to \mathfrak{N}_0$ and $h_0 : \mathfrak{C}_{n-1} \to \mathfrak{N}_0$ such that

$$g_0 \restriction U = h_0 \restriction U \quad \text{and} \quad g_0(\bar{a}) = h_0(\bar{c}_{n-1})\,.$$

Furthermore, by definition of $\sim$, we can find a structure $\mathfrak{N}_1 \in \mathcal{K}$, $\mathcal{K}$-substructures $\mathfrak{D} \leq_{\mathcal{K}} \mathfrak{C}_{n-1}$ and $\mathfrak{B}_0 \leq_{\mathcal{K}} \mathfrak{B}$ with $U \cup \bar{c}_{n-1} \subseteq D$ and $U \cup \bar{b} \subseteq B_0$, and $\mathcal{K}$-embeddings $g_1 : \mathfrak{D} \to \mathfrak{N}_1$ and $h_1 : \mathfrak{B}_0 \to \mathfrak{N}_1$ such that

$$g_1 \restriction U = h_1 \restriction U \quad \text{and} \quad g_1(\bar{c}_{n-1}) = h_1(\bar{b})\,.$$

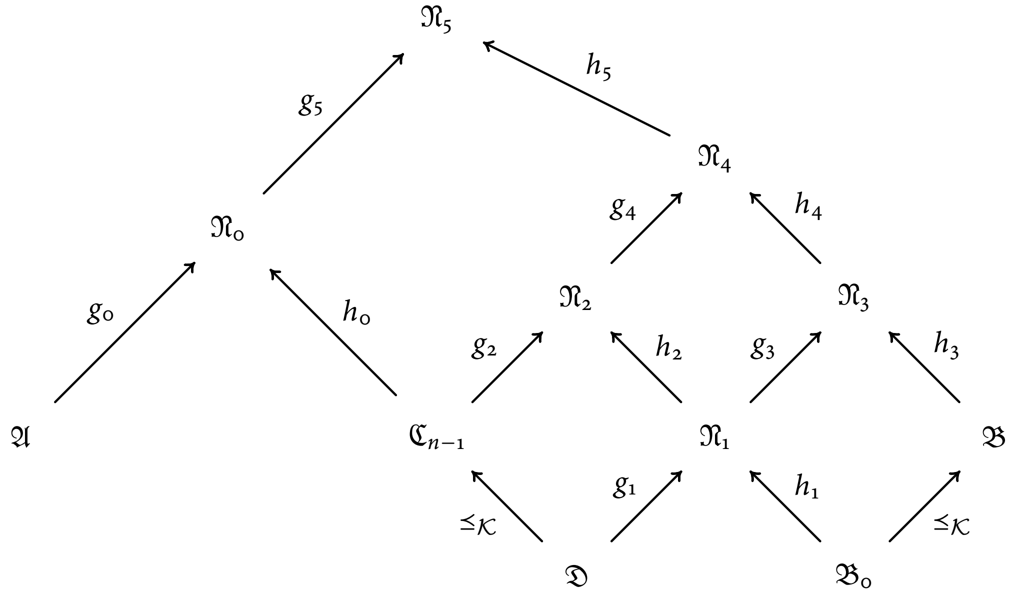

By the amalgamation property, there exist structures $\mathfrak{N}_2, \mathfrak{N}_3, \mathfrak{N}_4, \mathfrak{N}_5 \in \mathcal{K}$ such that we can complete the above diagram. Setting $g := g_5 \circ g_0$ and $h := h_5 \circ h_4 \circ h_3$ it follows that

$$g \restriction U = h \restriction U \quad \text{and} \quad g(\bar{a}) = h(\bar{b})\,.$$

Choosing a $\mathcal{K}$-substructure $\mathfrak{M} \leq_{\mathcal{K}} \mathfrak{N}_5$ of size $|M| \leq |A| \oplus |B| \oplus \ln(\mathcal{K})$ with $\operatorname{rng} g \cup \operatorname{rng} h \subseteq M$ the claim follows. □

Next, we introduce a notion of saturation for abstract elementary classes.

Definition 2.4. Let $\mathcal{K}$ be an abstract elementary class and let $\kappa \geq \ln(\mathcal{K})$ be a cardinal.

(a) A structure $\mathfrak{U} \in \mathcal{K}$ is *κ-universal* (for $\mathcal{K}$) if, for all $\mathfrak{A} \in \mathcal{K}_{<\kappa}$, there exists a $\mathcal{K}$-embedding $\mathfrak{A} \to \mathfrak{U}$. We call $\mathfrak{U}$ *$\mathcal{K}$-universal* if it is $|U|^+$-universal for $\mathcal{K}$.

(b) Similarly, we say that a structure $\mathfrak{U} \in \mathcal{K}$ is κ-universal *over a substructure* $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{U}$ if, for all $\mathcal{K}$-embeddings $f : \mathfrak{A} \to \mathfrak{B}$ with $|B| < \kappa$, there exists a $\mathcal{K}$-embedding $g : \mathfrak{B} \to \mathfrak{U}$ such that $g \circ f = \mathrm{id}_A$.

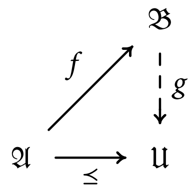

(c) A structure $\mathfrak{J} \in \mathcal{K}$ is *κ-injective* (for $\mathcal{K}$), or *κ-model homogeneous,* if, for all $\mathcal{K}$-embeddings $f : \mathfrak{A} \to \mathfrak{J}$ and $g : \mathfrak{A} \to \mathfrak{B}$ with $|A|, |B| < \kappa$, there exists a $\mathcal{K}$-embedding $h : \mathfrak{B} \to \mathfrak{J}$ with $h \circ g = f$.

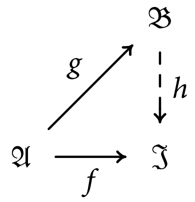

$\mathfrak{J}$ is called *$\mathcal{K}$-injective* if it is $|I|$-injective.

Remark. Note that a structure $\mathfrak{M}$ is κ-injective if and only if it is κ-universal over every substructure $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$ of size $|A| < \kappa$.

We can characterise κ-injective structures also by a back-and-forth condition.

Definition 2.5. Let $\mathcal{K}$ be an abstract elementary class and $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$.

(a) We denote by $I^\kappa_{\mathcal{K}}(\mathfrak{A}, \mathfrak{B})$ the set of all $\mathcal{K}$-embeddings $f : \mathfrak{A}_0 \to \mathfrak{B}_0$ between $\mathcal{K}$-substructures $\mathfrak{A}_0 \leq_{\mathcal{K}} \mathfrak{A}$ and $\mathfrak{B}_0 \leq_{\mathcal{K}} \mathfrak{B}$ of size $|A_0|, |B_0| < \kappa$.

(b) We write

$$\mathfrak{A} \sqsubseteq^\kappa_{\mathcal{K}} \mathfrak{B} \quad :\text{iff} \quad I^\kappa_{\mathcal{K}}(\mathfrak{A}, \mathfrak{B}) : \mathfrak{A} \sqsubseteq^\kappa_{\text{iso}} \mathfrak{B}\,,$$

$$\text{and} \quad \mathfrak{A} \cong^\kappa_{\mathcal{K}} \mathfrak{B} \quad :\text{iff} \quad I^\kappa_{\mathcal{K}}(\mathfrak{A}, \mathfrak{B}) : \mathfrak{A} \cong^\kappa_{\text{iso}} \mathfrak{B}\,.$$

In Lemma E1.2.2 we have characterised κ-saturated models in terms of the relation $\sqsubseteq_{\mathrm{FO}}^{\kappa}$. The next lemma gives a similar characterisation of κ-injective structures.

Lemma 2.6. *Let $\mathcal{K}$ be an abstract elementary class and $\kappa > \ln(\mathcal{K})$ a cardinal. A structure $\mathfrak{M} \in \mathcal{K}$ is κ-injective if and only if*

$$\mathfrak{A} \sqsubseteq_{\mathcal{K}}^{\kappa} \mathfrak{M}\,, \quad \textit{for all } \mathfrak{A} \in \mathcal{K} \textit{ with } I_{\mathcal{K}}^{\kappa}(\mathfrak{A}, \mathfrak{M}) \neq \varnothing\,.$$

Proof. ($\Leftarrow$) Suppose that $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}_{<\kappa}$ are structures with $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{B}$, and let $f : \mathfrak{A} \to \mathfrak{M}$ be a $\mathcal{K}$-embedding. Then $f \in I_{\mathcal{K}}^{\kappa}(\mathfrak{B}, \mathfrak{M})$. Since $|B| < \kappa$, we can use Lemma C4.4.9 (b) to find a $\mathcal{K}$-embedding $g \in I_{\mathcal{K}}^{\kappa}(\mathfrak{B}, \mathfrak{M})$ with $\operatorname{dom} g = B$ and $g \restriction A = f$.

($\Rightarrow$) By assumption, $I_{\mathcal{K}}^{\kappa}(\mathfrak{A}, \mathfrak{M})$ is nonempty. It has the forth property since $\mathfrak{M}$ is κ-injective. Furthermore, $I_{\mathcal{K}}^{\kappa}(\mathfrak{M}, \mathfrak{A})$ is $\ln(\mathcal{K})^+$-bounded. Finally, the closure of $\mathcal{K}$-embeddings under direct limits implies that $I_{\mathcal{K}}^{\kappa}(\mathfrak{A}, \mathfrak{M})$ is κ-complete. ☐

As usual we can use Lemma C4.4.9 to prove that, up to isomorphism, $\mathcal{K}$-injective structures are uniquely determined by their cardinality.

Proposition 2.7. *Let $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$ be two $\mathcal{K}$-injective structures with $|A| = |B|$. Then*

$$I_{\mathcal{K}}(\mathfrak{A}, \mathfrak{B}) \neq \varnothing \quad \textit{implies} \quad \mathfrak{A} \cong \mathfrak{B}\,.$$

The existence of κ-injective structures implies a weak form of the amalgamation property.

Lemma 2.8. *Let $\mathcal{K}$ be an abstract elementary class and suppose that $\mathfrak{M} \in \mathcal{K}$ is κ-injective, for some $\kappa > \ln(\mathcal{K})$.*

(a) *The class of all $\mathcal{K}$-substructures $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$ with $|A| < \kappa$ has the amalgamation property.*

(b) *If $\mathcal{K}$ has the joint embedding property, then $\mathfrak{M}$ is κ^+-universal.*

(c) *If $\mathcal{K}$ has the joint embedding property, then the subclass $\mathcal{K}_{<\kappa}$ has the amalgamation property.*

Proof. (a) Let $f : \mathfrak{A} \to \mathfrak{B}$ and $g : \mathfrak{A} \to \mathfrak{C}$ be $\mathcal{K}$-embeddings with $\mathfrak{A}, \mathfrak{B}, \mathfrak{C} \leq_{\mathcal{K}} \mathfrak{M}$ and $|A|, |B|, |C| < \kappa$. Replacing $\mathfrak{A}$ by an isomorphic copy, we may assume that $g = \mathrm{id}_A$. Since $\mathfrak{M}$ is κ-injective, there exists a $\mathcal{K}$-embedding $h : \mathfrak{B} \to \mathfrak{M}$ with $h \circ f = \mathrm{id}_A$. Let $\mathfrak{D} \leq_{\mathcal{K}} \mathfrak{M}$ be a substructure containing $C \cup \operatorname{rng} h$. Then we can use $h : \mathfrak{B} \to \mathfrak{D}$ and $\mathrm{id}_C : \mathfrak{C} \to \mathfrak{D}$ to complete the amalgamation diagram.

(b) As a first step, we show that $\mathfrak{M}$ is κ-universal. Let $\mathfrak{A}$ be some structure of size $|A| < \kappa$. We can use the joint embedding property to find $\mathcal{K}$-embeddings $f : \mathfrak{M} \to \mathfrak{N}$ and $g : \mathfrak{A} \to \mathfrak{N}$, for some $\mathfrak{N} \in \mathcal{K}$.

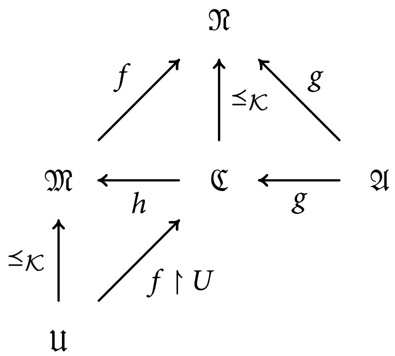

Choose a $\mathcal{K}$-substructure $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$ of size $|U| < \kappa$ and let $\mathfrak{C} \leq_{\mathcal{K}} \mathfrak{N}$ be a $\mathcal{K}$-substructure of size $|C| < \kappa$ with $f[U] \cup g[A] \subseteq C$. Since $\mathfrak{M}$ is κ-injective, there exists a $\mathcal{K}$-embedding $h : \mathfrak{C} \to \mathfrak{M}$ with $h \circ f \restriction U = \mathrm{id}_U$. The composition $h \circ g$ is a $\mathcal{K}$-embedding $\mathfrak{A} \to \mathfrak{M}$.

It remains to show that $\mathfrak{M}$ is even κ^+-universal. Let $\mathfrak{A}$ be a structure of size $|A| = \kappa$. Fix an increasing chain $(\mathfrak{C}_\alpha)_{\alpha<\kappa}$ of $\mathcal{K}$-substructures $\mathfrak{C}_\alpha \leq_{\mathcal{K}} \mathfrak{A}$ of size $|C_\alpha| < \kappa$ such that $\mathfrak{A} = \bigcup_{\alpha<\kappa} \mathfrak{C}_\alpha$. By induction on α, we construct $\mathcal{K}$-embeddings $f_\alpha : \mathfrak{C}_\alpha \to \mathfrak{M}$ such that $f_\beta \restriction C_\alpha = f_\alpha$, for all $\alpha \leq \beta$. We have already shown that $\mathfrak{M}$ is κ-universal. Hence, there exists a $\mathcal{K}$-embedding $f_0 : \mathfrak{C}_0 \to \mathfrak{M}$ which we can start our induction with. For limit ordinals δ, we set $f_\delta := \bigcup_{\alpha<\delta} f_\alpha$. For the successor step, suppose that we have already defined $f_\alpha : \mathfrak{C}_\alpha \to \mathfrak{M}$. Since $\mathfrak{M}$ is κ-injective, there exists a $\mathcal{K}$-embedding $f_{\alpha+1} : \mathfrak{C}_{\alpha+1} \to \mathfrak{M}$ such that $f_{\alpha+1} \restriction C_\alpha = f_\alpha$.

Having defined the family $(f_\alpha)_\alpha$ we can use the properties of a direct limit to find a $\mathcal{K}$-embedding $h : \bigcup_\alpha \mathfrak{C}_\alpha \to \mathfrak{M}$ such that $h \restriction C_\alpha = f_\alpha$, for all α. This is the desired $\mathcal{K}$-embedding $\mathfrak{A} \to \mathfrak{M}$.

(c) Let $f : \mathfrak{A} \to \mathfrak{B}$ and $g : \mathfrak{A} \to \mathfrak{C}$ be $\mathcal{K}$-embeddings with $|A|, |B|, |C| < \kappa$. By (b), we may assume that $\mathfrak{A}, \mathfrak{B}, \mathfrak{C} \leq_{\mathcal{K}} \mathfrak{M}$. Hence, we can use (a) to complete f and g to an amalgamation diagram. □

κ-injective structures generalise the characterisation of κ-saturated structures in terms of the relation $\sqsubseteq^{\kappa}_{\mathrm{FO}}$. We can also generalise the original definition of κ-saturation in terms of types. It turns out that, for amalgamation classes, these two notions coincide.

Definition 2.9. Let $\mathcal{K}$ be an abstract elementary class.

(a) A structure $\mathfrak{M} \in \mathcal{K}$ is *κ-Galois saturated* if it realises every Galois type in $S^{<\omega}_{\mathrm{Aut}}(U)$ where $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$ is a substructure of size $|U| < \kappa$. As usual we say that $\mathfrak{M}$ is *Galois saturated* if it is $|M|$-Galois saturated.

(b) $\mathcal{K}$ is *κ-Galois stable* if $|S^{<\omega}_{\mathrm{Aut}}(U)| \leq \kappa$, for all $\mathfrak{U} \in \mathcal{K}_{\leq\kappa}$.

Remark. Note that in the definition of κ-Galois stability we only count the Galois types over $\mathcal{K}$-substructures, not over arbitrary subsets. In general, this does make a difference.

The following lemma is the main ingredient in showing that κ-Galois saturated structures are κ-injective. We state it in a slightly more general form than needed here, since we will use it again in Section 3.

Lemma 2.10. *Let $\mathcal{K}$ be an amalgamation class and $\gamma \geq \ln(\mathcal{K})$ an ordinal. Suppose that $(\mathfrak{M}_\alpha)_{\alpha<\gamma}$ is an increasing chain such that each structure $\mathfrak{M}_{\alpha+1}$ realises every Galois type $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(U)$ where $\mathfrak{U} \leq \mathfrak{M}_\alpha$ is some substructure of size $|U| \leq |\gamma|$.*

Then the limit $\mathfrak{M} := \bigcup_{\alpha<\gamma} \mathfrak{M}_\alpha$ is $|\gamma|^+$-universal over every substructure $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}_0$ of size $|A| \leq |\gamma|$.

Proof. Let $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}_0$ be of size $|A| \leq |\gamma|$. To show that $\mathfrak{M}$ is $|\gamma|^+$-universal over $\mathfrak{A}$, we consider a $\mathcal{K}$-embedding $f : \mathfrak{A} \to \mathfrak{B}$ with $|B| \leq |\gamma|$. Set $\lambda := |B| \oplus \ln(\mathcal{K})$ and fix an enumeration $(b_\alpha)_{\alpha<\lambda}$ of B. We construct two increasing chains $(\mathfrak{A}_\alpha)_{\alpha<\lambda}$ and $(\mathfrak{C}_\alpha)_{\alpha<\lambda}$ of structures with $\mathfrak{B} \leq_{\mathcal{K}} \mathfrak{C}_\alpha$ and $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{A}_\alpha \leq_{\mathcal{K}} \mathfrak{M}_\alpha$, and an increasing chain $(h_\alpha)_{\alpha<\lambda}$ of $\mathcal{K}$-embeddings $h_\alpha : \mathfrak{A}_\alpha \to \mathfrak{C}_\alpha$ such that

$$|A_\alpha| \leq \lambda\,, \quad f \subseteq h_\alpha\,, \quad \text{and} \quad b_\alpha \in \operatorname{rng} h_{\alpha+1}\,.$$

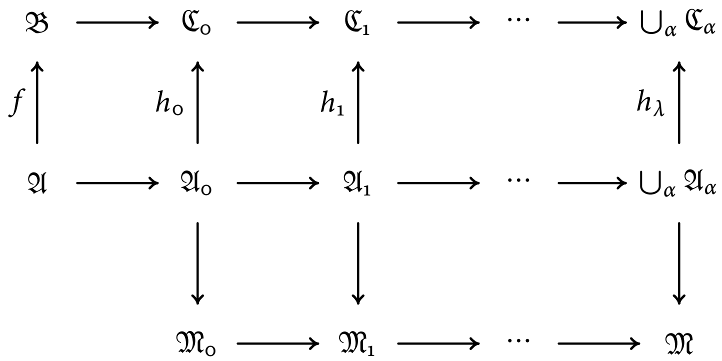

Then we obtain the desired embedding $g : \mathfrak{B} \to \mathfrak{M}$ by taking the limit $h_\lambda := \bigcup_{\alpha<\lambda} h_\alpha$ and setting $g := h_\lambda^{-1} \restriction B$.

We start with $\mathfrak{A}_0 := \mathfrak{A}$, $\mathfrak{C}_0 := \mathfrak{B}$, and $h_0 := f$. For limit ordinals δ, we take limits:

$$\mathfrak{A}_\delta := \bigcup_{\alpha<\delta} \mathfrak{A}_\alpha\,, \quad \mathfrak{C}_\delta := \bigcup_{\alpha<\delta} \mathfrak{C}_\alpha\,, \quad \text{and} \quad h_\delta := \bigcup_{\alpha<\delta} h_\alpha\,.$$

For the successor step, suppose that $h_\alpha : \mathfrak{A}_\alpha \to \mathfrak{C}_\alpha$ has already been defined. If $b_\alpha \in \operatorname{rng} h_\alpha$, we simply set $h_{\alpha+1} := h_\alpha$. Otherwise, we use amalgamation to find a $\mathcal{K}$-extension $\mathfrak{N} \geq_{\mathcal{K}} \mathfrak{M}$ and a $\mathcal{K}$-embedding $g : \mathfrak{C}_\alpha \to \mathfrak{N}$ with $g \circ h_\alpha = \mathrm{id}$.

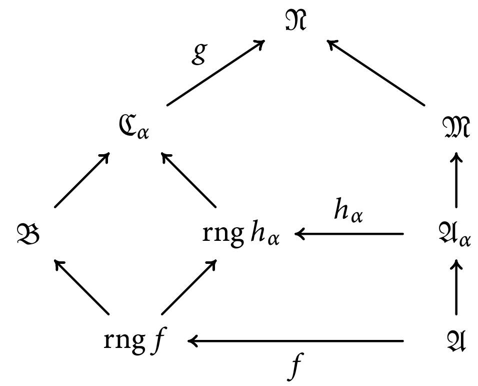

By assumption on $\mathfrak{M}_{\alpha+1}$, there is some element $c \in M_{\alpha+1}$ with

$$\operatorname{tp}_{\mathrm{Aut}}(c/\mathfrak{N}, A_\alpha) = \operatorname{tp}_{\mathrm{Aut}}(g(b_\alpha)/\mathfrak{N}, A_\alpha)\,.$$

By Lemma 2.3, this implies that there is a $\mathcal{K}$-extension $\mathfrak{N}^+ \geq_{\mathcal{K}} \mathfrak{N}$ and a $\mathcal{K}$-embedding $\sigma : \mathfrak{N} \to \mathfrak{N}^+$ such that

$$\sigma \restriction A_\alpha = \mathrm{id} \quad \text{and} \quad \sigma(g(b_\alpha)) = c\,.$$

We choose a $\mathcal{K}$-substructure $\mathfrak{A}_{\alpha+1} \leq_{\mathcal{K}} \mathfrak{M}_{\alpha+1}$ of size $|A_{\alpha+1}| \leq \lambda$ containing A_α and c. Let $\mathfrak{C}'_{\alpha+1} \leq_{\mathcal{K}} \mathfrak{N}^+$ be a $\mathcal{K}$-substructure containing $\mathrm{rng}(\sigma \circ g)$ and $A_{\alpha+1}$, and let $\mathfrak{C}_{\alpha+1}$ be the isomorphic copy of $\mathfrak{C}'_{\alpha+1}$ where each element of $\mathrm{rng}(\sigma \circ g)$ is replaced by its preimage. We denote the corresponding isomorphism $\mathfrak{C}'_{\alpha+1} \to \mathfrak{C}_{\alpha+1}$ by π. It follows that $\mathfrak{C}_\alpha \leq_{\mathcal{K}} \mathfrak{C}_{\alpha+1}$. We claim that the restriction $h_{\alpha+1} := \pi \restriction A_{\alpha+1}$ is the desired $\mathcal{K}$-embedding $\mathfrak{A}_{\alpha+1} \to \mathfrak{C}_{\alpha+1}$. Note that

$$b_\alpha = \pi((\sigma \circ g)(b_\alpha)) = \pi(c) \in \mathrm{rng}\, h_{\alpha+1}\,.$$

Furthermore, $\sigma \restriction A_\alpha = \mathrm{id}_{A_\alpha} = g \circ h_\alpha \restriction A_\alpha$ implies for $a \in A_\alpha$ that

$$h_{\alpha+1}(a) = \pi(a) = \pi(\sigma(a)) = \pi(\sigma((g \circ h_\alpha)(a))) = h_\alpha(a)\,.$$

Hence, $h_\alpha \subseteq h_{\alpha+1}$. □

Theorem 2.11. *Let $\mathcal{K}$ be an amalgamation class and $\kappa > \ln(\mathcal{K})$. A structure $\mathfrak{M} \in \mathcal{K}$ is κ-Galois saturated if and only if it is κ-injective.*

Proof. ($\Leftarrow$) Let $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$ be a substructure of size $|U| < \kappa$ and let $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(U)$ be a type. There exists an extension $\mathfrak{A} \geq_{\mathcal{K}} \mathfrak{U}$ realising $\mathfrak{p}$. We can choose $\mathfrak{A}$ of size $|A| \leq |U| \oplus \ln(\mathcal{K}) < \kappa$. Since $\mathfrak{M}$ is κ-injective, we can extend the $\mathcal{K}$-embedding $\mathfrak{U} \to \mathfrak{M}$ to a $\mathcal{K}$-embedding $\mathfrak{A} \to \mathfrak{M}$. Consequently, $\mathfrak{p}$ is realised in $\mathfrak{M}$.

($\Rightarrow$) Suppose that $f : \mathfrak{A} \to \mathfrak{B}$ is a $\mathcal{K}$-embedding with $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$ and $\lambda := |B| < \kappa$. For $\alpha < \lambda$, we set $\mathfrak{M}_\alpha := \mathfrak{M}$. Then $(\mathfrak{M}_\alpha)_{\alpha<\lambda}$ is an increasing chain satisfying the hypothesis of Lemma 2.10. It follows that the limit $\bigcup_{\alpha<\lambda} \mathfrak{M}_\alpha = \mathfrak{M}$ is λ^+-universal over $\mathfrak{A}$. Consequently, there exists a $\mathcal{K}$-embedding $g : \mathfrak{B} \to \mathfrak{M}$ with $g \circ f \restriction A = \mathrm{id}$. □

The next lemma shows that Galois saturated structures are strongly homogeneous.

Lemma 2.12. *Let $\mathcal{K}$ be an amalgamation class, suppose that $\mathfrak{M} \in \mathcal{K}$ is a Galois saturated structure of size $|M| = \kappa$, and let $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$ be a substructure of size $\ln(\mathcal{K}) \leq |U| < \kappa$. For $\bar{a}, \bar{b} \in M^{<\kappa}$, we have*

$$\mathrm{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{M}, U) = \mathrm{tp}_{\mathrm{Aut}}(\bar{b}/\mathfrak{M}, U)$$

if and only if there exists an automorphism $\pi \in \mathrm{Aut}\,\mathfrak{M}$ with $\pi \upharpoonright U = \mathrm{id}_U$ and $\pi(\bar{a}) = \bar{b}$.

Proof. It is sufficient to find an embedding $p \in \mathcal{I}^{\kappa}_{\mathcal{K}}(\mathfrak{M}, \mathfrak{M})$ with $p \upharpoonright U = \mathrm{id}_U$ and $p(\bar{a}) = \bar{b}$. Since $\mathfrak{M} \cong^{\kappa}_{\mathcal{K}} \mathfrak{M}$ we can then use Lemma C4.4.9 to extend p to the desired isomorphism $\pi : \mathfrak{M} \to \mathfrak{M}$.

Fix $\mathcal{K}$-substructures $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$ and $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{B} \leq_{\mathcal{K}} \mathfrak{M}$ of size $|A|, |B| < \kappa$ with $\bar{a} \subseteq A$ and $\bar{b} \subseteq B$. Since

$$\mathrm{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{M}, U) = \mathrm{tp}_{\mathrm{Aut}}(\bar{b}/\mathfrak{M}, U)\,,$$

we can use Lemma 2.3 to find $\mathcal{K}$-embeddings $f, g : \mathfrak{M} \to \mathfrak{N}$ with $f \upharpoonright U = g \upharpoonright U$ and $f(\bar{a}) = g(\bar{b})$.

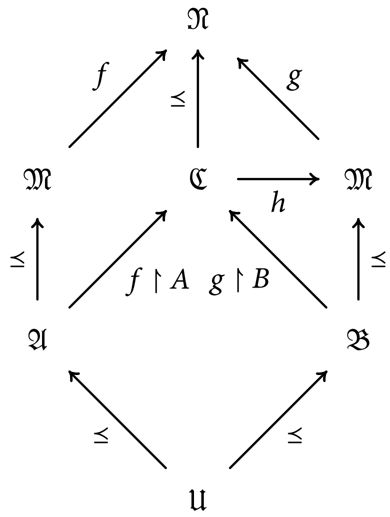

Let $\mathfrak{C} \leq_{\mathcal{K}} \mathfrak{N}$ be a $\mathcal{K}$-substructure of size $|C| < \kappa$ with $f[A] \cup g[B] \subseteq C$. Since $\mathfrak{M}$ is κ-injective, there exists a $\mathcal{K}$-embedding $h : \mathfrak{C} \to \mathfrak{M}$ with $h \circ g \upharpoonright B = \mathrm{id}_B$. Setting $p := h \circ f \upharpoonright A$ we have

$$p \upharpoonright U = h \circ f \upharpoonright U = h \circ g \upharpoonright U = \mathrm{id}_U\,,$$

and $\quad p(\bar{a}) = h(f(\bar{a})) = h(g(\bar{b})) = \bar{b}\,.$ □

When amalgamation is available we can construct κ-Galois saturated structures in the same way as κ-saturated ones. The main step in the inductive construction is the following lemma.

Lemma 2.13. *Let $\mathcal{K}$ be an amalgamation class. Every $\mathfrak{M} \in \mathcal{K}$ has an extension $\mathfrak{M}^+ \geq_{\mathcal{K}} \mathfrak{M}$ that realises every Galois type over $\mathfrak{M}$. If $\mathcal{K}$ is κ-stable, for $\kappa := |M| \oplus \ln(\mathcal{K})$, then we can choose $\mathfrak{M}^+$ of size $|M^+| \leq \kappa$.*

Proof. Let $(\mathfrak{p}_i)_{i<\lambda}$ be an enumeration of $S^{<\omega}_{\text{Aut}}(M)$. For every $i < \lambda$, we can find an extension $\mathfrak{A}_i \geq_{\mathcal{K}} \mathfrak{M}$ of size $|A_i| \leq |M| \oplus \ln(\mathcal{K}) = \kappa$ realising $\mathfrak{p}_i$. We construct $\mathfrak{M}^+$ as the limit of an increasing chain $(\mathfrak{B}_i)_{i<\lambda}$ where the structure $\mathfrak{B}_\alpha$ realises all types $\mathfrak{p}_i$ with $i < \alpha$. We start with $\mathfrak{B}_0 := \mathfrak{M}$. For limit ordinals δ, we set $\mathfrak{B}_\delta := \bigcup_{i<\delta} \mathfrak{B}_i$. For successor ordinals $\alpha = \beta + 1$, we use the amalgamation property to find an extension $\mathfrak{B}_\alpha \geq_{\mathcal{K}} \mathfrak{B}_\beta$ of size $|B_\alpha| \leq |B_\beta| \oplus |A_\beta| \oplus \ln(\mathcal{K})$ such that there exists a $\mathcal{K}$-embedding $h : \mathfrak{A}_\beta \to \mathfrak{B}_\alpha$ with $h \restriction M = \text{id}$.

We obtain the desired extension of $\mathfrak{M}$ by setting $\mathfrak{M}^+ := \bigcup_{i<\lambda} \mathfrak{B}_i$. By induction on α, it follows that $|B_\alpha| \leq \kappa \otimes |\alpha + 1|$. In particular, $|M^+| \leq \kappa \otimes \lambda$. Hence, if $\mathcal{K}$ is κ-stable then we have $\lambda \leq \kappa$ and $|M^+| = \kappa$. □

Iterating the construction of the preceding lemma, we obtain the desired Galois saturated extension. For the proof that the limit really is Galois saturated, we need the following technical lemma.

Definition 2.14. Let $\mathfrak{p} \in S^{<\omega}_{\text{Aut}}(B)$ be a Galois type and let $f : \mathfrak{A} \to \mathfrak{B}$ be a $\mathcal{K}$-embedding. We define the *restriction* $\mathfrak{p}|_f$ of $\mathfrak{p}$ along f as follows.

Fix a structure $\mathfrak{N} \geq_{\mathcal{K}} \mathfrak{B}$ containing a tuple $\bar{a} \subseteq N$ with

$$\mathfrak{p} = \text{tp}_{\text{Aut}}(\bar{a}/\mathfrak{N}, B)\,.$$

Let $\mathfrak{M}$ be the isomorphic copy of $\mathfrak{N}$ obtained by replacing all elements of $\text{rng}\, f$ by their preimages in A, and let $\pi : \mathfrak{N} \to \mathfrak{M}$ be the corresponding isomorphism. We set

$$\mathfrak{p}|_f := \text{tp}_{\text{Aut}}(\pi(\bar{a})/\mathfrak{M}, A)\,.$$

If $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{B}$ and $f : \mathfrak{A} \to \mathfrak{B}$ is the inclusion map, then we also write $\mathfrak{p}|_A$ for $\mathfrak{p}|_f$.

Lemma 2.15. *Let $\mathcal{K}$ be an amalgamation class and $f : \mathfrak{A} \to \mathfrak{B}$ a $\mathcal{K}$-embedding. For every Galois type $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(A)$, there is a Galois type $\mathfrak{q} \in S^{<\omega}_{\mathrm{Aut}}(B)$ with $\mathfrak{q}|_f = \mathfrak{p}$.*

Proof. We fix an extension $\mathfrak{C} \geq_{\mathcal{K}} \mathfrak{A}$ and a tuple $\bar{a} \subseteq C$ such that $\mathfrak{p} = \mathrm{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{C}, A)$. By the amalgamation property, we can find an extension $\mathfrak{D} \geq_{\mathcal{K}} \mathfrak{B}$ such that there exists a $\mathcal{K}$-embedding $h : \mathfrak{C} \to \mathfrak{D}$ with $h \restriction A = f$. We can set $\mathfrak{q} := \mathrm{tp}_{\mathrm{Aut}}(h(\bar{a})/\mathfrak{D}, B)$. □

Lemma 2.16. *Let $\mathcal{K}$ be an amalgamation class, γ an ordinal, and suppose that $(\mathfrak{A}_\alpha)_{\alpha<\gamma}$ is an increasing chain of structures $\mathfrak{A}_\alpha \in \mathcal{K}$ such that $\mathfrak{A}_{\alpha+1}$ realises every type in $S^{<\omega}_{\mathrm{Aut}}(A_\alpha)$, for all α. Then their union $\bigcup_{\alpha<\gamma} \mathfrak{A}_\alpha$ is* $\mathrm{cf}(\gamma)$*-Galois saturated.*

Proof. Let $\mathfrak{U} \leq_{\mathcal{K}} \bigcup_{\alpha<\gamma} \mathfrak{A}_\alpha$ be a substructure of size $|U| < \mathrm{cf}(\gamma)$ and fix a type $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(U)$. There exists an index $\alpha < \gamma$ with $U \subseteq A_\alpha$. Hence, $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{A}_\alpha$ and, by Lemma 2.15, we can find a type $\mathfrak{q} \in S^{<\omega}_{\mathrm{Aut}}(A_\alpha)$ with $\mathfrak{q}|_U = \mathfrak{p}$. By construction, $\mathfrak{q}$ is realised in $\bigcup_{\alpha<\gamma} \mathfrak{A}_\alpha \geq_{\mathcal{K}} \mathfrak{A}_{\alpha+1}$. Hence, so is $\mathfrak{p}$. □

Proposition 2.17. *Let $\mathcal{K}$ be an amalgamation class and suppose that κ is a regular cardinal. Every structure $\mathfrak{M} \in \mathcal{K}$ has a κ-Galois saturated extension $\mathfrak{M}^+ \geq_{\mathcal{K}} \mathfrak{M}$.*

Proof. We construct an increasing chain $(\mathfrak{A}_\alpha)_{\alpha<\kappa}$ as follows. We start with $\mathfrak{A}_0 := \mathfrak{M}$. For limit ordinals δ, we set $\mathfrak{A}_\delta := \bigcup_{\alpha<\delta} \mathfrak{A}_\alpha$. For the successor step, we use Lemma 2.13 to find an extension $\mathfrak{A}_{\alpha+1} \geq_{\mathcal{K}} \mathfrak{A}_\alpha$ realising all Galois types over $\mathfrak{A}_\alpha$. By Lemma 2.16, it follows that the limit $\mathfrak{M}^+ := \bigcup_{\alpha<\kappa} \mathfrak{A}_\alpha$ is κ-Galois saturated. □

As usual the existence of Galois saturated structures depends on an additional hypothesis like stability.

Theorem 2.18. *Let $\mathcal{K}$ be a Jónsson class and suppose that κ is a regular cardinal with* $\ln(\mathcal{K}) \leq \kappa < \mathrm{hn}(\mathcal{K})$. *If $\mathcal{K}$ is κ-stable then every structure $\mathfrak{M} \in \mathcal{K}$ of size $|M| \leq \kappa$ has a Galois saturated $\mathcal{K}$-extension of size κ.*

Proof. We construct an increasing chain $(\mathfrak{A}_\alpha)_{\alpha<\kappa}$ of structures $\mathfrak{A}_\alpha \in \mathcal{K}$ of size $|A_\alpha| = \kappa$ as follows. Since $\kappa < \mathrm{hn}(\mathcal{K})$ we have $\mathcal{K}_\kappa \neq \emptyset$. Using amalgamation and the joint embedding property, we can find a structure $\mathfrak{A}_0 \in \mathcal{K}$ of size $|A_0| = \kappa$ with $\mathfrak{M} \leq_{\mathcal{K}} \mathfrak{A}_0$. For limit ordinals δ, we set $\mathfrak{A}_\delta := \bigcup_{\alpha<\delta} \mathfrak{A}_\alpha$. Note that $|A_\delta| \leq |\delta| \otimes \kappa = \kappa$. For the successor step, suppose that $\mathfrak{A}_\alpha$ has already been defined. We use Lemma 2.13 to find an extension $\mathfrak{A}_{\alpha+1} \geq_{\mathcal{K}} \mathfrak{A}_\alpha$ of size $|A_{\alpha+1}| = \kappa$ that realises all types over $\mathfrak{A}_\alpha$. By Lemma 2.16, it follows that the limit $\bigcup_{\alpha<\kappa} \mathfrak{A}_\alpha$ is κ-Galois saturated. □

3. *Limits of chains*

We have seen that we can inductively construct Galois saturated structures as limits of chains. In this section we take a close look at such chains. Our aim is Theorem 4.13, which states that, under certain conditions, the union of a chain of Galois saturated structures is again Galois saturated.

Definition 3.1. Let $\mathcal{K}$ be an abstract elementary class and γ an ordinal.

(a) An increasing chain $(\mathfrak{M}_\alpha)_{\alpha<\gamma}$ is a *weak γ-chain* if each $\mathfrak{M}_{\alpha+1}$ realises every Galois type over M_α. In this case we say that $\mathfrak{M} := \bigcup_\alpha \mathfrak{M}_\alpha$ is the *weak γ-limit* of the chain, or that $\mathfrak{M}$ is a *weak γ-limit* over $\mathfrak{M}_0$.

(b) An increasing chain $(\mathfrak{M}_\alpha)_{\alpha<\gamma}$ is a *strong γ-chain* if every $\mathfrak{M}_{\alpha+1}$ is $|M_{\alpha+1}|^+$-universal over $\mathfrak{M}_\alpha$. In this case we say that $\mathfrak{M} := \bigcup_\alpha \mathfrak{M}_\alpha$ is the *strong γ-limit* of the chain, or that $\mathfrak{M}$ is a *strong γ-limit* over $\mathfrak{M}_0$.

The following observation is just a restatement of Lemma 2.16.

Lemma 3.2. *Let $\mathcal{K}$ be an amalgamation class. Every weak γ-limit is* $\mathrm{cf}(\gamma)$-*Galois saturated.*

Lemma 3.3. *Suppose that $\mathcal{K}$ is an amalgamation class and* $\gamma \geq \ln(\mathcal{K})$ *an ordinal. Let $\mathfrak{M}$ be a weak γ-limit over $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$. Then $\mathfrak{M}$ is $|\gamma|^+$-universal over every $\mathcal{K}$-substructure $\mathfrak{A}_0 \leq_{\mathcal{K}} \mathfrak{A}$ of size $|A_0| \leq |\gamma|$.*

Proof. Let $(\mathfrak{M}_\alpha)_{\alpha<\gamma}$ be a weak γ-chain with limit $\mathfrak{M}$ and $\mathfrak{M}_0 = \mathfrak{A}$. This chain satisfies the hypothesis of Lemma 2.10. □

Corollary 3.4. *Suppose that $\mathcal{K}$ is an amalgamation class, let $\kappa \geq \ln(\mathcal{K})$ be a cardinal, and γ an ordinal. Let $(\mathfrak{M}_\alpha)_{\alpha<\kappa\gamma}$ be a weak $\kappa\gamma$-chain with $|\bigcup_{\alpha<\kappa\gamma} \mathfrak{M}_\alpha| \leq \kappa$. Then the subsequence $(\mathfrak{M}_{\kappa\alpha})_{\alpha<\gamma}$ is a strong γ-chain.*

Proof. Let $\alpha < \gamma$. The sequence $(\mathfrak{M}_{\kappa\alpha+\beta})_{\beta<\kappa}$ is a weak κ-chain over $\mathfrak{M}_{\kappa\alpha}$ with limit $\mathfrak{N} := \bigcup_{\beta<\kappa} \mathfrak{M}_{\kappa\alpha+\beta} \leq \mathfrak{M}_{\kappa(\alpha+1)}$. By the preceding lemma, $\mathfrak{N}$ is κ^+-universal over $\mathfrak{M}_{\kappa\alpha}$. Hence, so is its extension $\mathfrak{M}_{\kappa(\alpha+1)} \geq_{\mathcal{K}} \mathfrak{N}$. As $|M_{\kappa(\alpha+1)}| \leq \kappa$, the claim follows. □

Corollary 3.5. *Suppose that $\mathcal{K}$ is an amalgamation class. Let $\mathfrak{A} \in \mathcal{K}$ be a structure of size $\kappa := |A| \geq \ln(\mathcal{K})$ and let $\gamma < \kappa^+$ be an ordinal. If $\mathcal{K}$ is κ-Galois stable, then there exists a strong γ-limit $\mathfrak{M} \in \mathcal{K}$ over $\mathfrak{A}$ of size $|M| = \kappa$.*

Proof. By Corollary 3.4, it is sufficient to construct a weak $\kappa\gamma$-chain $(\mathfrak{M}_\alpha)_{\alpha<\kappa\gamma}$ over $\mathfrak{A}$ such that $|M_\alpha| = \kappa$, for all α. We define such a chain by induction on α starting with $\mathfrak{M}_0 := \mathfrak{A}$. For the inductive step, note that, given $\mathfrak{M}_\alpha$, we can use Lemma 2.13 to find a structure $\mathfrak{M}_{\alpha+1}$ with the desired properties. □

The next lemma implies that, in the definition of a strong γ-chain $(\mathfrak{M}_\alpha)_\alpha$, we could also require universality of $\mathfrak{M}_{\alpha+1}$ over every $\mathcal{K}$-substructure of $\mathfrak{M}_\alpha$.

Lemma 3.6. *Suppose that $\mathcal{K}$ is an amalgamation class and let $\mathfrak{A} \in \mathcal{K}$ be a structure of size $\ln(\mathcal{K}) \leq |A| < \kappa$. If $\mathfrak{M}$ is κ-universal over $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{M}$, then it is also κ-universal over every substructure $\mathfrak{A}_0 \leq_{\mathcal{K}} \mathfrak{A}$.*

Proof. Let $\mathfrak{A}_0 \leq_{\mathcal{K}} \mathfrak{A}$ and consider a $\mathcal{K}$-embedding $f : \mathfrak{A}_0 \to \mathfrak{C}$ with $|C| < \kappa$. By amalgamation, we can find a $\mathcal{K}$-extension $\mathfrak{C}_+ \geq_{\mathcal{K}} \mathfrak{C}$ of size $|C_+| = |C| \oplus |A| < \kappa$ and a $\mathcal{K}$-embedding $f_+ : \mathfrak{A} \to \mathfrak{C}_+$ such that $f_+ \upharpoonright A_0 = f$. As $\mathfrak{M}$ is κ-universal over $\mathfrak{A}$, there exists a $\mathcal{K}$-embedding $h_+ : \mathfrak{C}_+ \to \mathfrak{M}$ with $h_+ \circ f_+ = \mathrm{id}_A$. Setting $h := h_+ \upharpoonright C$ it follows that $h \circ f = h_+ \circ f_+ \upharpoonright A_0 = \mathrm{id}_{A_0}$, as desired. □

Lemma 3.7. *Let $\mathcal{K}$ be an amalgamation class. If a structure $\mathfrak{M} \in \mathcal{K}$ realises all Galois types over $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$, then it also realises all Galois type over every $\mathfrak{U}_0 \leq_{\mathcal{K}} \mathfrak{U}$.*

Proof. Let $\mathfrak{U}_0 \leq_{\mathcal{K}} \mathfrak{U}$ and $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(\mathfrak{U}_0)$. By Lemma 2.15, there exists a type $\mathfrak{q} \in S^{<\omega}_{\mathrm{Aut}}(\mathfrak{U})$ with $\mathfrak{q}|_{U_0} = \mathfrak{p}$. By assumption, $\mathfrak{M}$ realises $\mathfrak{q}$. Hence, it also realises $\mathfrak{p}$. □

We conclude this section with a result stating that a strong limit is unique up to isomorphism.

Theorem 3.8. *Let $\mathcal{K}$ be an amalgamation class, $\mathfrak{A}, \mathfrak{A}' \in \mathcal{K}$ structures of size $|A|, |A'| \geq \ln(\mathcal{K})$, and let δ, δ' be limit ordinals with $\mathrm{cf}(\delta) = \mathrm{cf}(\delta')$.*

If $\mathfrak{M}$ is a strong δ-limit over $\mathfrak{A}$ and $\mathfrak{M}'$ is a strong δ'-limit over $\mathfrak{A}'$ with $|M| = |M'|$, then we can extend every isomorphism $f : \mathfrak{A} \to \mathfrak{A}'$ to an isomorphism $\pi : \mathfrak{M} \to \mathfrak{M}'$.

Proof. Fix strong chains $(\mathfrak{M}_\alpha)_{\alpha<\delta}$ and $(\mathfrak{M}'_\alpha)_{\alpha<\delta'}$ such that

$$\bigcup_{\alpha<\delta} \mathfrak{M}_\alpha = \mathfrak{M}, \quad \bigcup_{\alpha<\delta'} \mathfrak{M}'_\alpha = \mathfrak{M}', \quad \mathfrak{M}_0 = \mathfrak{A}, \quad \mathfrak{M}'_0 = \mathfrak{A}'.$$

Set $\beta := \mathrm{cf}(\delta)$ and let $h : \beta \to \delta$ and $h' : \beta \to \delta'$ be strictly increasing functions with $h(0) = 0$ and $h'(0) = 0$. We can choose h and h' such that, for every $\alpha < \beta$, $h(\alpha + 1)$ and $h'(\alpha + 1)$ are successor ordinals.

Since $|M| = |M'|$ we can find increasing chains $(\mathfrak{N}_\alpha)_{\alpha<\beta}$ and $(\mathfrak{N}'_\alpha)_{\alpha<\beta}$ of $\mathcal{K}$-substructures $\mathfrak{N}_\alpha \leq_{\mathcal{K}} \mathfrak{M}_{h(\alpha)}$ and $\mathfrak{N}'_\alpha \leq_{\mathcal{K}} \mathfrak{M}'_{h'(\alpha)}$ such that

$$\bigcup_{\alpha<\beta} \mathfrak{N}_\alpha = \mathfrak{M}, \quad \bigcup_{\alpha<\beta} \mathfrak{N}'_\alpha = \mathfrak{M}', \quad \mathfrak{N}_0 = \mathfrak{A}, \quad \mathfrak{N}'_0 = \mathfrak{A}',$$
$$\text{and} \quad |N_\alpha| = |N'_\alpha| = \min\left\{|M_{h(\alpha)}|, |M'_{h'(\alpha)}|\right\}.$$

We construct an increasing chain $(p_\alpha)_{\alpha<\beta}$ of isomorphisms $p_\alpha : \mathfrak{B}_\alpha \to \mathfrak{B}'_\alpha$ such that

$$\mathfrak{N}_\alpha \leq_{\mathcal{K}} \mathfrak{B}_\alpha \leq_{\mathcal{K}} \mathfrak{M}_{h(\alpha)},$$
$$\mathfrak{N}'_\alpha \leq_{\mathcal{K}} \mathfrak{B}'_\alpha \leq_{\mathcal{K}} \mathfrak{M}'_{h'(\alpha)+1},$$

and $\quad |B_\alpha| = |N_\alpha|$.

Then the limit $\pi := \bigcup_{\alpha<\beta} p_\alpha$ is the desired isomorphism $\pi : \mathfrak{M} \to \mathfrak{M}'$.

We start with $p_0 := f : \mathfrak{A} \to \mathfrak{A}'$. For limit ordinals γ, we set $p_\gamma := \bigcup_{\alpha<\gamma} p_\alpha$. For the successor step, suppose that $p_\alpha : \mathfrak{B}_\alpha \to \mathfrak{B}'_\alpha$ has already been defined. We fix a substructure $\mathfrak{C}' \leq_{\mathcal{K}} \mathfrak{M}'_{h'(\alpha+1)}$ such that

$$N'_{\alpha+1} \cup B'_\alpha \subseteq C' \quad \text{and} \quad |C'| = |N'_{\alpha+1}|.$$

By Lemma 3.6, $\mathfrak{M}_{h(\alpha+1)}$ is $|M_{h(\alpha+1)}|^+$-universal over $\mathfrak{B}_\alpha \leq_{\mathcal{K}} \mathfrak{M}_{h(\alpha)}$. Since $|C'| \leq |M_{h(\alpha+1)}|$, it therefore follows that there is a $\mathcal{K}$-embedding $g : \mathfrak{C}' \to \mathfrak{M}_{h(\alpha+1)}$ such that $g \circ p_\alpha = \mathrm{id}_{B_\alpha}$. Fix a $\mathcal{K}$-substructure $\mathfrak{C} \leq_{\mathcal{K}} \mathfrak{M}_{h(\alpha+1)}$ such that

$$N_{\alpha+1} \cup \operatorname{rng} g \subseteq C \quad \text{and} \quad |C| = |N_{\alpha+1}|.$$

As above, $\mathfrak{M}'_{h'(\alpha+1)+1}$ is $|M'_{h'(\alpha+1)+1}|^+$-universal over $\mathfrak{C}' \leq_{\mathcal{K}} \mathfrak{M}'_{h'(\alpha+1)}$, and we have $|C| \leq |M'_{h'(\alpha+1)+1}|$. Hence, we can find a $\mathcal{K}$-embedding $g' : \mathfrak{C} \to \mathfrak{M}'_{h'(\alpha+1)+1}$ such that $g' \circ g = \mathrm{id}_{C'}$. We take this embedding g' for our isomorphism $p_{\alpha+1} : \mathfrak{B}_{\alpha+1} \to \mathfrak{B}'_{\alpha+1}$. Then

$$\mathfrak{N}_{\alpha+1} \leq_{\mathcal{K}} \mathfrak{B}_{\alpha+1} \leq_{\mathcal{K}} \mathfrak{M}_{h(\alpha+1)}\,,$$
$$\mathfrak{N}'_{\alpha+1} \leq_{\mathcal{K}} \mathfrak{B}'_{\alpha+1} \leq_{\mathcal{K}} \mathfrak{M}'_{h'(\alpha+1)+1}\,,$$

and $\quad |B_{\alpha+1}| = |N_{\alpha+1}|$.

Furthermore, for $a \in B_\alpha$, we have

$$\begin{aligned} p_{\alpha+1}(a) = g'(a) &= g'((g \circ p_\alpha)(a)) \\ &= (g' \circ g)(p_\alpha(a)) = p_\alpha(a)\,. \end{aligned}$$

Hence, $p_\alpha \subseteq p_{\alpha+1}$. ☐

Corollary 3.9. *Suppose that $\mathcal{K}$ is an amalgamation class with $\kappa \geq \ln(\mathcal{K})$, and let $\mathfrak{M}$ be a weak $\kappa\delta$-limit over $\mathfrak{A}$ of size $|M| = \kappa$ where δ is a limit ordinal with $\delta < \kappa^+$. Every strong $\kappa\delta$-limit over $\mathfrak{A}$ is isomorphic to $\mathfrak{M}$.*

Proof. By Corollary 3.4, $\mathfrak{M}$ is a strong δ-limit over $\mathfrak{A}$. Since δ is a limit ordinal we have $\mathrm{cf}(\delta) = \mathrm{cf}(\kappa\delta)$. Consequently, the claim follows from Theorem 3.8. ☐

4. Categoricity and stability

In this section we study the consequences of categoricity and stability for an abstract elementary class. We will see that Ehrenfeucht-Mostowski functors provide an invaluable tool in this context.

Lemma 4.1. *Let $\mathcal{K}$ be a κ-categorical abstract elementary class with the joint embedding property where $\kappa \geq \ln(\mathcal{K})$. The structure $\mathfrak{M} \in \mathcal{K}$ of size κ is $\mathcal{K}$-universal.*

Proof. Let $\mathfrak{A} \in \mathcal{K}$ be of size $|A| \leq \kappa$. By the joint embedding property, we can find $\mathcal{K}$-embeddings $f : \mathfrak{A} \to \mathfrak{N}$ and $g : \mathfrak{M} \to \mathfrak{N}$ into some structure $\mathfrak{N} \in \mathcal{K}$ of size $|N| \leq |M| \oplus |A| \oplus \ln(\mathcal{K}) = \kappa$. Since $\mathcal{K}$ is κ-categorical, there exits an isomorphism $\pi : \mathfrak{N} \to \mathfrak{M}$. It follows that $\pi \circ f$ is a $\mathcal{K}$-embedding $\mathfrak{A} \to \mathfrak{M}$. □

We start by showing that categoricity implies stability. This generalises Theorem E6.3.16.

Lemma 4.2. *Suppose that $\mathcal{K}$ is unbounded and κ-categorical, for $\kappa \geq \ln(\mathcal{K})$, and let $\mathfrak{M} \in \mathcal{K}$ be the structure of size $|M| = \kappa$. For every $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$, $\mathfrak{M}$ realises at most $|U| \oplus \ln(\mathcal{K})$ Galois types over U.*

Proof. By Corollary 1.10, there exists an Ehrenfeucht-Mostowski functor $F = \mathrm{pr}_\Sigma \circ F_+$ for $\mathcal{K}$. Then $|F(\kappa)| = \kappa$ implies $F(\kappa) \cong \mathfrak{M}$. W.l.o.g. we may assume that this isomorphism is the identity. Fix a substructure $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{M}$. There is some $I \subseteq \kappa$ of size $|I| \leq |U|$ such that $\mathfrak{U} \subseteq F(I)$. Every finite tuple $\bar{a} \subseteq M = F_+(\kappa)|_\Sigma$ is of the form $a_l = t_l[\bar{\imath}]$ where t_l is a term of the expansion $F_+(\kappa)$ with parameters $\bar{\imath} \subseteq \kappa$. By enlarging the tuples $\bar{\imath}$ we may assume that these parameters are the same for every a_l. If $a'_l = t_l[\bar{\imath}']$ are elements where $\bar{\imath}$ and $\bar{\imath}'$ have the same order type over I, then we can find a linear order L extending κ and an automorphism π of L that fixes I and maps $\bar{\imath}$ to $\bar{\imath}'$. Hence, $F_+(\pi)$ is an automorphism of $F_+(L)$ fixing U and mapping $\bar{a}$ to $\bar{a}'$. Consequently, $\mathrm{tp}_{\mathrm{Aut}}(\bar{a}/\mathfrak{M}, U) = \mathrm{tp}_{\mathrm{Aut}}(\bar{a}'/\mathfrak{M}, U)$.

It follows that the number of Galois types over $\mathfrak{U}$ realised in $\mathfrak{M}$ is bounded by the number of terms $t(\bar{x})$, times the number of order types of finite tuples $\bar{\imath} \subseteq \kappa$ over I. There are at most $\ln(\mathcal{K})$ such terms and, since κ is well-ordered, at most $|I|$ such order types. □

Theorem 4.3. *An unbounded κ-categorical Jónsson class $\mathcal{K}$ is λ-Galois stable, for every cardinal* $\ln(\mathcal{K}) \leq \lambda < \kappa$.

Proof. For a contradiction, suppose that $\mathcal{K}$ is not λ-Galois stable, for some $\ln(\mathcal{K}) \leq \lambda < \kappa$. Fix a structure $\mathfrak{U} \in \mathcal{K}$ of size $|U| = \lambda$ such that $|S^{<\omega}_{\mathrm{Aut}}(U)| > \lambda$. By Proposition 2.17, we can find a $\mathcal{K}$-extension $\mathfrak{A} \geq_{\mathcal{K}} \mathfrak{U}$ of size $|A| = \lambda^+$ realising λ^+ types from $S^{<\omega}_{\mathrm{Aut}}(U)$.

Let $\mathfrak{M} \in \mathcal{K}$ be a structure of size κ. We have seen in Lemma 4.1 that $\mathfrak{M}$ is κ^+-universal. Hence, there exists a $\mathcal{K}$-embedding $f : \mathfrak{A} \to \mathfrak{M}$. It follows that $\mathfrak{M}$ realises at least λ^+ Galois types over $f[U]$. This contradicts Lemma 4.2. □

Lemma 4.4. *Let $\mathcal{K}$ be an amalgamation class. If $\mathcal{K}$ is κ-categorical for $\kappa > \ln(\mathcal{K})$, then the structure $\mathfrak{M} \in \mathcal{K}$ of size κ is* $\mathrm{cf}(\kappa)$*-Galois saturated.*

Proof. Starting with an arbitrary structure $\mathfrak{A}_0 \in \mathcal{K}_{<\kappa}$ we use Lemma 2.13 to construct a strictly increasing chain $(\mathfrak{A}_\alpha)_{\alpha<\kappa}$ of structures $\mathfrak{A}_\alpha \in \mathcal{K}$ of size $|A_\alpha| < \kappa$ such that $\mathfrak{A}_{\alpha+1}$ realises every Galois type over A_α.

By Lemma 2.16, the union $\mathfrak{A}_\kappa := \bigcup_{\alpha<\kappa} \mathfrak{A}_\alpha$ is $\mathrm{cf}(\kappa)$-Galois saturated. Since $|A_\kappa| = \kappa$ and $\mathcal{K}$ is κ-categorical, we have $\mathfrak{A}_\kappa \cong \mathfrak{M}$. Hence, $\mathfrak{M}$ is $\mathrm{cf}(\kappa)$-Galois saturated. □

Corollary 4.5. *Let $\mathcal{K}$ be an unbounded Jónsson class. If $\mathcal{K}$ is κ-categorical, for $\kappa > \ln(\mathcal{K})$, then $\mathcal{K}$ contains Galois saturated structures of size λ, for every regular cardinal λ with* $\ln(\mathcal{K}) \leq \lambda \leq \kappa$.

Proof. For $\lambda = \kappa$, we have already proved the claim in Lemma 4.4. For $\lambda < \kappa$, it follows from Theorems 4.3 and 2.18. □

Next, we consider an analogue of the notion of an indiscernible sequence for abstract elementary classes. The following result is comparable to Theorem E5.3.13.

Lemma 4.6. *Let $\mathcal{K}$ be an amalgamation class and let $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ be an Ehrenfeucht-Mostowski functor for $\mathcal{K}$ with spine s. Suppose that I is a linear order, $\bar{\imath} \in [I]^{<\omega}$ a finite tuple, and $\sigma : \bar{\imath} \to \bar{\imath}$ a permutation such that*

$$\mathrm{tp}_{\mathrm{Aut}}\big(s_I(\bar{\imath}) \bigm/ F(I), \varnothing\big) \neq \mathrm{tp}_{\mathrm{Aut}}\big(s_I(\sigma(\bar{\imath})) \bigm/ F(I), \varnothing\big) .$$

Then $\mathcal{K}$ is not κ-stable, for any $\kappa \geq \ln(\mathcal{K})$.

Proof. We can write each permutation as a product of transpositions. Hence, suppose that $\sigma = \sigma_0 \circ \cdots \circ \sigma_n$, where each $\sigma_l : \bar{\imath} \to \bar{\imath}$ is a permutation of $\bar{\imath}$ interchanging two consecutive components of $\bar{\imath}$. There is at least one index l such that

$$\mathrm{tp}_{\mathrm{Aut}}\big(s_I(\bar{\imath}) \bigm/ F(I), \varnothing\big) \neq \mathrm{tp}_{\mathrm{Aut}}\big(s_I(\sigma_l(\bar{\imath})) \bigm/ F(I), \varnothing\big) ,$$

since, otherwise, we would have

$$\mathrm{tp}_{\mathrm{Aut}}\big(s_I(\bar{\imath}) \bigm/ F(I), \varnothing\big) = \mathrm{tp}_{\mathrm{Aut}}\big(s_I(\sigma(\bar{\imath})) \bigm/ F(I), \varnothing\big) .$$

Replacing σ by σ_l we may therefore assume that $\bar{\imath} = \bar{k} i j \bar{m}$ and $\sigma(\bar{\imath}) = \bar{k} j i \bar{m}$ where $\bar{k} < i < j < \bar{m}$.

Let J be a linear order of size $|J| > \kappa$ containing a dense subset $J_0 \subseteq J$ of size $|J_0| = \kappa$. Set $\mathfrak{M} := F(J)$ and $\mathfrak{U} := F(J_0)$. Since $|U| = \kappa$, it is sufficient to show that

$$\mathrm{tp}_{\mathrm{Aut}}(s_J(x)/\mathfrak{M}, U) \neq \mathrm{tp}_{\mathrm{Aut}}(s_J(y)/\mathfrak{M}, U) , \quad \text{for all } x \neq y \text{ in } J .$$

Fix elements $x < y$ in J. To prove that the Galois types of $s_J(x)$ and $s_J(y)$ over U are different, we choose indices $w, \bar{u}, \bar{v} \subseteq J_0$ such that $x < w < y$ and the tuples $\bar{u} x y \bar{v}$ and $\bar{k} i j \bar{m}$ have the same order type. It follows that

$$\begin{aligned}\mathrm{tp}_{\mathrm{Aut}}(s_J(x w \bar{u} \bar{v})/\mathfrak{M}, \varnothing) &= \mathrm{tp}_{\mathrm{Aut}}(s_I(i j \bar{k} \bar{m})/F(I), \varnothing) \\ &\neq \mathrm{tp}_{\mathrm{Aut}}(s_I(j i \bar{k} \bar{m})/F(I), \varnothing) \\ &= \mathrm{tp}_{\mathrm{Aut}}(s_J(y w \bar{u} \bar{v})/\mathfrak{M}, \varnothing) .\end{aligned}$$

Since $s_J(w \bar{u} \bar{v}) \subseteq U$ the claim follows. □

We have already seen that κ-categorical classes are stable and, therefore, they contain Galois saturated structures of all regular cardinals below κ. We conclude this section with some results about the existence of Galois saturated structures of *singular* cardinality.

Lemma 4.7. *Let $\mathcal{K}$ be a κ-categorical amalgamation class, let $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ be an Ehrenfeucht-Mostowski functor for $\mathcal{K}$, let $\lambda > \ln(\mathcal{K})$ be a cardinal, and set $\mathcal{C}_\lambda := \{ \mu^+ \mid \mu < \lambda \}$. Then $F(I)$ is λ-Galois saturated, for every $\mathcal{C}_\lambda$-universal linear order I of size $\lambda \leq |I| < \mathrm{cf}(\kappa)$.*

Proof. It is sufficient to show that $F(I)$ is μ^+-Galois saturated, for every $\mu < \lambda$. Since I is $\mathcal{C}_\lambda$-universal there is some embedding $h : \mu^+ \to I$. Set $A := \Downarrow \operatorname{rng} h$, $B := I \setminus A$, and $J := A + \kappa + B$. Then $|F(J)| = \kappa$. Since $\mu^+ < \mathrm{cf}(\kappa)$ it therefore follows by Lemma 4.4 that $F(J)$ is μ^+-Galois saturated.

To show that also $F(I)$ is μ^+-Galois saturated, we consider a substructure $\mathfrak{U} \leq_{\mathcal{K}} F(I)$ of size $|U| = \mu$ and a type $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(U)$. Let $\mathfrak{q} \in S^{<\omega}_{\mathrm{Aut}}(F(h)[U])$ be the type with $\mathfrak{q}|_{F(h)} = \mathfrak{p}$. Then $\mathfrak{q}$ is realised by some tuple $\bar{a} \subseteq F(J)$. Each a_l is denoted by a term $t_l[\bar{\imath}\bar{k}]$ (in the Skolem expansion) with parameters $\bar{\imath} \subseteq I$ and $\bar{k} \subseteq J \setminus I$. By enlarging the tuples of parameters we may assume without loss of generality that the parameters $\bar{\imath}\bar{k}$ are the same for every l. Let $J_0 \subseteq J$ be a set of size $|J_0| = \mu$ such that $F(h)[U] \cup \bar{\imath} \subseteq F(J_0)$. Since μ^+ is regular, there is some $\alpha < \mu^+$ such that $J_0 \cap A \subseteq h[\downarrow\alpha]$. Hence, there is some tuple $\bar{k}' \subseteq \operatorname{rng} h$ such that $\bar{k}$ and $\bar{k}'$ have the same order type over $J_0 \cup \bar{\imath}$. Setting $b_l := t_l[\bar{\imath}\bar{k}']$ it follows that $\mathrm{tp}_{\mathrm{Aut}}(\bar{b}/F(I), U) = \mathfrak{p}$. □

In the following $\lambda^{<\omega}$ denotes the linear order $\langle \lambda^{<\omega}, \leq_{\mathrm{lex}} \rangle$ where $\leq_{\mathrm{lex}}$ is the lexicographic order on $\lambda^{<\omega}$.

Proposition 4.8. *Let $\mathcal{K}$ be an unbounded amalgamation class that is κ-categorical, for some regular cardinal $\kappa > \ln(\mathcal{K})$. If $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ is an Ehrenfeucht-Mostowski functor for $\mathcal{K}$, then*

(a) *$F(\lambda)$ is Galois saturated, for every $\ln(\mathcal{K}) < \lambda \leq \kappa$;*

(b) *$F(\lambda^{<\omega}\alpha)$ is Galois saturated, for every cardinal* $\ln(\mathcal{K}) < \lambda \leq \kappa$ *and every ordinal* $\alpha < \lambda^+$.

Proof. For $\lambda < \kappa$, the claims follow from Lemma 4.7 since the orders $\lambda^{<\omega}\alpha$ and λ are both $\mathcal{C}_\lambda$-universal. For $\lambda = \kappa$, note that $F(\kappa^{<\omega}\alpha) \cong F(\kappa)$ is the only structure in $\mathcal{K}$ of size κ. This structure is Galois saturated by Corollary 4.5. □

We can use structures of the form $F(\lambda^{<\omega}\alpha)$ to build strong δ-chains. We start by proving an universality lemma for the order $\lambda^{<\omega}$.

Lemma 4.9. *Let λ be a cardinal. For every ordinal $\beta < \lambda^+$, there exists an embedding $g : \beta \to \lambda^{<\omega}$.*

Proof. We define g by induction on β. If $\beta \leq \lambda$ then we can set $g(\alpha) := \langle\alpha\rangle$, for all $\alpha < \beta$. For the successor step, suppose that $\beta = \gamma + 1$ and let $g_0 : \gamma \to \lambda^{<\omega}$ be the embedding obtained by inductive hypothesis. We define $g : \beta \to \lambda^{<\omega}$ by

$$g(\alpha) := \begin{cases} \langle 0\rangle \cdot g_0(\alpha) & \text{for } \alpha < \gamma\,, \\ \langle 1\rangle & \text{for } \alpha = \gamma\,. \end{cases}$$

If β is a limit ordinal, we fix an increasing chain $(\gamma_i)_{i<\lambda}$ of ordinals $\lambda \leq \gamma_i < \beta$ with $\sup_i \gamma_i = \beta$. By inductive hypothesis, there are embeddings $g_i : \gamma_i \to \lambda^{<\omega}$. We define $g : \beta \to \lambda^{<\omega}$ by

$$g(\alpha) := \langle i\rangle \cdot g_i(\alpha) \quad \text{where } i \text{ is the least index with } \alpha < \gamma_i\,. \qquad \square$$

Lemma 4.10. *Let $\mathcal{K}$ be a κ-categorical amalgamation class where κ is regular, let* $\ln(\mathcal{K}) \leq \lambda < \kappa$ *be a cardinal, and $\delta < \lambda^+$ a limit ordinal. Suppose that $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ is an Ehrenfeucht-Mostowski functor for $\mathcal{K}$.*

(a) *$(F(\lambda^{<\omega}\alpha))_{\alpha<\delta}$ is a strong δ-chain over $F(\lambda^{<\omega})$.*

(b) *If $\mathfrak{M}$ is a strong δ-limit over $F(\lambda^{<\omega})$ of size $|\mathfrak{M}| = \lambda$, then $\mathfrak{M} \cong F(\lambda^{<\omega}\delta)$.*

Proof. (b) follows immediately by (a) and Theorem 3.8.

(a) We have to show that $F(\lambda^{<\omega}(\alpha+1))$ is λ^+-universal over $F(\lambda^{<\omega}\alpha)$. Let $f : F(\lambda^{<\omega}\alpha) \to \mathfrak{C}$ be a $\mathcal{K}$-embedding with $|C| \leq \lambda$. Since $\mathcal{K}$ is κ-categorical, we know by Lemma 4.4 that $F(\lambda^{<\omega}\kappa)$ is Galois saturated. In particular, $F(\lambda^{<\omega}\kappa)$ is λ^+-universal over $F(\lambda^{<\omega}\alpha)$. Hence, we can find a $\mathcal{K}$-embedding $g : \mathfrak{C} \to F(\lambda^{<\omega}\kappa)$ such that $g \circ f = \mathrm{id}$. There exists a set $I \subseteq \lambda^{<\omega}\kappa$ of size $|I| = \lambda$ such that $\operatorname{rng} g \subseteq F(I)$. Setting $I_0 := I \cap \lambda^{<\omega}\alpha$ and $I_1 := I \setminus \lambda^{<\omega}\alpha$, we obtain a partition $I = I_0 \cup I_1$ with $I_0 < I_1$. Since I_1 is a well-order with $\operatorname{ord}(I_1) < \lambda^+$, we can apply Lemma 4.9 to find an embedding $\sigma_1 : I_1 \to \lambda^{<\omega}$. Using σ_1, we define an embedding $\sigma : I \to \lambda^{<\omega}(\alpha+1)$ by

$$\sigma(i) := \begin{cases} i & \text{if } i \in I_0\,, \\ \lambda^{<\omega}\alpha + \sigma_1(i) & \text{if } i \in I_1\,. \end{cases}$$

Setting $h := F(\sigma) \circ g$ we obtain a $\mathcal{K}$-embedding $h : \mathfrak{C} \to F(\lambda^{<\omega}(\alpha+1))$ with

$$h \circ f = F(\sigma) \circ g \circ f = F(\sigma) \circ \mathrm{id}_{F(\lambda^{<\omega}\alpha)} = \mathrm{id}_{F(\lambda^{<\omega}\alpha)}\,. \qquad \square$$

Using these technical results about Ehrenfeucht-Mostowski functors we can prove the following two theorems on the existence of Galois saturated structures.

Theorem 4.11. *Suppose that $\mathcal{K}$ is an unbounded κ-categorical Jónsson class where κ is regular. Let $\mathfrak{A} \in \mathcal{K}$ be a structure of size $|A| = \lambda$ where $\ln(\mathcal{K}) < \lambda < \kappa$, and let $\delta < \lambda^+$ be a limit ordinal. Every strong δ-limit $\mathfrak{M}$ over $\mathfrak{A}$ of size $|M| = \lambda$ is Galois saturated.*

Proof. Let $F : \mathfrak{Lin} \to \mathfrak{Emb}(\mathcal{K})$ be an Ehrenfeucht-Mostowski functor for $\mathcal{K}$ and let $(\mathfrak{M}_\alpha)_{\alpha<\delta}$ be a strong δ-chain over $\mathfrak{A}$ with limit $\mathfrak{M}$. According to Proposition 4.8, the structure $F(\lambda^{<\omega})$ is Galois saturated and has size λ. By Lemma 2.8 (b), $F(\lambda^{<\omega})$ is λ^+-universal. Hence, there exists a $\mathcal{K}$-embedding $f : \mathfrak{A} \to F(\lambda^{<\omega})$. Since $\mathfrak{M}_1$ is λ^+-universal over $\mathfrak{M}_0 = \mathfrak{A}$, there also exists a $\mathcal{K}$-embedding $g : F(\lambda^{<\omega}) \to \mathfrak{M}_1$ with $g \circ f = \mathrm{id}_A$.

Replacing the sequence $(\mathfrak{M}_\alpha)_\alpha$ by isomorphic copies, we may therefore assume that

$$\mathfrak{A} \leq_{\mathcal{K}} F(\lambda^{<\omega}) \leq_{\mathcal{K}} \mathfrak{M}_1 .$$

Since $\mathfrak{M}_2$ is λ^+-universal over $\mathfrak{M}_1$, it is also λ^+-universal over $F(\lambda^{<\omega})$. Let $(\mathfrak{M}'_\alpha)_{\alpha<\delta}$ be the sequence obtained from $(\mathfrak{M}_\alpha)_{\alpha<\delta}$ by replacing the first two entries $\mathfrak{M}_0, \mathfrak{M}_1$ by the single entry $F(\lambda^{<\omega})$. Then $(\mathfrak{M}'_\alpha)_{\alpha<\delta}$ is also a strong δ-chain with limit $\mathfrak{M}$. By Lemma 4.10 (b), we have $\mathfrak{M} \cong F(\lambda^{<\omega}\delta)$. Since $\lambda^{<\omega}\delta$ is $\mathcal{C}_\lambda$-universal, it follows by Lemma 4.7 that $\mathfrak{M}$ is λ-Galois saturated. □

Using the fact that Galois saturated structures of the same cardinality are isomorphic, we obtain the following strengthening of Theorem 3.8.

Corollary 4.12. *Suppose that $\mathcal{K}$ is an unbounded Jónsson class that is κ-categorical, for some regular cardinal κ. Let λ be a cardinal with* $\ln(\mathcal{K}) < \lambda < \kappa$ *and let $\delta, \delta' < \lambda^+$ be limit ordinals. If $\mathfrak{M}, \mathfrak{M}', \mathfrak{A}, \mathfrak{A}' \in \mathcal{K}$ are structures of size λ such that $\mathfrak{M}$ is a strong δ-limit over $\mathfrak{A}$ and $\mathfrak{M}'$ is a strong δ'-limit over $\mathfrak{A}'$, then $\mathfrak{M} \cong \mathfrak{M}'$.*

Our final theorem concerns unions of Galois saturated structures. One can show that we can do without the assumption that λ is a limit cardinal, but the proof is much more involved for regular cardinals λ.

Theorem 4.13. *Let $\mathcal{K}$ be an unbounded κ-categorical Jónsson class where κ is regular, and let λ be a limit cardinal with* $\ln(\mathcal{K}) < \lambda < \kappa$. *If $(\mathfrak{M}_\alpha)_{\alpha<\delta}$ is an increasing chain of Galois saturated structures $\mathfrak{M}_\alpha \in \mathcal{K}$ of size $|M_\alpha| = \lambda$ with $\delta < \lambda^+$, then the union $\bigcup_{\alpha<\delta} \mathfrak{M}_\alpha$ is also Galois saturated.*

Proof. Let $\mathfrak{N} := \bigcup_{\alpha<\delta} \mathfrak{M}_\alpha$ be the limit. Then $|N| \leq |\delta| \otimes \lambda = \lambda$. To show that $\mathfrak{N}$ is Galois saturated fix a structure $\mathfrak{U} \leq_{\mathcal{K}} \mathfrak{N}$ of size $\mu := |U| < \lambda$ and some type $\mathfrak{p} \in S^{<\omega}_{\mathrm{Aut}}(U)$. W.l.o.g. we may assume that $\mu \geq \ln(\mathcal{K})$. Note that λ being a limit implies that $\mu^{++} < \lambda$.

The set $I := \{\, \alpha < \delta \mid (M_{\alpha+1} \smallsetminus M_\alpha) \cap U \neq \emptyset \,\}$ has size $|I| \leq |U| = \mu$. Consequently, there exists a cofinal strictly increasing map $f : \mu_0 \to I$

where $\mu_0 := \mathrm{cf}(\mu) \leq \mu$. We construct a strong μ_0-chain $(\mathfrak{N}_\alpha)_{\alpha<\mu_0}$ where each $\mathfrak{N}_\alpha \leq_{\mathcal{K}} \mathfrak{M}_{f(\alpha)}$ has size $|N_\alpha| = \mu^+$ and, for all $\alpha < \mu_0$, we have

$$U \cap M_{f(\alpha+1)} \subseteq N_{\alpha+1} \subseteq M_{f(\alpha+1)} .$$

We define $\mathfrak{N}_\alpha$ by induction on α. We start with an arbitrary structure $\mathfrak{N}_0 \leq_{\mathcal{K}} \mathfrak{M}$ of size $|N_0| = \mu^+$. For limit ordinals γ, we set $\mathfrak{N}_\gamma := \bigcup_{\alpha<\gamma} \mathfrak{N}_\alpha$.

For the successor step, suppose that $\mathfrak{N}_\alpha$ has already been defined. We construct a weak μ^+-chain $(\mathfrak{B}_\beta)_{\beta<\mu^+}$ with $|B_\beta| = \mu^+$ as follows. We start with an arbitrary structure $\mathfrak{B}_0 \leq_{\mathcal{K}} \mathfrak{M}_{f(\alpha+1)}$ of size $|B_0| = \mu^+$ such that $N_\alpha \cup (U \cap M_{f(\alpha+1)}) \subseteq B_0$. Then we use Lemma 2.13 to inductively define $\mathfrak{B}_\beta$, for $0 < \beta < \mu^+$. Since $\mathcal{K}$ is μ^+-Galois stable, we can choose all $\mathfrak{B}_\beta$ of size $|B_\beta| = \mu^+$. Since $\mathfrak{M}_{f(\beta+1)}$ is μ^{++}-Galois saturated, we can further choose $\mathfrak{B}_\beta$ such that $\mathfrak{B}_\beta \leq_{\mathcal{K}} \mathfrak{M}_{f(\beta+1)}$. Let $\mathfrak{N}_{\alpha+1} := \bigcup_{\beta<\mu^+} \mathfrak{B}_\beta$ be the limit. By Lemma 3.3, $\mathfrak{N}_{\alpha+1}$ is μ^{++}-universal over $\mathfrak{B}_0 \geq_{\mathcal{K}} \mathfrak{N}_\alpha$.

We have constructed a strong μ_0-chain $(\mathfrak{N}_\alpha)_{\alpha<\mu_0}$ whose limit $\mathfrak{A} := \bigcup_{\alpha<\mu_0} \mathfrak{N}_\alpha$ has size $|A| = \mu_0 \otimes \mu^+ = \mu^+$. Since $|N_0| = \mu^+$ it follows by Theorem 4.11 that $\mathfrak{A}$ is Galois saturated. Consequently, $\mathfrak{p}$ is realised in $\mathfrak{A} \leq_{\mathcal{K}} \mathfrak{N}$. $\square$

Recommended Literature

Set theory

M. D. Potter, *Sets. An Introduction,* Oxford University Press 1990.

A. Lévy, *Basic Set Theory,* Springer 1979, Dover 2002.

K. Kunen, *Set Theory. An Introduction to Independence Proofs,* North-Holland 1983.

T. J. Jech, *Set Theory,* 3rd ed., Springer 2003.

Algebra and Category Theory

G. M. Bergman, *An Invitation to General Algebra and Universal Constructions,* 2nd ed., Springer 2015.

P. M. Cohn, *Universal Algebra,* 2nd ed., Springer 1981.

P. M. Cohn, *Basic Algebra,* Springer 2003.

S. Lang, *Algebra,* 3rd ed., Springer 2002.

F. Borceux, *Handbook of Categorical Algebra,* Cambridge University Press 1994.

S. MacLane, *Categories for the Working Mathematician,* 2nd ed., Springer 1998.

J. Adámek, J. Rosický, and M. Vitale, *Algebraic Theories,* Cambridge University Press 2011.

J. Adámek and J. Rosický, *Locally Presentable and Accessible Categories,* Cambridge University Press 1994.

Topology and lattice theory

R. Engelking, *General Topology,* 2nd ed., Heldermann 1989.

C.-A. Faure, A. Frölicher, *Modern Projective Geometry,* Kluwer 2000.

P. T. Johnstone, *Stone Spaces,* Cambridge University Press 1982.

G. Gierz, K. H. Hofmann, K. Keimel, J. D. Lawson, M. Mislove, and D. S. Scott, *Continuous Lattices and Domains,* Cambridge University Press 2003.

Model theory

K. Tent and M. Ziegler, *A Course in Model Theory,* Cambridge University Press 2012.

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

B. Poizat, *A Course in Model Theory,* Springer 2000.

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

General model theory

J. Barwise and S. Feferman, eds., *Model-Theoretic Logics,* Springer 1985.

J. T. Baldwin, *Categoricity,* AMS 2010.

R. Diaconescu, *Institution-Independent Model Theory,* Birkhäuser 2008.

H.-D. Ebbinghaus and J. Flum, *Finite Model Theory,* Springer 1995.

Stability theory

S. Buechler, *Essential Stability Theory,* Springer 1996.

E. Casanovas, *Simple Theories and Hyperimaginaries,* Cambridge University Press 2011.

A. Pillay, *Geometric Stability Theory,* Oxford Science Publications 1996.

F. O. Wagner, *Simple Theories,* Kluwer Academic Publishers 2000.

S. Shelah, *Classification Theory,* 2nd ed., North-Holland 1990.

Symbol Index

Chapter A1

Chapter A2

Chapter A3

Chapter A4

Chapter B1

Chapter B2

Chapter B3

Chapter B4

Chapter B5

Chapter B6

Chapter C1

Chapter C2

Chapter C3

Chapter C4

Chapter C5

Chapter D1

Chapter D2

Chapter D3

Chapter E2

Chapter E3

Chapter E4

Chapter E5

Chapter E6

Chapter E7

Chapter F1

Chapter F2

Chapter F3

Chapter F4

Chapter F5

Chapter G1

Chapter G2

Index

The Roman and Fraktur alphabets

A	*a*	𝔄	𝔞	*N*	*n*	𝔑	𝔫
B	*b*	𝔅	𝔟	*O*	*o*	𝔒	𝔬
C	*c*	ℭ	𝔠	*P*	*p*	𝔓	𝔭
D	*d*	𝔇	𝔡	*Q*	*q*	𝔔	𝔮
E	*e*	𝔈	𝔢	*R*	*r*	ℜ	𝔯
F	*f*	𝔉	𝔣	*S*	*s*	𝔖	ſ 𝔰
G	*g*	𝔊	𝔤	*T*	*t*	𝔗	𝔱
H	*h*	ℌ	𝔥	*U*	*u*	𝔘	𝔲
I	*i*	ℑ	𝔦	*V*	*v*	𝔙	𝔳
J	*j*	𝔍	𝔧	*W*	*w*	𝔚	𝔴
K	*k*	𝔎	𝔨	*X*	*x*	𝔛	𝔵
L	*l*	𝔏	𝔩	*Y*	*y*	𝔜	𝔶
M	*m*	𝔐	𝔪	*Z*	*z*	ℨ	𝔷

The Greek alphabet

Α	*α*	alpha	*Ν*	*ν*	nu
Β	*β*	beta	*Ξ*	*ξ*	xi
Γ	*γ*	gamma	*Ο*	*ο*	omicron
Δ	*δ*	delta	*Π*	*π*	pi
Ε	*ε*	epsilon	*Ρ*	*ρ*	rho
Ζ	*ζ*	zeta	*Σ*	*σ*	sigma
Η	*η*	eta	*Τ*	*τ*	tau
Θ	*ϑ*	theta	*Υ*	*υ*	upsilon
Ι	*ι*	iota	*Φ*	*ϕ*	phi
Κ	*κ*	kappa	*Χ*	*χ*	chi
Λ	*λ*	lambda	*Ψ*	*ψ*	psi
Μ	*μ*	mu	*Ω*	*ω*	omega