INTRODUCTION TO MODEL-THEORETIC STABILITY

L. VAN DEN DRIES

1. Introduction

Stability theory is a set of ideas and techniques in model theory that originated in Morley's proof in the 60's of the Łoś conjecture: if a theory in a countable language is categorical in one uncountable power, then it is categorical in all uncountable powers. Morley discovered that such a theory must be *omega-stable*, which means that its models and their definable sets have some highly restrictive and therefore useful properties. Shelah then took up the ambitious problem of describing explicitly all functions $I(T,-)$ for complete theories T in countable languages, where $I(T,\kappa)$ is the number of models of T of size κ (up to isomorphism) with κ ranging over infinite cardinals. On general grounds (which?) there can only be a limited number of possibilities for such functions. A full solution was achieved by Shelah in a remarkable and extensive body of work, but more important than the precise answer are the intrinsic dividing lines between (complete) theories exposed in this way. Here "intrinsic" means, roughly speaking, "invariant under bi-interpretability". One such dividing line is

$$\textit{stable} \qquad \text{versus} \qquad \textit{unstable},$$

and Shelah was able to associate dimension-like quantities to types and definable sets in models of stable theories.

Later applications in algebra and number theory are closely connected to the new insights concerning definable sets gained in this way.

It may be surprising that a problem on how many models a theory has of a given size can be relevant for the structure of definable sets in a given model. The way this happens is via *types*: elementary extensions of a given model $\mathcal{M}$ are constructed from $\mathcal{M}$ by realizing types over $\mathcal{M}$; in particular, the size of various type spaces over $\mathcal{M}$ determines more or less how many elementary extensions of $\mathcal{M}$ of various kinds there can be. But a type space over $\mathcal{M}$ is just the Stone dual of a boolean algebra of definable sets in $\mathcal{M}$, and the *size* of a type space more or less reflects the *complexity* of the corresponding boolean algebra of definable sets. (I am intentionally vague about "size" and "complexity"; there are various ways to make this precise.) Perhaps the remarkable properties of definable sets in stable structures could have been discovered in another way. As a fact of history, however, they came

Date: November 2005.

to light via the study of type spaces forced by questions on the number of models of a theory.

Throughout, m, n (sometimes decorated with subscripts or accents) range over the set $\mathbb{N} = \{0, 1, 2, \dots\}$ of natural numbers. We let κ, κ', etc., range over *infinite* cardinals. For a set X we let $|X|$ be the size (cardinal) of X.

Given sets P, Q and $R \subseteq P \times Q$, we define for $a \in P$,

$$R(a) := \{b \in Q : (a, b) \in R\} \subseteq Q \qquad \text{(the section of } R \text{ above } a\text{)},$$

and we view R as describing the family $\big(R(a)\big)_{a\in P}$ of subsets of Q. (Of course, this notation is only justified when P and Q are clear from the context.) A *partial* map f from the set P to the set Q (notation: $f : P \rightharpoonup Q$) is a map $f : P' \to Q$ with $P' \subseteq P$.

For an equivalence relation E on a set P we let P/E be the quotient set (its elements are the E-classes $E(p)$ with $p \in P$). Abusing language, we call E *finite* if P/E is finite.

2. Boolean algebras with rank

The *Cantor rank*[1] of a definable set X in a model is, roughly speaking, an ordinal that measures to what extent X can be split up in smaller definable sets. It behaves as a kind of *dimension* of X, and

$$\text{Cantor rank} \ = \ \text{Morley rank}$$

if the ambient model is $\aleph_0$-saturated, as we shall see later.

A key point is that the Cantor rank of X can be defined purely in terms of the boolean algebra of definable subsets of X. This suggests introducing such a notion of rank for elements of any boolean algebra.

We also intend this section as a review of the Stone representation of a boolean algebra, which underlies our later use of types.

Let Or be the class of ordinals. For convenience we add two extra elements $-\infty$ and $+\infty$ to Or, and extend the usual linear ordering on Or to a linear ordering on $\mathrm{Or}_\infty := \mathrm{Or} \cup \{-\infty, +\infty\}$ by letting $-\infty < \lambda < +\infty$ for each ordinal λ. Below we let α, β, λ range over ordinals.

Terminology and notations concerning boolean algebras. Recall that a boolean algebra B is a set with distinguished elements 0 and 1, binary operations $\vee$ and $\wedge$ (*join* and *meet*), and a unary operation $-$ (*complement*), such that certain equational laws are satisfied.

Let B be a boolean algebra. For $a, b \in B$ we define

$$a \leq b \ :\Longleftrightarrow \ a = a \wedge b,$$

which makes B into a poset (partially ordered set), and we say that a and b are *disjoint* if $a \wedge b = 0$. An *atom* of B is an $a > 0$ in B such that there is no $b \in B$ with $a > b > 0$. We let $\mathrm{at}(B)$ be the set of atoms of B.

[1]The term *Cantor-Bendixson rank* is more common.

Given any set X, the (boolean) *algebra of subsets of* X has the subsets of X as its elements, with $\emptyset$ and X as its 0 and 1, the binary operations of taking union and intersection as its join and meet operations, and the unary operation of taking the complement relative to X as its complement operation. The atoms of this boolean algebra are the singletons $\{x\}$ with $x \in X$. These atoms generate the (boolean) subalgebra whose elements are the finite and cofinite subsets of X. (A *cofinite* subset of X is the complement in X of a finite subset of X.) Part of the Stone representation theorem says that any boolean algebra embeds into the boolean algebra of subsets of X, for some set X.

We now fix a boolean algebra B, and let a, b, c (with or without subscripts) denote elements of B. By "$b = b_1 + \cdots + b_k$" we mean that $b = b_1 \vee \cdots \vee b_k$ with pairwise disjoint $b_1, \dots, b_k$. We also put $a - b := a \wedge (-b)$.

An *ideal* of B is a set $I \subseteq B$ such that $0 \in I$, and for all a, b,

$$a \leq b \in I \Rightarrow a \in I, \qquad a, b \in I \Rightarrow a \vee b \in I.$$

Let I be an ideal of B. This yields an equivalence relation $=_I$ on B by setting

$$a =_I b \iff a \vee i = b \vee i \text{ for some } i \in I.$$

Let a/I be the equivalence class of a, and make the set $B/I := \{a/I : a \in B\}$ into a boolean algebra by requiring that $a \mapsto a/I : B \to B/I$ is a boolean algebra homomorphism.

Given any set $H \subseteq B$, the ideal (H) of B generated by H (i.e. the smallest ideal of B, under set inclusion, that contains H) is given by

$$(H) := \{x \in B : x \leq h_1 \vee \cdots \vee h_n \text{ for some } h_1, \dots, h_n \in H\}.$$

Thus $\big(\mathrm{at}(B)\big)$ is the ideal generated by the atoms of B, and its elements are the $a_1 + \cdots + a_m$ with atoms $a_1, \dots, a_m$, with $m = 0$ yielding the element $0 \in B$. Each $a \in \big(\mathrm{at}(B)\big)$ has a unique representation as a sum of atoms: if

$$a = a_1 + \cdots + a_m = b_1 + \cdots + b_n$$

with atoms a_i and b_j, then $m = n$ and $\{a_1, \dots, a_m\} = \{b_1, \dots, b_n\}$.

Defining Cantor rank. By transfinite recursion we assign to each ordinal λ an ideal I_λ of B such that $I_\lambda \subseteq I_\mu$ for $\lambda \leq \mu$:

(1) $I_0 := \big(\mathrm{at}(B)\big)$, the ideal generated by the atoms of B;
(2) for $\lambda > 0$, assume inductively that I_α is an ideal of B for all $\alpha < \lambda$, and that $I_\alpha \subseteq I_\beta$ whenever $\alpha \leq \beta < \lambda$; put $I_{<\lambda} := \bigcup_{\alpha<\lambda} I_\alpha$; then I_λ is defined to be the ideal of B that contains $I_{<\lambda}$ and whose image in $B/I_{<\lambda}$ is $\big(\mathrm{at}(B/I_{<\lambda})\big)$, the ideal generated by the atoms of $B/I_{<\lambda}$.

For convenience we put $I_{<0} := \{0\}$, the trivial ideal of B.

Lemma 2.1. *If $b \notin I_\lambda$, then there is an infinite subset of B whose elements are pairwise disjoint, $< b$, and outside $I_{<\lambda}$.*

Proof. We first do the case $\lambda = 0$. Let $b \notin I_0$. If there are infinitely many atoms $< b$, then the set of such atoms has the desired property. If there are only finitely many atoms $< b$, then by subtracting these atoms from b we reduce to the case that there are no atoms $< b$. In this case we take some nonzero $b_0 < b$, then some nonzero $b_1 < b - b_0$, next some nonzero $b_2 < b - b_0 - b_1$, and so on. Then the set of all b_n has the desired property. The general case follows likewise by considering $B/I_{<\lambda}$ instead of B. $\square$

Note that if $I_\lambda = I_{\lambda+1}$, then $I_\lambda = I_\mu$ for all ordinals $\mu \geq \lambda$.

We now assign to each b its *Cantor rank*, an element of Or_∞ and denoted by $\mathrm{CR}(b)$, and by $\mathrm{CR}_B(b)$ if we wish to indicate the dependence on B. We set $\mathrm{CR}(0) := -\infty$; if $b \neq 0$ and $b \in I_\lambda$ for some λ, then we let $\mathrm{CR}(b)$ be the least such λ; if $b \neq 0$ and $b \notin I_\lambda$ for all λ, then $\mathrm{CR}(b) := +\infty$. Thus

$$I_\lambda = \{b : \mathrm{CR}(b) \leq \lambda\}.$$

For example, if X is infinite and B is the algebra of finite and cofinite subsets of X, then the nonempty finite subsets of X have Cantor rank 0, and the cofinite subsets of X have Cantor rank 1. Note that the lemma above has the following reformulation:

$$\begin{aligned}\mathrm{CR}(b) > \lambda \iff &\text{ there is an infinite sequence } a_0, a_1, a_2, \dots \text{ with } a_n < b \text{ and}\\ &\mathrm{CR}(a_n) \geq \lambda \text{ for all } n, \text{ and } a_m \wedge a_n = 0 \text{ for all } m \neq n.\end{aligned}$$

This equivalence is constantly and tacitly used in inductive proofs below.
Our rank function clearly satisfies

(i) $\mathrm{CR}(b) = -\infty \iff b = 0$,
(ii) $a \leq b \implies \mathrm{CR}(a) \leq \mathrm{CR}(b)$,
(iii) $\mathrm{CR}(a \vee b) = \max(\mathrm{CR}(a), \mathrm{CR}(b))$,
(iv) if $\alpha < \mathrm{CR}(b) < +\infty$, then $\alpha = \mathrm{CR}(a)$ for some a.

Property (iv) says that the ordinals that occur as Cantor ranks of elements of B form an initial segment of Or. In particular, these ordinals are $< |B|^+$, the least cardinal $> |B|$.

We say that b is *ranked* if $-\infty < \mathrm{CR}(b) < +\infty$, that is, $\mathrm{CR}(b)$ is an ordinal. We define the *Cantor degree* $\mathrm{CD}(b)$ (or $\mathrm{CD}_B(b)$) of a ranked element b to be the largest $d \geq 1$ such that there are $b_1, \dots, b_d$ of the same Cantor rank as b with $b = b_1 + \cdots + b_d$; in other words, if $\mathrm{CR}(b) = \lambda$, then $\mathrm{CD}(b)$ is the number of atoms of $B/I_{<\lambda}$ that are $\leq b/I_{<\lambda}$.

Lemma 2.2. *Cantor rank and degree are related as follows:*

(v) *If* $-\infty < \mathrm{CR}(a) = \mathrm{CR}(b) < +\infty$ *and* $a \leq b$, *then* $\mathrm{CD}(a) \leq \mathrm{CD}(b)$.
(vi) *If* $-\infty < \mathrm{CR}(a) = \mathrm{CR}(b) < +\infty$, *then* $\mathrm{CD}(a \vee b) \leq \mathrm{CD}(a) + \mathrm{CD}(b)$ *with equality if* $a \wedge b = 0$,
(vii) *If* $\mathrm{CR}(a) < \mathrm{CR}(b) < +\infty$, *then* $\mathrm{CD}(a \vee b) = \mathrm{CD}(b)$.

Call b *Cantor irreducible* if b is ranked and $\mathrm{CD}(b) = 1$; in that case there are no b_1 and b_2 of the same Cantor rank as b such that $b = b_1 + b_2$. Clearly:

(viii) $(\mathrm{CR}(a)=0,\ \mathrm{CD}(a)=1) \Longleftrightarrow a$ is an atom;

(ix) $\mathrm{CR}(b)=0 \Longleftrightarrow b=a_1+\cdots+a_d$ for some $d \geq 1$ and atoms $a_1,\ldots,a_d$;

(x) if b is ranked of Cantor degree d and $b=b_1+\cdots+b_d$ with each b_i of the same Cantor rank as b, then each b_i is Cantor irreducible.

Cantor rank and degree behave as expected under morphisms:

Proposition 2.3. *Let $\phi : B \to C$ be a boolean algebra morphism and $b \in B$.*

(1) *If ϕ is injective, then $\mathrm{CR}(b) \leq \mathrm{CR}(\phi(b))$, and in case of equality with b ranked we have $\mathrm{CD}(b) \leq \mathrm{CD}(\phi(b))$.*

(2) *If ϕ is surjective, then $\mathrm{CR}(b) \geq \mathrm{CR}(\phi(b))$, and in case of equality with b ranked we have $\mathrm{CD}(b) \geq \mathrm{CD}(\phi(b))$.*

Proof. Suppose ϕ is injective. Then an easy induction shows that $\mathrm{CR}(b) > \beta$ implies $\mathrm{CR}(\phi(b)) > \beta$, from which the first part of (1) follows. The second part of (1) is then immediate from the definition of degree. We prove (2) by induction on $\mathrm{CR}(b)$. The case $\mathrm{CR}(b) = -\infty$ is trivial, so let $\mathrm{CR}(b)=\beta$, and assume that the desired result holds for smaller values of $\mathrm{CR}(b)$. Let $c := \phi(b)$, and suppose $c = c_1+\cdots+c_k$ with $k \geq 1$ and $\mathrm{CR}(c_i) \geq \beta$ for $i=1,\ldots,k$. Take $b_1,\ldots,b_k \in B$ with $\phi(b_i)=c_i$. Replacing b_i by $b_i \wedge b$ if necessary, we may assume $b_i \leq b$ for all i. Note that

$$\phi\big(b_i - \bigvee_{j\neq i} b_j\big) = \phi(b_i) - \bigvee_{j\neq i} \phi(b_j) = c_i - \bigvee_{j\neq i} c_j = c_i,$$

so replacing each b_i by $b_i - \bigvee_{j\neq i} b_j$ we may assume $b_1,\ldots,b_k$ are pairwise disjoint. If $\mathrm{CR}(b_i) < \beta$, then by the inductive hypothesis $\mathrm{CR}(c_i) < \beta$, a contradiction. Hence $\mathrm{CR}(b_i) = \beta$ for all i, and thus $k \leq \mathrm{CD}(b)$. This bound on k shows that $\mathrm{CR}(c)$ cannot be larger than β, and also implies the inequality on the degrees in case $\mathrm{CR}(b) = \mathrm{CR}(c)$. □

Let $B|b := \{a | a \leq b\}$, and consider $B|b$ as a boolean algebra in its own right, by restricting $\vee$ and $\wedge$ to $B|b$. Note that $B|b$ has b as largest element, so $B|b$ is not a boolean subalgebra of B when $b \neq 1$, but the map $a \mapsto a \wedge b : B \to B|b$ is a (surjective) morphism of boolean algebras. The following is almost obvious from the definitions.

Lemma 2.4. *Let $a \leq b$. Then $\mathrm{CR}_B(a) = \mathrm{CR}_{B|b}(a)$. If in addition a is ranked, then $\mathrm{CD}_B(a) = \mathrm{CD}_{B|b}(a)$.*

In combination with a previous remark this lemma implies:

If $\alpha < \mathrm{CR}(b) < +\infty$, then there is $a \leq b$ with $\mathrm{CR}(a) = \alpha$.

When is every nonzero element ranked? To answer this question, we first review filters and the Stone representation. A *filter* of B is just the dual of an ideal of B: it is a set $F \subseteq B$ such that $1 \in F$, and for all a, b we have $a \geq b \in F \Rightarrow a \in F$, and $a, b \in F \Rightarrow a \wedge b \in F$. Equivalently, it is a set $F \subseteq B$ such that $-F := \{-b : b \in F\}$ is an ideal of B. Each a determines a filter $F_a := \{b : b \geq a\}$ of B.

A filter F of B is said to be *proper* if $0 \notin F$, that is, $F \neq B$. An *ultrafilter* of B is a proper filter F of B such that for each a, either $a \in F$ or $-a \in F$. Note that F_a is an ultrafilter of B iff a is an atom. Using Zorn it is easy to see that every proper filter of B is contained in a maximal proper filter of B. One also checks easily that the maximal proper filters of B are exactly the ultrafilters of B. Finally, each proper filter F of B is the intersection of the ultrafilters of B that contain F. These facts easily yield the Stone representation, to which we now turn. The Stone space $\mathrm{St}(B)$ of B is the set of ultrafilters of B. To a we assign

$$[a] := \{F \in \mathrm{St}(B) : a \in F\}, \quad \text{a subset of } \mathrm{St}(B),$$

and this assignment satisfies $[0] = \emptyset$, $[1] = \mathrm{St}(B)$, and

$$[a] \cup [b] = [a \vee b], \quad [a] \cap [b] = [a \wedge b], \quad [-a] = \mathrm{St}(B) \setminus [a].$$

Thus $a \mapsto [a]$ is a morphism of B into the boolean algebra of subsets of $\mathrm{St}(B)$. It is in fact an embedding, called the *Stone representation of* B. We use the word "Stone space" because $\mathrm{St}(B)$ is given the so-called Stone topology, which has the sets $[a]$ as basis. It makes $\mathrm{St}(B)$ a compact hausdorff space whose clopen sets are exactly the sets $[a]$. Thus the Stone representation of B maps B isomorphically onto the algebra of clopen subsets of $\mathrm{St}(B)$.

Next a combinatorial notion: an *infinite binary tree in* B is a family (a_j) of nonzero elements of B, where $j = j_1, \dots, j_n$ ranges over the finite sequences of zeros and ones, with $a_\emptyset = 1$ (with $\emptyset$ the empty sequence), and such that for each j as above we have $a_j = a_{j,0} + a_{j,1}$.

This combinatorial notion is related to an algebraic notion: B is said to be *atomless* if $1 \neq 0$ but B has no atom. Note that if (a_j) is an infinite binary tree in B, then the elements a_j generate an atomless subalgebra of B. Conversely, if B has an atomless subalgebra, then B has an infinite binary tree.

Example of an atomless boolean algebra: take the subsets of $[0,1) \subseteq \mathbb{R}$ that are finite unions of intervals $[a,b)$ with $0 \leq a < b \leq 1$; these subsets are the elements of an atomless subalgebra of the algebra of subsets of $[0,1)$.

We can now state an answer to the question above:

Theorem 2.5. *The following four conditions are equivalent:*

(i) *B has no infinite binary tree,*
(ii) *B has no atomless subalgebra,*
(iii) *each non-zero element of B is ranked,*
(iv) $\mathrm{St}(A)$ *is countable for each countable subalgebra A of B.*

Moreover, (i) implies (v), and is equivalent to (v) if B is countable:

(v) $|\mathrm{St}(B)| \leq |B|$.

Towards the proof of this result we first observe:

Lemma 2.6. *Suppose* $\mathrm{CR}(a) = +\infty$. *Then there are a_1, a_2 such that*

$$a = a_1 + a_2, \quad \mathrm{CR}(a_1) = \mathrm{CR}(a_2) = +\infty.$$

Proof. Take an ordinal $\lambda > |B|^+$. From $\mathrm{CR}(a) > \lambda$ it follows that $a = a_1 + a_2$ with $\mathrm{CR}(a_1), \mathrm{CR}(a_2) \geq \lambda$. Then $\mathrm{CR}(a_1) = \mathrm{CR}(a_2) = +\infty$. $\square$

For Cantor irreducible a we put $F(a) := \{b : \mathrm{CR}(a \wedge b) = \mathrm{CR}(a)\}$, hence also $F(a) = \{b : \mathrm{CR}(a - b) < \mathrm{CR}(a)\}$. Clearly $F(a)$ is then an ultrafilter on B, and if a is an atom, then $F(a) = F_a$. For Cantor irreducible a, b we have

$$F(a) = F(b) \iff \mathrm{CR}(a) = \mathrm{CR}(b) = \mathrm{CR}(a \wedge b).$$

We say that an ultrafilter on B is *ranked* if it contains a ranked element. Given a ranked ultrafilter F on B, take an $a \in F$ of minimal Cantor rank, and of minimal Cantor degree of that rank. Then a is Cantor irreducible, and one checks easily that then $F = F(a)$. So $a \mapsto F(a)$ maps the set of Cantor irreducible elements of B onto the set of ranked ultrafilters of B.

Proof of the Theorem. We already observed the equivalence of (i) and (ii). To prove that (i)$\Rightarrow$(iii), assume (i) and suppose B has a nonranked non-zero element. Then $\mathrm{CR}(1) = +\infty$, so by Lemma 2.6 we get an infinite binary tree (a_j) in B such that all a_j have Cantor rank $+\infty$, contradicting (i). To prove (iii)$\Rightarrow$(ii), assume (iii), and let A be a subalgebra of B. We can assume $1 \neq 0$. Since 1 is ranked in B, it is ranked in A by (1) of Proposition 2.3, so $\mathrm{CR}_A(a) = 0$ for some $a \in A$, so A has an atom, and is therefore not atomless. We have now shown that (i), (ii) and (iii) are equivalent.

Next we prove (iii)$\Rightarrow$(v). Assume (iii). Then each ultrafilter of B is ranked, hence of the form $F(a)$ for some irreducible $a \in B$, so (v) holds.

Assume B is countable. To show (v)$\Rightarrow$(i), we prove the contrapositive: Suppose we have an infinite binary tree (a_j) in B. Each infinite sequence $j_1, j_2, j_3, \ldots$ of zeros and ones leads to a subset $\{1, a_{j_1}, a_{j_1,j_2}, a_{j_1,j_2,j_3}, \ldots\}$ of B that is contained in some ultrafilter on B, and different such infinite sequences give rise in this way to necessarily different ultrafilters. Thus B has $2^{\aleph_0}$ many ultrafilters.

The last two arguments also give the equivalence of (i) with (iv) since condition (i) is inherited by subalgebras. This concludes the proof of the theorem.

In the presence of a set of generators. Let our boolean algebra B be equipped with a distinguished subset G that generates B. It turns out that we can recover the notion of (un)stability in this setting in a very natural way. The combinatorial fact behind this is the following:

Lemma 2.7. *Suppose H is an infinite set and $\mathcal{C}$ is a collection of subsets of H such that $|\mathcal{C}| > |H|$. Then for each n there are distinct elements $h_1, \ldots, h_n \in H$ and sets $C_0, \ldots, C_n \in \mathcal{C}$ such that*

$$\{h_1, \ldots, h_n\} \cap C_j = \{h_1, \ldots, h_j\} \quad \text{for } j = 0, \ldots, n,$$

in other words: $\quad h_i \in C_j \iff i \leq j, \qquad$ *for* $i = 1, \ldots, n,\ j = 0, \ldots, n.$

Proof. For given n the property holds for $\mathcal{C}$ iff it holds for the collection $H - \mathcal{C} := \{H \setminus C : C \in \mathcal{C}\}$: if $h_1, \dots, h_n \in H$ and $C_0, \dots, C_n$ witness this property, then $h_n, \dots, h_1$ and $H \setminus C_n, \dots, H \setminus C_0$ witness it as well. We now proceed by induction on n. The case $n = 0$ is obvious. Suppose the desired result holds for a certain n, for all H and $\mathcal{C}$ satisfying the hypothesis of the lemma. Take a set $C \in \mathcal{C}$ and put $\kappa := |H|$. We make the following case distinction:

Case 1. $\mathcal{C} \cap C := \{X \cap C : X \in \mathcal{C}\}$ has size $> \kappa$. Then there is a $c \in C$ such that more than κ sets $Y \in \mathcal{C} \cap C$ do not contain c: otherwise, each $c \in C$ would be outside at most κ many $Y \in \mathcal{C} \cap C$, hence $C \neq Y$ for at most κ many $Y \in \mathcal{C} \cap C$, contradicting the assumption of case 1. Take such a c, and put $\mathcal{D} := \{Y \in \mathcal{C} \cap C : c \notin Y\}$, a collection of subsets of C with $|\mathcal{D}| > \kappa$. Hence by the inductive hypothesis applied to C and $\mathcal{D}$ there are distinct $h_1, \dots, h_n \in C$ and sets $Y_0, \dots, Y_n \in \mathcal{D}$ such that

$$\{h_1, \dots, h_n\} \cap Y_j = \{h_1, \dots, h_j\} \quad \text{for } j = 0, \dots, n.$$

We have $Y_j = C_j \cap C$ with $C_j \in \mathcal{C}$, $j = 0, \dots, n$. Put $h_{n+1} := c$ and $C_{n+1} := C$. Then

$$\{h_1, \dots, h_{n+1}\} \cap C_j = \{h_1, \dots, h_j\} \quad \text{for } j = 0, \dots, n+1.$$

Case 2. $\mathcal{C} \cap C := \{X \cap C : X \in \mathcal{C}\}$ has size $\leq \kappa$. Then $\mathcal{C} \cap (H \setminus C)$ has size $> \kappa$, so $(H - \mathcal{C}) \cap (H \setminus C)$ has size $> \kappa$. Then case 1 applies to $H - \mathcal{C}$ and $H - C$ in place of $\mathcal{C}$ and C, and we are done by the remark at the beginning of the proof. $\square$

This result of Erdös-Makkai is a perfect example how the *uncountable* can reflect combinatorial complexity of a *finite* nature. (Is there a finite version of this result?)

To apply Lemma 2.7, we note that different ultrafilters of B have different intersections with G, because each ultrafilter of B equals $\{b : \phi(b) = 1\}$ for some boolean algebra morphism $\phi : B \to \{0, 1\}$, and each such morphism is uniquely determined by its restriction to G.

Corollary 2.8. *Suppose* $\mathrm{CR}(1) = +\infty$. *Then for each n there are distinct elements $g_1 \dots, g_n$ and ultrafilters $F_0, \dots, F_n$ of B such that*

$$\{g_1, \dots, g_n\} \cap F_j = \{g_1, \dots, g_j\} \textit{ for } j = 0, \dots, n.$$

Proof. By the assumption B has a countable atomless subalgebra. Hence we have a countable subset H of G such that the subalgebra A generated by H has an atomless subalgebra. Thus A has uncountably many ultrafilters. Different ultrafilters of A intersect H in different subsets, and each ultrafilter of A extends to an ultrafilter of B. Now apply the lemma above to H and the collection $\mathcal{C}$ of intersections of ultrafilters of B with H. $\square$

3. Many-sorted Structures

Model theory is traditionally developed for one-sorted structures, but it is now generally recognized that the many-sorted setting is better suited for expressing various basic constructions like adding imaginaries. Also, natural mathematical structures are often many-sorted to begin with. A *many-sorted* (or multi-sorted) structure is a family of sets $(M_s)_{s\in S}$ equipped with relations

$$R \subseteq M_{s_1} \times \cdots \times M_{s_m}, \qquad (s_1,\dots,s_m \in S)$$

and functions

$$f : M_{s_1} \times \cdots \times M_{s_n} \to M_{s_{n+1}}, \qquad (s_1,\dots,s_{n+1} \in S).$$

The elements of the index set S are called *sorts*, and M_s is the underlying set of sort s. For example, an incidence geometry is a two-sorted structure, consisting of a set P whose elements are called *points*, a set Q whose elements are called *lines*, and a relation $R \subseteq P \times Q$ between points and lines. (Depending on the situation this *incidence relation* R is also assumed to satisfy certain axioms such as "for any two distinct points p_1, p_2 there is exactly one line q such that $R(p_1,q)$ and $R(p_2,q)$".)

We shall need a bit of syntax in connection with many-sorted structures even though we wish to deal with issues that transcend our choice of syntax. A many-sorted language L is a triple $(S, L^{\mathrm{r}}, L^{\mathrm{f}})$ consisting of

(1) a set S whose elements are the *sorts of* L,
(2) a set L^{r} whose elements are the *relation symbols of* L,
(3) a set L^{f} whose elements are the *function symbols of* L,

where L^{r} and L^{f} are disjoint, each $R \in L^{\mathrm{r}}$ is equipped with an arity $(s_1,\dots,s_m) \in S^m$, and each $f \in L^{\mathrm{f}}$ is equipped with an arity $(s_1,\dots,s_{n+1}) \in S^{n+1}$. A function symbol of L of arity (s) with $s \in S$ is also called a *constant symbol of* L *of sort* s. The elements of $L^{\mathrm{r}} \cup L^{\mathrm{f}}$ are also called *nonlogical symbols of* L.

Let from now on L be a language, with S as its set of sorts, unless specified otherwise. The *size* of L is the cardinal

$$|L| := \max\{\aleph_0, |S|, |L^{\mathrm{r}} \cup L^{\mathrm{f}}|\},$$

and we say that L is *countable* if $|L| = \aleph_0$.

An L-*structure* is a many-sorted structure

$$\mathcal{M} = \big(M; (R^{\mathcal{M}})_{R\in L^{\mathrm{r}}}, (f^{\mathcal{M}})_{f\in L^{\mathrm{f}}}\big), \qquad \text{where } M = (M_s)_{s\in S},$$

such that for $R \in L^{\mathrm{r}}$ of arity $(s_1,\dots,s_m)$ its *interpretation* $R^{\mathcal{M}}$ *in* $\mathcal{M}$ is a subset of $M_{s_1} \times \cdots \times M_{s_m}$, and for $f \in L^{\mathrm{f}}$ of arity $(s_1,\dots,s_{n+1})$ its *interpretation* $f^{\mathcal{M}}$ *in* $\mathcal{M}$ is a function $M_{s_1} \times \cdots \times M_{s_n} \to M_{s_{n+1}}$. For a constant symbol c of L of sort s the corresponding M_s-valued function $c^{\mathcal{M}}$ is identified with its unique value in M_s, so $c^{\mathcal{M}} \in M_s$.

From now on $\mathcal{M}$ denotes an L-structure $(M;\cdots)$ with $M = (M_s)_{s\in S}$, unless specified otherwise. If $\mathcal{M}$ is understood from the context we often omit the superscript $\mathcal{M}$ in denoting the interpretation in $\mathcal{M}$ of a nonlogical symbol of L. Elements of M_s are also called "elements of $\mathcal{M}$ of sort s". Given any tuple (family) $\vec{s} = (s_i)_{i\in I}$ of sorts in S, we let $M_{\vec{s}}$ denote the corresponding product set:

$$M_{\vec{s}} := \prod_{i\in I} M_{s_i}.$$

Variables and Terms. We assume that we have available infinitely many symbols, called *unsorted variables*; these are given once and for all, independent of the language L we are dealing with. Given any sort s, the *variables of sort* s are the pairs $v^s := (v, s)$ where v is an unsorted variable. This convention is to guarantee that if s and s' are different sorts, then no variable of sort s is a variable of sort s'. We also assume that no variable of any sort is a nonlogical symbol of any language.

A *variable of* L is a variable of sort s for some $s \in S$, and a *multivariable of* L is a tuple $(x_i)_{i\in I}$ of *distinct* variables of L, where "distinct" means that $x_i \neq x_j$ for $i \neq j$. The size of the index set I is called the *size of* x, and the x_i are called the *variables in* x. Often the index set I is finite, say $I = \{1, \dots, n\}$, so that $x = (x_1, \dots, x_n)$. Given a multivariable $x = (x_i)_{i\in I}$ of L, with x_i of sort s_i for $i \in I$, we define the *x-set of* $\mathcal{M}$ to be the product set

$$M_x := M_{\vec{s}} = \prod_i M_{s_i}, \text{ with } \vec{s} = (s_i)_{i\in I},$$

and we think of x as a variable running over M_x.

Multivariables $x = (x_i)_{i\in I}$ and $y = (y_j)_{j\in J}$ of L are said to be *disjoint* if $x_i \neq y_j$ for all $i \in I$ and $j \in J$, and in that case we put $M_{x,y} := M_x \times M_y$. If in addition $I = J$ and x_i and y_i have the same sort for all $i \in I$ (so that $M_x = M_y$), then we call x and y *disjoint and similar*. From now on x and y denote multivariables of L, unless specified otherwise.

Instead of "x has finite size" we also say "x is finite". A *finite* tuple a in $\mathcal{M}$ is a tuple $a \in M_x$ for a finite x.

We define L-terms to be words on the alphabet consisting of the function symbols and variables of L, obtained recursively as follows: each variable of sort $s \in S$ is an L-term (of sort s) when viewed as a word of length 1, and if f is a function symbol of L of arity $(s_1, \dots, s_n, s_{n+1})$ and $t_1, \dots, t_n$ are L-terms of sort $s_1, \dots, s_n$ respectively, then $ft_1 \dots t_n$ is an L-term of sort s_{n+1}. (In practice we often write $f(t_1, \dots, t_n)$ to denote $ft_1 \dots t_n$, and use similar devices to increase readability.) Each variable-free L-term t determines in the usual way an element $t^{\mathcal{M}}$ of $\mathcal{M}$ of the same sort as t, and often we drop the superscript $\mathcal{M}$ in $t^{\mathcal{M}}$ if it is clear from the context that we mean to refer to an element of $\mathcal{M}$.

An (L, x)-term is a pair (t, x) where t is an L-term such that each variable in t is a variable in x. Such an (L, x)-term is also written more suggestively as $t(x)$, and referred to as "the L-term $t(x)$".

We extend L to the language L_M by adding for each $s \in S$ and $a \in M_s$ a new constant symbol $c(a,s)$ of arity s, and we expand $\mathcal{M}$ to an L_M-structure, also denoted by $\mathcal{M}$ for convenience, by setting $c(a,s)^{\mathcal{M}} := a$. When s and $\mathcal{M}$ are clear from context, then we call $c(a,s)$ the *name of* a and denote it also just by a. Given an L_M-term $t(x)$ and a tuple $a \in M_x$ we obtain in this way by substitution a variable-free L_M-term $t(a)$ of the same sort as t, so $t(a)$ defines an element $t(a)^{\mathcal{M}}$ of $\mathcal{M}$ of the same sort as t, and denoted also by $t(a)$ when $\mathcal{M}$ is clear from context. In other words, we obtain a function

$$a \mapsto t(a) \; : \; M_x \to M_s, \qquad (t \text{ of sort } s).$$

Formulas. There is a slightly annoying problem with variables: to define, say, the theory $\mathrm{Th}(\mathcal{M})$ of $\mathcal{M}$ as a *set*, our usual set theory ZFC seems to require the variables to be the elements of a set. But for some purposes we do not want to limit how many variables there are, or even assume that all variables are elements of a single set. The way out is to assume that bound (quantified) variables are always taken from a fixed countable set (for each sort), while free (unquantified) variables are not limited in this way. More precisely, we assume from now on as given a countably infinite set of unsorted variables whose elements are called *unsorted quantifiable variables.* For each sort s the *quantifiable variables of sort* s are the variables v^s where v is an unsorted quantifiable variable.

We now fix once and for all the usual eight logical symbols:

$$\top \qquad \bot \qquad \neg \qquad \wedge \qquad \vee \qquad = \qquad \exists \qquad \forall$$

These are assumed to be distinct from all relation symbols, function symbols and variables of every language.

The *L-formulas* are words on the alphabet consisting of the nonlogical symbols of L, the variables of L, and the eight logical symbols, and are defined recursively in the usual way from the atomic L-formulas with the proviso that every occurrence of $\exists$ and $\forall$ in a formula is followed immediately by a quantifiable variable. The *atomic L-formulas* are the words $= t_1 t_2$ (usually written as $t_1 = t_2$ for readability) where t_1 and t_2 are L-terms of the *same sort*, together with the words $Rt_1 \dots t_m$ where $R \in L^{\mathrm{r}}$ has arity $(s_1, \dots, s_m)$, and $t_1, \dots, t_m$ are L-terms of sorts $s_1, \dots, s_m$, respectively. An (L,x)-formula is a pair ϕ, x where ϕ is an L-formula all whose free variables are in x. Such an (L,x)-formula is also written more suggestively as $\phi(x)$, and referred to as "the L-formula $\phi(x)$". Sometimes we separate the free variables in a formula into two disjoint multivariables: for example, when referring to an L-formula $\phi(x,y)$ we really mean a triple ϕ, x, y consisting of an L-formula ϕ and disjoint multivariables x and y such that each free variable of ϕ is in x or in y.

A formula without free variables is called a *sentence*, so all variables in a sentence are quantifiable variables. Thus the set of L-sentences has size $|L|$.

Truth and Definability. We define in the usual recursive way what it means for an L_M-sentence σ to be *true in* $\mathcal{M}$ (notation: $\mathcal{M} \models \sigma$). The *theory of* $\mathcal{M}$, denoted by $\mathrm{Th}(\mathcal{M})$, is the set of all L-sentences true in $\mathcal{M}$. If we wish to express the dependence on L we call it the L-theory of $\mathcal{M}$ and denote it by $\mathrm{Th}_L(\mathcal{M})$. Given an L_M-formula $\phi(x)$ with $x = (x_i)$ and a tuple $a \in M_x$ we obtain, by substituting for each i the name of a_i for the free occurrences of x_i in ϕ, an L_M-sentence $\phi(a)$. Thus each L_M-formula $\phi(x)$ defines the subset

$$\phi(M_x) := \{a \in M_x : \mathcal{M} \models \phi(a)\}$$

of M_x in $\mathcal{M}$. We also use this notation for a *set* $\Phi(x)$ of L_M-formula $\phi(x)$:

$$\Phi(M_x) := \{a \in M_x : \mathcal{M} \models \phi(a)\text{for all } \phi \in \Phi\} = \bigcap_{\phi \in \Phi} \phi(M_x).$$

The *definable subsets of* M_x *in* $\mathcal{M}$ are the subsets of M_x defined in $\mathcal{M}$ by an L_M-formula $\phi(x)$, and are exactly the elements of a boolean algebra $\mathrm{Def}_x(\mathcal{M})$ of subsets of M_x. The 0-*definable subsets of* M_x *in* $\mathcal{M}$ are exactly the sets defined in $\mathcal{M}$ by an L-formula $\phi(x)$, and are exactly the elements of a boolean algebra $\mathrm{Def}_x(\mathcal{M}|0)$ of subsets of M_x.

A partial map $h : M_x \rightharpoonup M_y$ is said to be *definable in* $\mathcal{M}$ if its graph is a definable subset of $M_x \times M_y$ in $\mathcal{M}$. Also, 0-definability of such a map refers likewise to 0-definability of its graph.

A *definable set in* $\mathcal{M}$ is a pair $X, \vec{s}$ with $\vec{s} \in S^n$ for some n and X a definable subset of $M_{\vec{s}}$ in $\mathcal{M}$. The role of $\vec{s}$ is just to specify the intended ambient set $M_{\vec{s}}$ of X. The reason for this convention should be clear: it may happen that, for example, $M_s = M_{s'}$ where s and s' are different sorts. In that case, a set $X \subseteq M_s$ could be a definable subset of M_s in $\mathcal{M}$, according to the definition above, while not being a definable subset of $M_{s'}$ in $\mathcal{M}$. Usually we refer to a definable set $X, \vec{s}$ in $\mathcal{M}$ just by its first component X, and $\vec{s}$ is clear from the context or left implicit.

For an L_M-formula $\phi(x)$, $\mathcal{M} \models \phi(x)$ means that $\mathcal{M} \models \phi(a)$ for all $a \in M_x$. Thus if $\psi(x)$ is a second L_M-formula, then $\mathcal{M} \models \phi(x) \to \psi(x)$ means $\phi(M_x) \subseteq \psi(M_x)$, and $\mathcal{M} \models \phi(x) \leftrightarrow \psi(x)$ means $\phi(M_x) = \psi(M_x)$.

The *primitives of* $\mathcal{M}$ (that is, the interpretations in $\mathcal{M}$ of the nonlogical symbols of L) are secondary in the kind of model theory we are going to do; their role is just to generate the 0-definable sets in $\mathcal{M}$. (The definable sets in $\mathcal{M}$ can be defined purely in terms of the 0-definable sets in $\mathcal{M}$; how?) To elaborate on this, define a *structure on* $M = (M_s)_{s \in S}$ to be a family $\big(D(\vec{s})\big)_{\vec{s}}$ indexed by the finite sequences $\vec{s} = s_1, \dots, s_n$ in S, such that for each such $\vec{s} = s_1, \dots, s_n$:

(1) $D(\vec{s})$ is a boolean algebra of subsets of $M_{\vec{s}} := M_{s_1} \times \cdots \times M_{s_n}$;
(2) whenever $X \in D(\vec{s})$ and $s' \in S$, then

$$M_{s'} \times X \in D(s', \vec{s}), \qquad X \times M_{s'} \in D(\vec{s}, s');$$

(3) whenever $s' \in S$, then $\{(a, b, a) : a \in M_{s'}, b \in M_{\vec{s}}\} \in D(s', \vec{s}, s')$;

(4) whenever $s' \in S$ and $X \in D(\vec{s}, s')$, then $\pi(X) \in D(\vec{s})$ where $\pi : M_{\vec{s}} \times M_{s'} \to M_{\vec{s}}$ is the obvious projection map.

It is easy to check that the family $\big(\mathrm{Def}_{\vec{s}}(\mathcal{M}|0)\big)$ is a structure on M that contains the primitives of $\mathcal{M}$, and that it is the smallest structure on M containing these primitives.

Parameter sets. A *parameter set in* $\mathcal{M}$ is a family $A = (A_s)$ where $A_s \subseteq M_s$ for each s. (We omit "in $\mathcal{M}$" if $\mathcal{M}$ is clear from the context.) Such A gives rise to a sublanguage L_A of L_M in the obvious way. If $X \subseteq M_x$ is defined by an L_A-formula $\phi(x)$, then we also say that X is *A-definable in $\mathcal{M}$*, or *definable over A in $\mathcal{M}$*. The subsets of M_x that are A-definable in $\mathcal{M}$ are the elements of a boolean subalgebra $\mathrm{Def}_x(\mathcal{M}|A)$ of $\mathrm{Def}_x(\mathcal{M})$. A partial map $h : M_x \rightharpoonup M_y$ is said to be *A-definable in $\mathcal{M}$* or *definable over A in $\mathcal{M}$* if its graph is an A-definable subset of $M_x \times M_y$ in $\mathcal{M}$. Given parameter sets $A = (A_s)$ and $B = (B_s)$ in $\mathcal{M}$ we define

$$\begin{aligned} A \subseteq B \ &:\Longleftrightarrow\ A_s \subseteq B_s \text{ for all } s, \\ AB &:= (A_s \cup B_s) \quad \text{(a parameter set in } \mathcal{M}), \\ A \cap B &:= (A_s \cap B_s) \quad \text{(a parameter set in } \mathcal{M}). \end{aligned}$$

Given a parameter set A as above, an A-tuple (in $\mathcal{M}$) is a pair $a, \vec{s}$ with $\vec{s} = (s_i) \in S^I$ and $a \in M_{\vec{s}}$, such that $a_i \in A_{s_i}$ for all $i \in I$; this A-tuple is also referred to as "the A-tuple $a \in M_{\vec{s}}$", or even "the A-tuple a".

Let $A = (A_s)$ be a parameter set in $\mathcal{M}$. One checks easily that for finite y, a set $Y \subseteq M_y$ is A-definable iff for some finite x disjoint from y, some A-tuple $a \in M_x$ and some 0-definable $Z \subseteq M_{x,y}$ we have $Y = Z(a)$.

A finite tuple $a \in M_x$ is said to be *A-definable in $\mathcal{M}$* (or *definable over A in $\mathcal{M}$*) if $\{a\}$ is an A-definable subset of M_x, and is said to be *A-algebraic in $\mathcal{M}$* (or *algebraic over A in $\mathcal{M}$*) if a belongs to a finite A-definable subset of M_x. (We can omit "in $\mathcal{M}$" if $\mathcal{M}$ is clear from the context.) We let $\mathrm{dcl}(A)$ be the parameter set such that $\mathrm{dcl}(A)_s = \{a \in M_s : a \text{ is definable over } A\}$ for each s, and we define the parameter set $\mathrm{acl}(A)$ likewise, with "algebraic" instead of "definable". We call $\mathrm{dcl}(A)$ (respectively, $\mathrm{acl}(A)$) the *definable closure of A in $\mathcal{M}$* (respectively, the *algebraic closure of A in $\mathcal{M}$*). We say that A is *definably closed in $\mathcal{M}$* (respectively, *algebraically closed in $\mathcal{M}$*) if $\mathrm{dcl}(A) = A$ (respectively, $\mathrm{acl}(A) = A$). It is easy to check that $\mathrm{dcl}(A)$ is definably closed in $\mathcal{M}$, and that $\mathrm{acl}(A)$ is algebraically closed in $\mathcal{M}$.

The *size* of A is the cardinal sum

$$|A| := \sum_s |A_s|,$$

so $|A| < \kappa$ if $|S| < \kappa$ and $|A_s| < \kappa$ for all s. If L, A and x all have size $< \kappa$, then the set of (L_A, x)-formulas has size $< \kappa$. We say that A is *finite* if $|A|$ is finite, *infinite* if $|A|$ is infinite, and *countable* if $|A| \leq \aleph_0$.

For $x = (x_i)_{i \in I}$ where each x_i is of sort s_i, we put

$$A_x := \prod_{i \in I} A_{s_i}.$$

The parameter set A with $A_s = \emptyset$ for all s is also denoted by 0, in accordance with how we have been using terminology like "0-definable". Sometimes we shall use a tuple $a \in M_{\vec{s}}$ with $s = (s_i) \in S^I$ as if it were a parameter set, for example, when notations like $\mathrm{dcl}(a)$ or $\mathrm{acl}(a)$ occur. In such a context a stands for the parameter set (A_s) such that A_s is the set of all a_i with $s_i = s$, for each s. Thus if B is a parameter set and $a \in M_{\vec{s}}$ as above, then aB and Ba denote again a parameter set: $aB = Ba = AB$ with A the parameter set for which a stands. The terminology "a-definable" and "a-algebraic" is used in the same vein.

Exercises and Definitions. Let $a \in M_x$ and $b \in M_y$ be finite tuples. Show that b is a-definable iff there is a 0-definable partial function $h : M_x \rightharpoonup M_y$ with a in its domain such that $h(a) = b$. Show that b is a-algebraic iff there is a 0-definable $Z \subseteq M_x \times M_y$ such that $b \in Z(a)$ and $Z(a)$ is finite.

Below, notions such as "a and b are interdefinable" are understood to be with respect to the ambient structure $\mathcal{M}$, and we add "in $\mathcal{M}$" if we wish to specify $\mathcal{M}$ as the ambient structure. We say that a and b are *interdefinable* if a is b-definable and b is a-definable. Show the equivalence of (i), (ii), (iii):

(i) $\mathrm{dcl}(a) = \mathrm{dcl}(b)$;
(ii) a and b are interdefinable;
(iii) there is an injective 0-definable partial function $h : M_x \rightharpoonup M_y$ with a in its domain such that $h(a) = b$.

Let in addition A be a parameter set. Then b is said to be *a-definable over A* if b is a-definable in the L_A-structure $\mathcal{M}$. Show that b is a-definable over A iff b is Aa-definable. Also, a and b are said to be *interdefinable over A* if a and b are interdefinable in the L_A-structure $\mathcal{M}$. Show that a and b are interdefinable over A iff $\mathrm{dcl}(Aa) = \mathrm{dcl}(Ab)$.

Likewise, b is said to be *a-algebraic over A* if b is a-algebraic in the L_A-structure $\mathcal{M}$. Show that b is a-algebraic over A iff b is Aa-algebraic. Finally, a and b are said to be *interalgebraic over A* if a and b are interalgebraic in the L_A-structure $\mathcal{M}$. Show that a and b are interalgebraic over A iff $\mathrm{acl}(Aa) = \mathrm{acl}(Ab)$.

Types. Let A be a parameter set in $\mathcal{M}$. A *partial x-type over A in $\mathcal{M}$* is a set $\Phi(x)$ of L_A-formulas $\phi(x)$, each finite subset of which is realized in $\mathcal{M}$ by some element of M_x. An *x-type over A in $\mathcal{M}$* is a partial x-type $p(x)$ over A in $\mathcal{M}$ such that for each L_A-formula $\phi(x)$, either $\phi(x) \in p(x)$ or $\neg\phi(x) \in p(x)$. For $b \in M_x$ we let $\mathrm{tp}_x^{\mathcal{M}}(b|A)$ be the x-type over A in $\mathcal{M}$ realized by b (and we leave out the superscript $\mathcal{M}$ or subscript x if $\mathcal{M}$ or x are clear from context). The set of x-types over A in $\mathcal{M}$ is denoted by

$\mathrm{St}_x^{\mathcal{M}}(A)$, or $\mathrm{St}_x(A)$ if $\mathcal{M}$ is clear from context, and we have a bijection

$$p(x) \mapsto \{\phi(M_x) : \phi(x) \in p(x)\} \ : \ \mathrm{St}_x(A) \longrightarrow \mathrm{St}\left(\mathrm{Def}_x(\mathcal{M}|A)\right)$$

of this set of x-types onto the Stone space of the boolean algebra $\mathrm{Def}_x(\mathcal{M}|A)$. Given parameter sets A and B in $\mathcal{M}$ such that $A \subseteq B$, we have the surjective restriction map

$$p \mapsto p\restriction A \ : \ \mathrm{St}_x(B) \longrightarrow \mathrm{St}_x(A)$$

given by $p\restriction A := \{\phi(x) \in p : \phi(x) \text{ is an } L_A\text{-formula}\}$.

We declare $\mathcal{M}$ to be *κ-saturated* if for each parameter set A in $\mathcal{M}$ of size $< \kappa$ and each variable v of L, each v-type over A in $\mathcal{M}$ can be realized in $\mathcal{M}$. The following is proved as in the one-sorted case:

Lemma 3.1. *If $\mathcal{M}$ is κ-saturated, A a parameter set of size $< \kappa$ and x of size $\leq \kappa$, then each partial x-type over A in $\mathcal{M}$ can be realized in $\mathcal{M}$, that is, $\Phi(M_x) \neq \emptyset$ for each partial x-type over A in $\mathcal{M}$.*

Note also that with the assumptions of this lemma, if X is a definable subset of M_x contained in a union $\bigcup_{j\in J} Y_j$ of A-definable subsets Y_j of M_x, then $X \subseteq \bigcup_{j\in J_0} Y_j$ for some finite $J_0 \subseteq J$.

Exercise. With the assumptions of the last lemma, suppose

$$\bigcap_{i\in I} X_i \subseteq \bigcup_{j\in J} Y_j$$

where the X_i and Y_j are A-definable subsets of M_x. Show that there are finite $I_0 \subseteq I$ and $J_0 \subseteq J$ such that $\bigcap_{i\in I_0} X_i \subseteq \bigcup_{j\in J_0} Y_j$.

Comment. The result of the last exercise can sometimes be used to prove that a set $X \subseteq M_x$ with finite x is A-definable, where $\mathcal{M}$ is κ-saturated and $|A| < \kappa$: if one manages to represent X both as an intersection $X = \bigcap_{i\in I} X_i$ and as a union $X = \bigcup_{j\in J} Y_j$ of A-definable subsets X_i and Y_j of M_x, then by this exercise there are finite $I_0 \subseteq I$ and $J_0 \subseteq J$ such that

$$\bigcap_{i\in I_0} X_i \ = \ X \ = \ \bigcup_{j\in J_0} Y_j,$$

so X is A-definable.

Elementary maps. Let $\mathcal{N} = \left((N_s); \cdots\right)$ be a second L-structure, and f a *partial map from $\mathcal{M}$ to $\mathcal{N}$*. The latter means that f is a family (f_s) with $f_s : A_s \to N_s$ and $A_s \subseteq M_s$ for each $s \in S$; the parameter set $A = (A_s)$ is called the *domain* of f. Given an A-tuple $a \in M_{\vec{s}}$, $\vec{s} = (s_i) \in S^I$, we define $fa = f(a) := \left(f_{s_i}(a_i)\right) \in N_{\vec{s}}$. We call f a *partial elementary map from $\mathcal{M}$ to $\mathcal{N}$* if

$$\mathcal{M} \models \phi(a) \iff \mathcal{N} \models \phi(fa)$$

for every finite A-tuple $a \in M_x$ and L-formula $\phi(x)$. Note that then each $f_s : A_s \to N_s$ is injective.

Let f be a partial elementary map from $\mathcal{M}$ to $\mathcal{N}$ as above. If $A_s = M_s$ for all s, then we call f an *elementary embedding from $\mathcal{M}$ into $\mathcal{N}$*. If $A_s = M_s$

and $f_s(M_s) = N_s$ for all s, then we call f an isomorphism from $\mathcal{M}$ onto $\mathcal{N}$, and if in addition $\mathcal{M} = \mathcal{N}$, we call f an *automorphism of* $\mathcal{M}$.

We define $\mathcal{M}$ to be *strongly κ-homogeneous* if any partial elementary map from $\mathcal{M}$ to itself with domain of size $< \kappa$ can be extended to an automorphism of $\mathcal{M}$.

The automorphism group. The automorphisms of $\mathcal{M}$ form a group $\mathrm{Aut}(\mathcal{M})$ under the obvious composition operation. A parameter set A in $\mathcal{M}$ yields the subgroup $\mathrm{Aut}(\mathcal{M}|A)$ of $\mathrm{Aut}(\mathcal{M})$ consisting of the automorphisms f of $\mathcal{M}$ such that $f_s(a) = a$ for each $s \in S$ and $a \in A_s$. Note that if $\mathcal{M}$ is strongly κ-homogeneous, $|A| < \kappa$ and $a, b \in M_x$ have the same type over A, with x of size $< \kappa$, then there is an $f \in \mathrm{Aut}(\mathcal{M}|A)$ such that $fa = b$. For $X \subseteq M_x$ and $f \in \mathrm{Aut}(\mathcal{M})$ we define $f(X) := \{fa : a \in X\} \subseteq M_x$, and we note that if $X \in \mathrm{Def}_x(\mathcal{M}|A)$, then $f(X) \in \mathrm{Def}_x(\mathcal{M}|fA)$, where $fA = \big(f_s(A_s)\big)$. Thus we have a group action

$$(f, X) \mapsto f(X) \ : \ \mathrm{Aut}(\mathcal{M}|A) \times \mathrm{Def}_x(\mathcal{M}|A) \longrightarrow \mathrm{Def}_x(\mathcal{M}|A).$$

For an automorphism f of $\mathcal{M}$ and a map $\gamma : X \to M_y$ with $X \subseteq M_x$ we define the map $f(\gamma) : f(X) \to M_y$ by $f(\gamma)(fa) = f(\gamma(a))$ for $a \in X$. In other words, $f(\mathrm{graph}(\gamma)) = \mathrm{graph}(f(\gamma))$.

Automorphisms also act on types: for $f \in \mathrm{Aut}(\mathcal{M})$ and $p \in \mathrm{St}_x(A)$ we define $f(p) \in \mathrm{St}_x(fA)$ by

$$f(p) := \{\phi(x, fa) : \ \phi(x, y) \text{ an } L\text{-formula, } a \in A_y, \ \phi(x, a) \in p\}.$$

If we identify these types $p \in \mathrm{St}_x(A)$ and $f(p) \in \mathrm{St}_x(fA)$ with the corresponding ultrafilters of the boolean algebras $\mathrm{Def}(\mathcal{M}|A)$ and $\mathrm{Def}(\mathcal{M}|fA)$, this means $f(p) = \{f(X) : X \in p\}$. In particular, we have a group action

$$(f, p) \mapsto f(p) \ : \ \mathrm{Aut}(\mathcal{M}|A) \times \mathrm{St}_x(A) \longrightarrow \mathrm{St}_x(A).$$

Expansions and Reducts. An *extension* of $L = (S, L^{\mathrm{r}}, L^{\mathrm{f}})$ is a language $L' = (S', L'^{\mathrm{r}}, L'^{\mathrm{f}})$ such that $S \subseteq S'$, $L^{\mathrm{r}} \subseteq L'^{\mathrm{r}}$, and $L^{\mathrm{f}} \subseteq L'^{\mathrm{f}}$. In that case, an L'-structure $\mathcal{M}' = \big((M'_s)_{s \in S'}; \cdots\big)$ is said to be an *L'-expansion of* $\mathcal{M}$ (and $\mathcal{M}$ is said to be the *L-reduct of* $\mathcal{M}'$) if $M_s = M_{s'}$ for all $s \in S$, and each nonlogical symbol of L has the same interpretation in $\mathcal{M}$ as in $\mathcal{M}'$. Given such an L'-expansion $\mathcal{M}'$ of $\mathcal{M}$ we sometimes abuse notation by letting a parameter set $A = (A_s)$ in $\mathcal{M}$ denote also the parameter set $(A_s)_{s \in S'}$ in $\mathcal{M}'$, where $A_s = \emptyset$ for $s \in S' \setminus S$.

Exercise. Let L^1 and L^2 be languages with disjoint sets of sorts S^1 and S^2, and let L be the disjoint union of L^1 and L^2, defined in the obvious way, in particular, with $S = S^1 \cup S^2$ as set of sorts. Let $\mathcal{M}^1 = (M^1; \dots)$ be an L^1-structure and $\mathcal{M}^2 = (M^2; \dots)$ an L^2-structure. Let $\mathcal{M} = (M; \dots)$ be the L-structure such that

- $M_s = M^1_s$ for $s \in S^1$, and $M_s = M^2_s$ for $s \in S^2$;
- each nonlogical symbol of L^1 (respectively, L^2) has the same interpretation in $\mathcal{M}$ as in $\mathcal{M}^1$ (respectively, as in $\mathcal{M}^2$).

(So $\mathcal{M}$ is like a disjoint sum of $\mathcal{M}^1$ and $\mathcal{M}^2$.) Let x be a multivariable of L^1 and y a multivariable of L^2 (so x and y are disjoint multivariables of L), and let $Z \subseteq M_{x,y}$. Show: Z is 0-definable in $\mathcal{M}$ iff Z is a finite union of cartesian products $X \times Y$ where $X \subseteq M_x$ is 0-definable in $\mathcal{M}^1$ and $Y \subseteq M_y$ is 0-definable in $\mathcal{M}^2$. (When y is the empty multivariable of L^2, this says that a subset of M_x is 0-definable in $\mathcal{M}$ iff it is 0-definable in $\mathcal{M}^1$.)

4. Imaginaries

Let E be a 0-definable equivalence relation on $M_{\vec{s}}$, where $\vec{s} = (s_1, \dots, s_n) \in S^n$, that is, E is an equivalence relation on $M_{\vec{s}}$, and is 0-definable as a subset of $M_{\vec{s}} \times M_{\vec{s}}$. It is often useful to treat the equivalence classes of E on an equal footing with the elements of $\mathcal{M}$. This can be done as follows. Let $f : M_{\vec{s}} \to M_E$ be a surjective map onto a set M_E such that for all $a, b \in M_{\vec{s}}$,

$$f(a) = f(b) \iff aEb.$$

(For example, take $M_E = M_{\vec{s}}/E$ and $f(a) = E(a)$, the E-equivalence class of $a \in M_{\vec{s}}$.) Extend L by a new sort $s' \notin S$ and a new function symbol f' of arity $(s_1, \dots, s_n, s')$ to give the language L', and expand $\mathcal{M}$ to an L'-structure $\mathcal{M}' = \big((M'_s)_{s \in S'}; \dots\big)$ by setting $M'_s = M_E$ and interpreting f' in $\mathcal{M}'$ as the function f.

For our purposes this is an innocuous way of expanding $\mathcal{M}$, and to explain this, let $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ be disjoint with x_i and y_i quantifiable of sort s_i for $i = 1, \dots, n$, and let $\phi(x, y)$ be an L-formula that defines $E \subseteq M_{\vec{s}} \times M_{\vec{s}}$ in $\mathcal{M}$.

Let $T := \mathrm{Th}(\mathcal{M})$ and let T' be the L'-theory axiomatized by T together with the L'-sentence

$$\forall v \exists x \big(v = f'(x)\big) \ \wedge\ \forall x \forall y \big(\phi(x, y) \longleftrightarrow f'(x) = f'(y)\big),$$

where v is a quantifiable variable of sort s'.

Lemma 4.1. *Let $\mathcal{N}$ be a model of T. Then*

1. *$\mathcal{N}$ has an expansion to a model $\mathcal{N}'$ of T';*
2. *if $\mathcal{N}'$ is an expansion of $\mathcal{N}$ to a model of T', then any isomorphism $\mathcal{M} \to \mathcal{N}$ expands uniquely to an isomorphism $\mathcal{M}' \to \mathcal{N}'$;*
3. *for any L'-formula $\theta(u, v_1, \dots, v_k)$ with u a multivariable of L and each variable v_j of sort s', there is an L-formula $\theta'(u, v^1, \dots, v^k)$ such that $T' \models \theta\big(u, f'(v^1), \dots, f'(v^k)\big) \longleftrightarrow \theta'(u, v^1, \dots, v^k)$ where $v^j = (v^j_1, \dots, v^j_n)$ and v^j_i is of sort s_i for $i = 1, \dots, n$ and $j = 1, \dots k$.*
4. *if $\mathcal{N}'$ is an expansion of $\mathcal{N}$ to a model of T', then any elementary embedding $\mathcal{M} \to \mathcal{N}$ expands uniquely to an elementary embedding $\mathcal{M}' \to \mathcal{N}'$.*

Items (1) and (2) are almost obvious, (3) follows by an easy induction on formulas, and (4) is a consequence of (3). As a special case of (3), each L'-sentence is T'-equivalent to an L-sentence. Thus T' is complete. Another consequence of (3) is that $\mathcal{M}'$ induces no new structure on $\mathcal{M}$: for any parameter set A in $\mathcal{M}$, a set $X \subseteq M_x$ is A-definable in $\mathcal{M}$ iff it is A-definable in $\mathcal{M}'$.

Expanding $\mathcal{M}$ to $\mathcal{M}^{\mathrm{eq}}$. We now carry out the above for all 0-definable equivalence relations simultaneously. For each $\vec{s} = (s_1, \dots, s_n) \in S^n$ and 0-definable equivalence relation E on $M_{\vec{s}}$, we pick an L-formula $\phi = \phi(x, y)$ with quantifiable x, y that defines E in $\mathcal{M}$, and introduce a new sort $s\phi$, put $M_{s\phi} := M_{\vec{s}}/E$, and let $f_E : M_{\vec{s}} \to M_{s\phi}$ be the function that assigns to each $a \in M_{\vec{s}}$ its E-equivalence class $f_E(a) = E(a)$. Let L^{eq} be the language obtained from L by adding, for each E and corresponding ϕ as above, $s\phi$ as a new sort and a function symbol f_ϕ of arity $(s_1, \dots, s_n, s\phi)$. Let $S^{\mathrm{eq}} \supseteq S$ be the set of sorts of L^{eq}. We then expand $\mathcal{M}$ to the L^{eq}-structure

$$\mathcal{M}^{\mathrm{eq}} := \big((M_s)_{s \in S^{\mathrm{eq}}}; \cdots\big)$$

by interpreting each new function symbol f_ϕ in $\mathcal{M}^{\mathrm{eq}}$ as the function f_E defined above. For each E and corresponding ϕ as above and v a quantifiable variable of sort $s\phi$ we call the sentence

$$\forall v \exists x \big(v = f_\phi(x)\big) \ \wedge\ \forall x \forall y \big(\phi(x,y) \longleftrightarrow f_\phi(x) = f_\phi(y)\big)$$

the *defining axiom* of f_ϕ. We let T^{eq} be the L^{eq}-theory axiomatized by T and the defining axioms for the new function symbols. Thus $\mathcal{M}^{\mathrm{eq}} \models T^{\mathrm{eq}}$, but L^{eq} and T^{eq} depend really only on the theory T of $\mathcal{M}$, rather than on the model $\mathcal{M}$ of T. In particular, *every* model $\mathcal{N}$ of T expands likewise to an L^{eq}-structure $\mathcal{N}^{\mathrm{eq}} \models T^{\mathrm{eq}}$.

Lemma 4.2. *Any expansion of $\mathcal{M}$ to a model of T^{eq} is isomorphic to $\mathcal{M}^{\mathrm{eq}}$ via a unique isomorphism that is the identity on $\mathcal{M}$. If $\mathcal{N}$ is a model of T, then any isomorphism $\mathcal{M} \to \mathcal{N}$ expands uniquely to an isomorphism $\mathcal{M}^{\mathrm{eq}} \to \mathcal{N}^{\mathrm{eq}}$.*

In particular, each automorphism f of $\mathcal{M}$ expands uniquely to an automorphism of $\mathcal{M}^{\mathrm{eq}}$, which we also denote by f. In this way the group $\mathrm{Aut}(\mathcal{M})$ gets identified with the group $\mathrm{Aut}(\mathcal{M}^{\mathrm{eq}})$.

Lemma 4.3. *Let $\theta(u, v_1, \dots, v_k)$ be an L^{eq}-formula with u a multivariable of L and each v_j a variable of sort $s\phi_j \in S^{\mathrm{eq}} \setminus S$ corresponding to a 0-definable equivalence relation $E(j)$ on $M_{\vec{s}(j)}$, $\vec{s}(j) = (s_1(j), \dots, s_{n(j)}(j)) \in S^{n(j)}$. Then there is an L-formula $\theta^{\mathrm{eq}}(u, v^1, \dots, v^k)$ such that*

$$T^{\mathrm{eq}} \models \theta\big(u, f_{\phi_1}(v^1), \dots, f_{\phi_k}(v^k)\big) \longleftrightarrow \theta^{\mathrm{eq}}(u, v^1, \dots, v^k)$$

where $v^j = (v^j_1, \dots, v^j_{n(j)})$ and v^j_i is of sort $s_i(j)$ for $j = 1, \dots k$ and $i = 1, \dots, n(j)$.

In particular, T^{eq} is complete. The lemma also yields that if $\mathcal{N}$ is a model of T, then any elementary embedding $\mathcal{M} \to \mathcal{N}$ expands uniquely to an elementary embedding $\mathcal{M}^{\mathrm{eq}} \to \mathcal{N}^{\mathrm{eq}}$. One more consequence of the lemma is that $\mathcal{M}^{\mathrm{eq}}$ induces no new structure on $\mathcal{M}$: for any parameter set A in $\mathcal{M}$, a set $X \subseteq M_x$ is A-definable in $\mathcal{M}$ iff it is A-definable in $\mathcal{M}^{\mathrm{eq}}$.

Elimination of Imaginaries and Coding Definable Sets. First some terminology in connection with an arbitrary map $g : P \to Q$. The *kernel of* g is the equivalence relation E_g on its domain P defined by

$$aE_g b \iff g(a) = g(b).$$

More generally, given an equivalence relation E on its codomain Q, the *pull back of* E *by* g is the equivalence relation $g^{-1}E$ on P defined by

$$a(g^{-1}E)b \iff g(a)Eg(b).$$

We say that T has *elimination of imaginaries* (EI, for short) if each 0-definable equivalence relation E on a set M_x with finite x is the kernel of some 0-definable map $f : M_x \to M_{\vec{s}}$, $\vec{s} = s_1, \dots, s_n$.

While this definition refers to the particular model $\mathcal{M}$ of T, any model of T instead of $\mathcal{M}$ would give the same notion of EI. We say that $\mathcal{M}$ has EI if its theory T has EI. If $\mathcal{M}$ has EI, then the eq-construction is superfluous (for most purposes), by the following lemma.

Lemma 4.4. *Suppose* $\mathcal{M}$ *has* EI. *Let* A *be a parameter set in* $\mathcal{M}$ *and* X *an* A*-definable set in* $\mathcal{M}^{\text{eq}}$. *Then there is an* A*-definable bijection* $\iota : X \to \iota X$ *between* X *and an* A*-definable set* ιX *in* $\mathcal{M}$.

Proof. Since X is given as a subset of a cartesian product $M_{E(1)} \times \cdots \times M_{E(k)}$ where $E(1), \dots, E(k)$ are 0-definable equivalence relations on the sets $M_{\vec{s}(1)}, \dots, M_{\vec{s}(k)}$, respectively, with $\vec{s}(i) \in S^{n(i)}$ for $i = 1, \dots, k$, it is enough to establish the following. Let E be a 0-definable equivalence relations on $M_{\vec{s}}$ with $\vec{s} \in S^n$; then there is a 0-definable bijection $\iota_E : M_E \to Y$ onto a 0-definable set Y in $\mathcal{M}$. Because $\mathcal{M}$ has EI, we have a 0-definable $f : M_{\vec{s}} \to M_x$ with finite x such that the kernel of f is E; setting $Y := f(M_{\vec{s}}) \subseteq M_x$ this yields a 0-definable bijection $f_E(a) \mapsto f(a) : \ M_E \to Y$ as desired. $\square$

Exercise. Show that if $\mathcal{M}$ has EI and A is a parameter set, then $\mathcal{M}$ as an L_A-structure also has EI.

Lemma 4.5. T^{eq} *has* EI.

Proof. Let $E(1), \dots, E(k)$ be 0-definable equivalence relations on the sets $M_{\vec{s}(1)}, \dots, M_{\vec{s}(k)}$, with $\vec{s}(i) \in S^{n(i)}$ for $i = 1, \dots, k$. Let E be an equivalence relation on $M_{\vec{s}} \times M_{E(1)} \times \cdots \times M_{E(k)}$ that is 0-definable in $\mathcal{M}^{\text{eq}}$. It suffices to find an $\vec{s*} \in S^N$ with $N \in \mathbb{N}$, an equivalence relation E^* on $M_{\vec{s*}}$ that is 0-definable in $\mathcal{M}$, and a map

$$h : M_{\vec{s}} \times M_{E(1)} \times \cdots \times M_{E(k)} \to M_{E^*}$$

that is 0-definable in $\mathcal{M}^{\text{eq}}$ and has E as kernel. Put

$$\vec{s*} := (\vec{s}, \vec{s}(1), \dots, \vec{s}(k)) \in S^{n+n(1)+\cdots+n(k)}, \text{ so}$$
$$M_{\vec{s*}} = M_{\vec{s}} \times M_{\vec{s}(1)} \times \cdots \times M_{\vec{s}(k)}.$$

Then the map $g : M_{\vec{s}*} \to M_{\vec{s}} \times M_{E(1)} \times \cdots \times M_{E(k)}$ given by

$$g(a, b_1, \ldots, b_k) := (a, E(1)(b_1), \ldots, E(k)(b_k))$$

is 0-definable in $\mathcal{M}^{\mathrm{eq}}$, and the pull back of the equivalence relation E by g is an equivalence relation E^* on $M_{\vec{s}*}$ that is 0-definable in $\mathcal{M}$. The map h with the domain and codomain indicated above and given by

$$h(a, E(1)(b_1), \ldots, E(k)(b_k)) = E^*(a, b_1, \ldots, b_k), \quad (a, b_1, \ldots, b_k) \in M_{\vec{s}*},$$

is easily seen to have the desired property. □

Given finite x, y, a tuple $a \in M_x$ is said to *code* the set $Y \subseteq M_y$ (in $\mathcal{M}$) if there is a 0-definable $Z \subseteq M_x \times M_y$ such that $Y = Z(a)$ and $Y \neq Z(b)$ for all $b \neq a$ in M_x. The reader should check that then any finite tuple $c \in M_z$ that is interdefinable with a (in $\mathcal{M}$) also codes Y.

Lemma 4.6. *If T has* EI*, then every definable set in $\mathcal{M}$ has a code.*

Proof. Let x and y be finite and $Y = Z(a) \subseteq M_y$ where $Z \subseteq M_x \times M_y$ is 0-definable. The equivalence relation E on M_x defined by

$$aEb \iff Z(a) = Z(b),$$

is 0-definable. Hence, if T has EI, then E is the kernel of a 0-definable $g : M_x \to M_{\vec{s}}$, $\vec{s} = (s_1, \ldots, s_n)$, and then $g(a)$ codes Y. □

The converse of this lemma is also true under some natural assumptions. One of these assumptions is that T has a home sort. By a *home sort of T* we mean a sort $s_0 \in S$ such that for every $s \in S$ there is an n and a surjective 0-definable map $M_{s_0}^n \to M_s$. While this definition refers to the particular model $\mathcal{M}$ of T, any model of T instead of $\mathcal{M}$ would give the same notion of homesort. In many cases, a many-sorted structure $\mathcal{M}$ is either one-sorted, or arises from a one-sorted structure by adding a few carefully chosen quotient sets, and then a homesort comes for free.

Lemma 4.7. *Let T have home sort s_0 with two distinct* 0*-definable elements in M_{s_0}. Then, given any* 0*-definable sets P and Q in $\mathcal{M}$, there is an $\vec{s} \in S^n$ and there are* 0*-definable injective maps $i : P \to M_{\vec{s}}$ and $j : Q \to M_{\vec{s}}$ with disjoint $i(P)$ and $j(Q)$.*

Proof. Let $P \subseteq M_x$ and $Q \subseteq M_y$ be 0-definable with finite x and y. The assumption on T yields 0-definable elements $p_0 \in M_x$, $q_0 \in M_y$, and $0, 1 \in M_{s_0}$ with $0 \neq 1$. Then define $i : P \to M_x \times M_y \times M_{s_0}$ and $j : Q \to M_x \times M_y \times M_{s_0}$ by $i(p) = (p, q_0, 0)$ and $j(q) = (p_0, q, 1)$. □

Lemma 4.8. *Suppose T has a home sort s_0 with two distinct* 0*-definable elements in M_{s_0}. Assume also that $\mathcal{M}$ is $\aleph_0$-saturated and every definable set in $\mathcal{M}$ has a code. Then T has* EI.

Proof. Let x be finite and E a 0-definable equivalence relation on M_x. For each equivalence class $E(a)$ there is a 0-definable $Z \subseteq M_{\vec{s}(a)} \times M_x$ with $\vec{s}(a) \in S^{n(a)}$ such that $E(a) = Z(c)$ for exactly one $c \in M_{\vec{s}(a)}$ (so this c codes $E(a)$). A standard saturation argument shows that as a varies over M_x we can choose such $\vec{s}(a)$ and Z from fixed finite collections, with the code c a piecewise 0-definable function of a. More precisely, there are 0-definable $X_1, \dots, X_k \subseteq M_x$ that cover M_x, tuples $\vec{s}(1) \in S^{n(1)}, \dots, \vec{s}(k) \in S^{n(k)}$, and 0-definable sets

$$Z_1 \subseteq M_{\vec{s}(1)} \times M_x,\ \dots,\ Z_k \subseteq M_{\vec{s}(k)} \times M_x$$

with for each $i \in \{1, \dots, k\}$ a 0-definable function $c_i : X_i \to M_{\vec{s}(i)}$, such that if $a \in X_i$, then $E(a) = Z_i(c_i(a))$ and $E(a) \neq Z_i(c)$ for all $c \neq c_i(a)$ in $M_{\vec{s}(i)}$. To keep notations simple we assume $k = 2$ in what follows. First we increase X_1 to the 0-definable set $X_1' \subseteq M_x$ consisting of the elements of M_x that are E-equivalent to some element of X_1, and extend c_1 to a 0-definable function $c_1' : X_1' \to M_{\vec{s}(1)}$ by $c_1'(a) := c_1(b)$ if $b \in X_1$ and aEb. We then decrease X_2 to $X_2' = M_x \setminus X_1'$ and let $c_2' : X_2' \to M_{\vec{s}(2)}$ be the restriction of c_2 to X_2'.

Up to this point we have not used the home sort assumption. This comes into play in producing from c_1' and c_2' a single 0-definable function

$$g : M_x \to M_{\vec{s}}, \qquad \vec{s} \in S^n,$$

such that E is the kernel of g. The previous lemma gives an $\vec{s} \in S^n$ and 0-definable injective maps $i : M_{\vec{s}(1)} \to M_{\vec{s}}$ and $j : M_{\vec{s}(2)} \to M_{\vec{s}}$ with disjoint images. Now define $g : M_x \to M_{\vec{s}}$ by $g(a) := i\big(c_1'(a)\big)$ if $a \in X_1'$ and $g(a) := j\big(c_2'(a)\big)$ if $a \in X_2'$. Then c has the desired property. $\square$

Exercise. Let x, y be finite, and suppose $a \in M_x$ codes an A-definable set $Y \subseteq M_y$. Show that the tuple a is A-definable.

5. The Monster Model

From now on we fix a complete L-theory T. By the completeness of T, any two models of T can be elementarily embedded into a third model of T. More generally, any commutative diagram of models of T and elementary embeddings between them can be realized inside a single model of T, where the embeddings become inclusions among elementary submodels. (Try to formulate this precisely, and prove it!) For our purpose there is indeed no serious loss of generality to have all action take place in a single "big" model of T, and there are advantages in doing so. For example, we can in this way import Galois-theoretic ideas.

We refer to basic texts in model theory for the fact that for any κ there is a model of T that is κ-saturated and strongly κ-homogeneous.

In what follows we fix a model $\mathbb{M} = \big((\mathbb{M}_s); \cdots\big) \models T$ and a cardinal $\kappa(\mathbb{M}) > |L|$ such that $\mathbb{M}$ is $\kappa(\mathbb{M})$-saturated and strongly $\kappa(\mathbb{M})$-homogeneous. We also call $\mathbb{M}$ the *monster model of* T. In particular, every model of T of size $\leq \kappa(\mathbb{M})$ has an elementary embedding into $\mathbb{M}$.

The following conventions are in force, unless specified otherwise. Notions of "definable" and "algebraic" for sets and tuples will be relative to $\mathbb{M}$, in the same way they were relative to $\mathcal{M}$ in earlier sections. "Small" will mean "of size $< \kappa(\mathbb{M})$". Multivariables are assumed to be small and in L, and A, A', B denote small parameter sets in $\mathbb{M}$. An elementary submodel of $\mathbb{M}$ is completely determined by its underlying family of sets, and this family is also a parameter set in $\mathbb{M}$. We let M, M' and N denote *small* parameter sets underlying an elementary submodel of $\mathbb{M}$, and we denote this elementary submodel also by M, M' and N; usually we signal this by phrasing like "the model M". On the other hand, if we use the more elaborate notation $\mathcal{M}$, then $\mathcal{M}$ can be any structure, not necessarily a small elementary submodel of $\mathbb{M}$, and *in that context* M will denote the underlying family of sets of $\mathcal{M}$, as before. Likewise with $\mathcal{N}$ and N. (We hope this does not confuse the reader.)

Thus a partial x-type over A is a small set of formulas according to these conventions. We shall write "$\models \theta(x)$" to indicate that $\theta(x)$ is an $L_{\mathbb{M}}$-formula and $\mathbb{M} \models \theta(x)$. Likewise, "$\Phi(x) \models \theta(x)$" will mean that $\Phi(x)$ is a small set of $L_{\mathbb{M}}$-formulas $\phi(x)$, and $\theta(x)$ is an $L_{\mathbb{M}}$-formula such that every $a \in \mathbb{M}_x$ realizing $\Phi(x)$ also realizes $\theta(x)$. Note that in that case there is a finite subset $\Phi_0(x)$ of $\Phi(x)$ such that $\Phi_0(x) \models \theta(x)$. (This is how compactness is built into the monster model via saturation.)

Note that $\mathbb{M}^{\mathrm{eq}}$ is also $\kappa(\mathbb{M})$-saturated and strongly $\kappa(\mathbb{M})$-homogeneous, and $\kappa(\mathbb{M}) > |L^{\mathrm{eq}}| = |L|$, so $\mathbb{M}^{\mathrm{eq}}$ can serve as the corresponding monster model of T^{eq}. Likewise, given any A, our $\mathbb{M}$ remains $\kappa(\mathbb{M})$-saturated and strongly $\kappa(\mathbb{M})$-homogeneous as an L_A-structure, so $\mathbb{M}$ can serve as monster model of T_A, the theory of the L_A-structure $\mathbb{M}$.

Relative definability and automorphisms. Given a partial type $\Phi(x)$ over A, its set $\Phi(\mathbb{M}_x)$ of realizations is a *nonempty* subset of $\mathbb{M}_x$. Note that if $p(x)$ is a type over A realized by $a \in \mathbb{M}_x$, then

$$p(\mathbb{M}_x) = \{fa : f \in \mathrm{Aut}(\mathbb{M}|A)\}$$

that is, $p(\mathbb{M}_x)$ is the orbit of a under the action of $\mathrm{Aut}(\mathbb{M}|A)$ on $\mathbb{M}_x$.

Lemma 5.1. *Given A and a finite tuple a, we have:*

(1) *a is A-definable if and only if $fa = a$ for each $f \in \mathrm{Aut}(\mathbb{M}|A)$;*
(2) *a is A-algebraic if and only if $\{fa : f \in \mathrm{Aut}(\mathbb{M}|A)\}$ is finite;*
(3) *a is A-algebraic if and only if $\{fa : f \in \mathrm{Aut}(\mathbb{M}|A)\}$ is small.*

Proof. Let $a \in \mathbb{M}_x$, and $p(x) := \mathrm{tp}(a|A)$. Suppose $fa = a$ for all $f \in \mathrm{Aut}(\mathbb{M}|A)$. Then $p(x)$ has a as its only realization, so with y disjoint from and similar to x, we have $p(x) \cup p(y) \models x = y$, so there is a $\phi(x) \in p(x)$ such that $\phi(x) \wedge \phi(y) \models x = y$, so $\phi(\mathbb{M}_x) = \{a\}$. The other direction of (1) is obvious.

Next, suppose that $\{fa : f \in \mathrm{Aut}(\mathbb{M}|A)\}$ is finite, say of size m. Then

$$p(x) \cup p(y^1) \cup \cdots \cup p(y^m) \models x = y^1 \vee x = y^2 \vee \cdots \vee x = y^m,$$

where $x, y^1, \ldots, y^m$ are pairwise disjoint and similar. Hence there is a $\phi(x) \in p(x)$ such that

$$\phi(x) \wedge \phi(y^1) \wedge \cdots \wedge \phi(y^m) \models x = y^1 \vee \cdots \vee x = y^m,$$

so $\phi(\mathbb{M}_x)$ is finite, A-definable, and $a \in \phi(\mathbb{M}_x)$. The other direction of (2) is obvious. We leave (3) as an exercise. $\square$

By coding, this lemma yields an analogue for definable sets (instead of finite tuples). To see this, suppose $a \in \mathbb{M}_x$ codes the set $Y \subseteq \mathbb{M}_y$, so we have a 0-definable $Z \subseteq \mathbb{M}_x \times \mathbb{M}_y$ such that $Y = Z(a)$ and $Y \neq Z(b)$ for all $b \neq a$ in $\mathbb{M}_x$. For $f \in \mathrm{Aut}(\mathbb{M})$ we have $f(Y) = Z(fa)$ with $f(Y) \neq Z(b)$ for all $b \neq fa$ in $\mathbb{M}_x$, so fa codes $f(Y)$, in particular, $fa = a$ if and only if $f(Y) = Y$. A definable set $Y \subseteq \mathbb{M}_y$ with finite y is said to be A-*algebraic* if there is a finite A-definable equivalence relation E on $\mathbb{M}_y$ such that Y is a union of E-classes.

Corollary 5.2. *Given A and definable $Y \subseteq \mathbb{M}_y$ with finite y, we have:*

(1) *Y is A-definable if and only if $f(Y) = Y$ for each $f \in \mathrm{Aut}(\mathbb{M}|A)$;*
(2) *Y is A-algebraic if and only if $\{f(Y) : f \in \mathrm{Aut}(\mathbb{M}|A)\}$ is finite;*
(3) *Y is A-algebraic if and only if $\{f(Y) : f \in \mathrm{Aut}(\mathbb{M}|A)\}$ is small.*

Proof. Assume without loss of generality (why?) that T has EI. Then (1) follows from part (1) of Lemma 5.1 by using a code of Y. In (2), the "if" direction requires a little argument: Assume $\{f(Y) : f \in \mathrm{Aut}(\mathbb{M}|A)\}$ is finite. Then this finite set generates a finite boolean algebra F of subsets of $\mathbb{M}_y$. The sets in F are definable subsets of $\mathbb{M}_y$, so the equivalence relation E on $\mathbb{M}_y$ whose equivalence classes are the atoms of F is definable. Since each $f \in \mathrm{Aut}(\mathbb{M}|A)\}$ permutes these atoms, we have $f(E) = E$ for such f, so E is A-definable by (1). Now use that Y is a union of atoms of F.

We leave (3) to the reader. $\square$

Corollary 5.3. *Suppose T has* EI *and let $Y \subseteq \mathbb{M}_y$, with finite y. Then Y is A-algebraic if and only if Y is* $\mathrm{acl}(A)$*-definable.*

Proof. Take a code a for Y. Then for all $f \in \mathrm{Aut}(\mathbb{M}|A)$ we have the equivalence $f(Y) = g(Y) \Leftrightarrow f(a) = g(a)$. By our earlier criteria,

$$\begin{aligned} Y \text{ is } A\text{-algebraic} &\iff a \text{ is } A\text{-algebraic} \iff a \text{ is } \mathrm{acl}(A)\text{-definable} \\ &\iff Y \text{ is } \mathrm{acl}(A)\text{-definable}, \end{aligned}$$

which contains the desired equivalence. $\square$

The next characterization of $\mathrm{acl}(A)$ plays no role in this course, but its proof nicely illustrates how to use automorphisms in the monster model:

$$\mathrm{acl}(A) = \bigcap_{M \supseteq A} M.$$

To prove this, let $a \in \mathbb{M}_s$ and $a \notin \mathrm{acl}(A)_s$. It suffices to show that then there is a model $M \supseteq A$ such that $a \notin M_s$. We take any model $N \supseteq A$, and observe that by Lemma 5.1, (3), there is an $f \in \mathrm{Aut}(\mathbb{M}|A)$ such that $fa \notin N$. Then $M := f^{-1}(N)$ has the desired property.

Strong types. For a parameter set C in $\mathbb{M}^{\mathrm{eq}}$, let $\mathrm{dcl}^{\mathrm{eq}}(C)$ and $\mathrm{acl}^{\mathrm{eq}}(C)$ denote its definable and algebraic closure in $\mathbb{M}^{\mathrm{eq}}$. The *strong type* $\mathrm{stp}(a|A)$ of a finite tuple $a \in \mathbb{M}_x$ over A, is the x-type of a over $\mathrm{acl}^{\mathrm{eq}}(A)$ in $\mathbb{M}^{\mathrm{eq}}$:

$$\mathrm{stp}(a|A) := \mathrm{tp}\left(a|\,\mathrm{acl}^{\mathrm{eq}}(A)\right).$$

Corollary 5.4. *Given A, and finite tuples $a, b \in \mathbb{M}_x$, we have:* $\mathrm{stp}(a|A) = \mathrm{stp}(b|A)$ *if and only if $E(a) = E(b)$ for each finite A-definable equivalence relation E on $\mathbb{M}_x$.*

Proof. By the above it suffices to prove the following claim: Let E be a finite A-definable equivalence relation E on $\mathbb{M}_x$, and let $f \in \mathrm{Aut}\left(\mathbb{M}^{\mathrm{eq}}|\,\mathrm{acl}^{\mathrm{eq}}(A)\right)$. Then $E(fa) = E(a)$.

To prove this claim, take a 0-definable $Z \subseteq \mathbb{M}^{\mathrm{eq}}_{\vec{s}} \times \mathbb{M}_x$ ($\vec{s} \in (S^{\mathrm{eq}})^n$) such that $E(a) = Z(c)$ for a unique $c \in \mathbb{M}^{\mathrm{eq}}_{\vec{s}}$. It follows that for this c its orbit $\{gc : g \in \mathrm{Aut}(\mathbb{M}^{\mathrm{eq}}|A)\}$ is finite, so c is algebraic over A in $\mathbb{M}^{\mathrm{eq}}$, hence $fc = c$, and thus $E(fa) = E(a)$. $\square$

Exercises. Let $Y \subseteq \mathbb{M}_y$ with finite y be definable and let c be a finite tuple. Show: c codes Y if and only if for all $f \in \mathrm{Aut}(\mathbb{M})$, $f(Y) = Y \Longleftrightarrow f(c) = c$. (This equivalence is used frequently.)

Let $h : P \to Q$ be 0-definable and surjective, with $P \subseteq \mathbb{M}_x$ and $Q \subseteq \mathbb{M}_y$ and finite x, y, let $Y \subseteq Q$, and let c be a finite tuple. Show that c codes $Y \subseteq \mathbb{M}_y$ if and only if c codes $h^{-1}(Y) \subseteq \mathbb{M}_x$.

Show that if c and d are finite tuples coding the *nonempty* sets $X \subseteq \mathbb{M}_x$ and $Y \subseteq \mathbb{M}_y$, then (c, d) codes $X \times Y \subseteq \mathbb{M}_{x,y}$.

Let $g : X \to \mathbb{M}_y$ and $h : X \to \mathbb{M}_z$ be definable, with $X \subseteq \mathbb{M}_x$ and finite x, y, z. Show: g and h have codes if and only if $(g, h) : X \to \mathbb{M}_{y,z}$ has a code.

Application to EI. The next lemma basically shows that coding definable relations can be reduced to coding *unary* definable functions. (This may be reminiscent of the fact that proving QE reduces to eliminating a single existential quantifier in front of a conjunction of atoms and negated atoms.) In the proof we use the results stated in the exercises above.

Lemma 5.5. *Suppose T has a home sort s_0, every definable subset of $\mathbb{M}_{s_0}$ has a code, and every definable partial function $\mathbb{M}_{s_0} \rightharpoonup \mathbb{M}_s$, $s \in S$, has a code. Then every definable set in $\mathbb{M}$ has a code.*

Proof. Assume as an inductive hypothesis that for a certain $n > 0$ all definable subsets of $\mathbb{M}^n_{s_0}$ have codes. Let $Y \subseteq \mathbb{M}^{1+n}_{s_0}$ be definable. We shall prove that Y has a code. For each $p \in \mathbb{M}_{s_0}$ the section $Y(p) \subseteq \mathbb{M}^n_{s_0}$ has a code. Hence, by saturation, there are tuples $\vec{s}(1) \in S^{m(1)}, \dots, \vec{s}(k) \in S^{m(k)}$ and 0-definable relations

$$Z_1 \subseteq \mathbb{M}_{\vec{s}(1)} \times \mathbb{M}_{s_0}^n, \ \dots, \ Z_k \subseteq \mathbb{M}_{\vec{s}(k)} \times \mathbb{M}_{s_0}^n$$

such that for each $p \in \mathbb{M}_{s_0}$ there is an $i \in \{1, \dots, k\}$ for which there is exactly one $c \in \mathbb{M}_{\vec{s}(i)}$ with $Y(p) = Z_i(c)$. Define $X_1, \dots, X_k \subseteq \mathbb{M}_{s_0}$ by

$$X_i := \{p \in \mathbb{M}_{s_0} : Y(p) = Z_i(c) \text{ for exactly one } c \in \mathbb{M}_{\vec{s}(i)}\},$$

so we have a definable map $\gamma_i : X_i \to \mathbb{M}_{\vec{s}(i)}$ such that if $p \in X_i$, then $Y(p) = Z_i\big(\gamma_i(p)\big)$ and $Y(p) \neq Z_i(c)$ for all $c \in \mathbb{M}_{\vec{s}(i)}$ with $c \neq \gamma_i(p)$. Note that $X_1, \dots, X_k$ cover $\mathbb{M}_{s_0}$. Let now $f \in \mathrm{Aut}(\mathbb{M})$, and identify γ_i with its graph, a subset of $\mathbb{M}_{s_0\vec{s}(i)}$.

Claim. $f(Y) = Y \iff f(X_i) = X_i$ and $f(\gamma_i) = \gamma_i$ for $i = 1, \dots, k$.

To prove this claim, let $q \in M_{s_0}$ and put $p := f^{-1}(q)$, so $q = f(p)$. Then

$$\begin{aligned} q \in f(X_i) \iff p \in X_i &\iff \exists^! c\big(Y(p) = Z_i(c)\big) \\ &\iff \exists^! c\big(f(Y(p)) = Z_i(fc)\big) \\ &\iff \exists^! d\big(f(Y)(q) = Z_i(d)\big) \end{aligned}$$

for $i = 1, \dots, k$, with c and d ranging over $M_{\vec{s}(i)}$. Note also that if $q \in f(X_i)$, then $f(Y)(q) = Z_i(f(\gamma_i)(q))$. The claim now follows easily.

By the assumption of the lemma, the sets $X_1, \dots, X_k \subseteq M_{s_0}$ have codes $a^1, \dots, a^k$. This assumption and one of the exercises above also yield that the maps $\gamma_1, \dots, \gamma_k$ have codes $b^1, \dots, b^k$. By the claim and one of the exercises above it follows that then Y has code $(a^1, \dots, a^k, b^1, \dots, b^k)$.

We have now shown that for every n, every definable subset of $\mathbb{M}_{s_0}^n$ has a code. An earlier exercise then yields the conclusion of the lemma. $\square$

Corollary 5.6. *Suppose T has a home sort s_0 with two distinct 0-definable elements in $\mathbb{M}_{s_0}$, and that every definable partial function $\mathbb{M}_{s_0} \rightharpoonup \mathbb{M}_s$, $s \in S$, has a code. Then T has* EI.

Decomposing a type space. We finish this section with a fact on type spaces. Let u and v be disjoint (small) multivariables. We wish to describe u,v-types in terms of u-types and v-types. In detail: Each u,v-type $p(u,v)$ over A yields a v-type $p_u(v)$ over A, such that if $(a,b) \in \mathbb{M}_u \times \mathbb{M}_v$ realizes $p(u,v)$, then b realizes $p_u(v)$. (This is clear if one thinks of types as orbits.) Choose for each type $q = q(v) \in \mathrm{St}_v(A)$ a realization $b_q \in \mathbb{M}_v$. Then we can define a map

$$\mathrm{St}_{u,v}(A) \to \bigcup_{q \in \mathrm{St}_v(A)} \mathrm{St}_u(Ab_q) \quad \text{(disjoint union)}$$

as follows: given $p = p(u,v) \in \mathrm{St}_{u,v}(A)$, pick a realization (a,b) such that $b = b_q$ for $q = p_u(v)$, and assign to p the element $\mathrm{tp}(a|Ab) \in \mathrm{St}_u(Ab)$.

Lemma 5.7. *The above map*

$$\mathrm{St}_{u,v}(A) \to \bigcup_{q\in \mathrm{St}_v(A)} \mathrm{St}_u(Ab_q)$$

is injective.

This injectivity is easily verified, and is useful later in bounding the size of the type space $\mathrm{St}_{u,v}(A)$ in terms of the size of simpler type spaces.

6. Definable Types

Definability of types is a consequence of stability, as we shall see in the next section. Definable types, however, also occur significantly in unstable settings. For example, all types over the field of real numbers are definable, and this fact has interesting consequences for limit sets of semialgebraic families. Therefore it makes good sense to introduce definable types before discussing stability.

Throughout this section, x and y are *small* multivariables, although there would be no loss of generality in assuming y to be finite.

Let $\delta(x,y)$ be an L-formula. Below we abuse notation by having δ stand for $\delta(x,y)$. By a *δ-instance over B* we mean a formula $\delta(x,b)$ with $b \in B_y$, and by a *δ-formula over B* we mean an L_B-formula $\phi(x)$ that is equivalent to a boolean combination of δ-instances over B. The subsets of $\mathbb{M}_x$ defined by δ-formulas over B form a boolean subalgebra $\mathrm{Def}_\delta(\mathbb{M}|B)$ of $\mathrm{Def}_x(\mathbb{M}|B)$. By a *partial δ-type over B* we mean a partial x-type $\Phi(x)$ over B consisting of δ-formulas over B. By a *δ-type over B* we mean a partial δ-type $p(x)$ over B such that for each δ-formula $\phi(x)$ over B, either $\phi(x) \in p$ or $\neg\phi(x) \in p$. Note that a δ-type $p(x)$ over B is uniquely determined by its subset

$$\{\delta(x,b) \in p : \ b \in B_y\}$$

of δ-instances. Let $\mathrm{St}_\delta(B)$ be the set of δ-types over B. Then we have a bijection

$$p(x) \mapsto \{\phi(\mathbb{M}_x) : \phi(x) \in p(x)\} \ : \ \mathrm{St}_\delta(B) \to \mathrm{St}\left(\mathrm{Def}_\delta(\mathbb{M}|B)\right)$$

of $\mathrm{St}_\delta(B)$ with an actual Stone space. Given $a \in \mathbb{M}_x$ we let $\mathrm{tp}_\delta(a|B)$ be the δ-type over B realized by a.

Let $p(x)$ be a δ-type over B and $\psi(y)$ an L_B-formula. We say that $\psi(y)$ *defines* $p(x)$ if for all $b \in B_y$,

$$\delta(x,b) \in p(x) \iff \models \psi(b).$$

Note that $\psi(y)$ can define at most one δ-type over B. We say that $p(x)$ *is definable* if some L_B-formula $\psi(y)$ defines $p(x)$. The following is immediate from the definitions:

Lemma 6.1. *Let $a \in \mathbb{M}_x$ and let M be a model. Then the set $\delta(a, M_y) \subseteq M_y$ is definable in M if and only if $\mathrm{tp}_\delta(a|M)$ is definable.*

Note that $\delta(a, M_y)$ is defined with a parameter a possibly outside M!

Let $p(x)$ be again a δ-type over B. When $A \subseteq B$ we say that $p(x)$ *is definable over* A if some L_A-formula $\psi(y)$ defines $p(x)$, that is, for all $b \in B_y$,

$$\delta(x, b) \in p(x) \iff \models \psi(b).$$

When B is a model M, then an L_M-formula $\psi(y)$ that defines $p(x)$ is clearly determined up to equivalence in the L_M-structure $\mathbb{M}$ by p and $\delta(x, y)$, and we let $d_p x\delta(x, y)$ denote such a formula $\psi(y)$; one may think of $d_p x$ as "for x realizing p" and should view the x in $d_p x\delta(x, y)$ as a bound multivariable and y as free.

Consider now a type $p(x) \in \mathrm{St}_x(B)$. Given an L-formula $\delta(x, y)$ we define

$$p(x){\restriction}\delta := \{\phi(x) \in p(x) : \phi(x) \text{ is a } \delta\text{-formula over } B\},$$

which is a δ-type over B. Note that $p(x)$ is the union of the $p(x){\restriction}\delta$ with $\delta(x, y)$ ranging over the L-formulas of the indicated form. A *defining scheme* for $p(x)$ is a map assigning to each L-formula $\delta(x, y)$ an L_B-formula $\psi(y)$ that defines $p(x){\restriction}\delta$, that is,

$$\psi(B_y) = \{b \in B_y : \delta(x, b) \in p(x)\}.$$

We say that $p(x)$ is *definable* if it has a defining scheme, in other words, $p(x)$ is definable if and only if $p(x){\restriction}\delta$ is definable for each L-formula $\delta(x, y)$. When B is a model M, then an L_M-formula $\psi(y)$ that defines $p \restriction \delta$ is determined up to equivalence in the L_M-structure $\mathbb{M}$ by p and $\delta(x, y)$, and we let $d_p x\delta(x, y)$ denote such a formula $\psi(y)$.

Lemma 6.2. *Let x be finite. Suppose that each x-type over B is definable and $|B| \geq |L|$. Then*

$$|\mathrm{St}_x(B)| \leq |B|^{|L|}.$$

Proof. Use that each x-types over B has a defining scheme. $\square$

This upperbound $|B|^{|L|}$ is generally much smaller than the trivial upper bound $2^{|B|}$ for $|\mathrm{St}_x(B)|$ when $|B| \geq |L|$.

Let $p(x) \in \mathrm{St}_x(B)$, and let $A \subseteq B$. Then we say that p is *definable over* A if for each L-formula $\delta(x, y)$ there is an L_A-formula $\psi(y)$ such that

$$\delta(x, b) \in p(x) \iff \models \psi(b), \qquad \text{for all } b \in B_y,$$

in other words, for each L-formula $\delta(x, y)$ the δ-type $p(x) \restriction \delta$ is definable over A. The following is immediate from the definitions

Lemma 6.3. *If $p(x) \in \mathrm{St}_x(B)$ is definable over $A \subseteq B$, and $f \in \mathrm{Aut}(\mathbb{M}|A)$ satisfies $f(B) = B$, then $f(p) = p$.*

Extending definable types. In this subsection we fix M and $B \supseteq M$.

Let $\delta(x,y)$ be an L-formula and $p(x)$ a definable δ-type over M. We extend p to a definable δ-type $q = p \upharpoonleft B$ in $\mathrm{St}_\delta(B)$ as follows. Pick an L_M-formula $\psi(y)$ defining p, and declare

$$\delta(x,b) \in q \iff \models \psi(b), \quad \text{for all } b \in B_y.$$

Let us check that this declaration determines indeed a δ-type q over B. Let $b_1,\dots,b_n \in B_y$ be such that

$$\models \psi(b_1) \wedge \cdots \wedge \psi(b_m) \wedge \neg\psi(b_{m+1}) \wedge \cdots \wedge \neg\psi(b_n).$$

It is enough to show that then

$$\models \exists x\big(\delta(x,b_1) \wedge \cdots \wedge \delta(x,b_m) \wedge \neg\delta(x,b_{m+1}) \wedge \cdots \wedge \neg\delta(x,b_n)\big).$$

Take disjoint $y_1,\dots,y_n$ similar to y and disjoint from x, y. Then

$$\begin{aligned}
M &\models \vec{\psi}(\vec{y}) \to \exists x \vec{\delta}(x,\vec{y}), \quad \text{where } \vec{y} = (y_1,\dots,y_n),\\
\vec{\psi}(\vec{y}) &:= \psi(y_1) \wedge \cdots \wedge \psi(y_m) \wedge \neg\psi(y_{m+1}) \wedge \cdots \wedge \neg\psi(y_n),\\
\vec{\delta}(x,\vec{y}) &:= \delta(x,y_1) \wedge \cdots \wedge \delta(x,y_m) \wedge \neg\delta(x,y_{m+1}) \wedge \cdots \wedge \neg\delta(x,y_n).
\end{aligned}$$

Hence $\models \vec{\psi}(\vec{y}) \to \exists x \vec{\delta}(x,\vec{y})$, as desired. Thus the equivalence above does determine a δ-type q over B. Moreover, $q \upharpoonright M = p$, and any L_M-formula defining p also defines q.

The considerations above can easily be adapted to definable types $p \in \mathrm{St}_x(M)$: each such p extends to a definable type $q = p \upharpoonleft B$ in $\mathrm{St}_x(B)$ by requiring that any defining scheme for p is also a defining scheme for q, in other words,

$$q \upharpoonright \delta \;=\; (p \upharpoonright \delta) \upharpoonleft B \qquad \text{for each } L\text{-formula } \delta(x,y).$$

Lemma 6.4. *Let $p \in \mathrm{St}_x(M)$ be definable. Then p has a unique extension to a type $q \in \mathrm{St}_x(B)$ that is definable over M.*

Proof. The type $q := p \upharpoonleft B$ extends p and is definable over M. Suppose $q' \in \mathrm{St}_x(B)$ also extends p and is definable over M. Given an L-formula $\delta(x,y)$, let $\psi(y)$ and $\psi'(y)$ be L_M-formulas that define $q \upharpoonright \delta$ and $q' \upharpoonright \delta$. Then $\psi(M_y) = \psi'(M_y)$, since q and q' have the same restriction p to M. Hence $\psi(\mathbb{M}_y) = \psi'(\mathbb{M}_y)$, so $q \upharpoonright \delta$ and $q' \upharpoonright \delta$ coincide. As δ is arbitrary, we conclude $q = q'$. $\square$

Global types and canonical bases. Let $\mathrm{St}_x(\mathbb{M})$ be the set of x-types over the parameter set $(\mathbb{M}_s)$ in $\mathbb{M}$. (This parameter set is in general not small!) A type in $\mathrm{St}_x(\mathbb{M})$ is also called a *global* x-type, and, depending on the context, is either viewed as a set of formulas, or as the corresponding ultrafilter of the boolean algebra $\mathrm{Def}_x(\mathbb{M})$. A global x-type $\mathbf{p}$ is said to be definable if for each L-formula $\delta(x,y)$ there is an $L_{\mathbb{M}}$-formula $\psi(y)$ such that for all $b \in \mathbb{M}_y$,

$$\delta(x,b) \in \mathbf{p} \iff \models \; \psi(b).$$

If in addition we can always choose such $\psi(y)$ to be an L_A-formula, then $\mathbf{p}$ is said to be definable over A. For definable $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$ and any L-formula $\phi(x,y)$ we let $d_{\mathbf{p}}x\delta(x,y)$ denote an $L_{\mathbb{M}}$-formula $\psi(y)$ as above.

Exercise. Suppose the (global) type $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$ is definable. Show that $\mathbf{p}$ is definable over some A. Show that for any A,

$$\mathbf{p} \text{ is definable over } A \iff f(\mathbf{p}) = \mathbf{p} \text{ for all } f \in \mathrm{Aut}(\mathbb{M}|A).$$

Definition. Let $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$. A *canonical base* for $\mathbf{p}$ is a parameter set A such that for all $f \in \mathrm{Aut}(\mathbb{M})$ we have

$$f(\mathbf{p}) = \mathbf{p} \iff f \in \mathrm{Aut}(\mathbb{M}|A).$$

If A is a canonical base for $\mathbf{p}$, then B is a canonical base for $\mathbf{p}$ if and only if $\mathrm{dcl}(A) = \mathrm{dcl}(B)$. If $\mathbf{p}$ has a canonical base, then it has exactly one definably closed canonical base, which shall be referred to as *the* canonical base of $\mathbf{p}$ and denoted by $\mathrm{cb}(\mathbf{p})$.

Lemma 6.5. *Suppose T has* EI *and the global type $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$ is definable. Then $\mathbf{p}$ has a canonical base, and $\mathrm{cb}(\mathbf{p})$ is the smallest definably closed parameter set over which $\mathbf{p}$ is definable.*

Proof. Here we restrict to *finite* y. For each L-formula $\phi(x,y)$, let c_ϕ code the definable set $d_{\mathbf{p}}x\phi(x,\mathbb{M}_y)$. Then we have for $f \in \mathrm{Aut}(\mathbb{M})$:

$$f(\mathbf{p}) = \mathbf{p} \iff f(c_\phi) = c_\phi \text{ for all } L\text{-formulas } \phi(x,y).$$

Therefore, if A consists of the components of all the c_ϕ, then A is a canonical base of $\mathbf{p}$. The claim about $\mathrm{cb}(\mathbf{p})$ now follows from the exercise at the end of Section 4 and the exercise in this subsection. $\square$

Heirs. Let $p = p(x) \in \mathrm{St}_x(M)$. Given $M \subseteq B$, a *heir* of p over B is a son $q \in \mathrm{St}_x(B)$ of p such that for each L_M-formula $\phi(x,y)$, if $\phi(x,b) \in q$ for some $b \in B_y$, then $\phi(x,a) \in p$ for some $a \in M_y$.

Let $M \subseteq B \subseteq C$, and suppose $r \in \mathrm{St}_x(C)$ is a son of $q \in \mathrm{St}_x(B)$ and q is a son of p. Then it follows easily that if r is a heir of p, then q is a heir of p, and also that if q is a heir of p and $B = N$ is an elementary submodel of $\mathbb{M}$ and r is a heir of q, then r is a heir of p.

Lemma 6.6. *Let $p \in \mathrm{St}_x(M)$, and $M \subseteq B \subseteq C$. Then*

(1) *p has a heir over B;*
(2) *if q is a heir of p over B, then q extends to a heir of p over C;*
(3) *if p is definable, then p has exactly one heir over B, namely $p{\upharpoonright}B$.*

Proof. Since (1) is a special case of (2) we prove (2). Let q be a heir of p over B. Claim 1: let $\phi(x,y)$ be an L_M-formula and let $c \in C_y$ be such that $q(x) \models \phi(x,c)$; then $\phi(x,a) \in p$ for some $a \in M_y$. To see this, note first that $q(x) \cup \mathrm{tp}_y(c|B) \models \phi(x,y)$, since for any realization $(d,e) \in \mathbb{M}_x \times \mathbb{M}_y$ of this partial type over B in $\mathbb{M}$ there is an $f \in \mathrm{Aut}(\mathbb{M}|B)$ with $f(e) = c$.

Compactness yields an L_M-formula $\theta(y,z)$ and a tuple $b \in B_z$ such that $\theta(y,b) \in \mathrm{tp}_y(c|B)$ and $q(x) \cup \{\theta(y,b)\} \models \phi(x,y)$. Hence

$$q(x) \models \big(\exists y \theta(y,b)\big) \wedge \forall y \big(\theta(y,b) \to \phi(x,y)\big).$$

Since q is a heir of p this gives $b' \in M_z$ such that

$$p(x) \models \big(\exists y \theta(y,b')\big) \wedge \forall y \big(\theta(y,b') \to \phi(x,y)\big).$$

Hence $M \models \exists y \theta(y,b')$, so we have $a \in M_y$ with $\models \theta(a,b')$, and thus $\phi(x,a) \in p(x)$. This proves Claim 1.

Let $\Phi(x)$ be the set of L_C-formulas $\neg\phi(x,c)$ where $\phi(x,y)$ is an L_M-formula such that $\neg\phi(x,a) \in p(x)$ for every $a \in \mathcal{M}_y$, and $c \in C_y$. Claim 2: $q(x) \cup \Phi(x)$ is a partial x-type over C. If Claim 2 holds, then we can take $r(x) \in \mathrm{St}_x(C)$ containing this partial type, and then r is heir of p that extends q. Suppose Claim 2 is false. This gives L_M-formulas $\phi_1(x,y), \dots, \phi_n(x,y)$ and $c \in C_y$ such that

(1) $\neg\phi_i(x,a) \in p(x)$ for $i = 1, \dots, n$ and all $a \in \mathcal{M}_y$;
(2) $q(x) \cup \{\neg\phi_1(x,c), \dots, \neg\phi_n(x,c)\}$ cannot be realized over C.

Then $q(x) \models \phi_1(x,c) \vee \cdots \vee \phi_n(x,c)$, and thus by Claim 1 there is an $a \in M_y$ such that $\phi_1(x,a) \vee \cdots \vee \phi_n(x,a) \in p(x)$, a contradiction. $\square$

Coheirs. Let $p \in \mathrm{St}_x(M)$. Given $B \supseteq M$, a *coheir* of p over B is a son $q \in \mathrm{St}_x(B)$ of p such that every formula in $q(x)$ is realized by some $a \in M_x$. Note that if $a \in \mathbb{M}_x$ and $b \in \mathbb{M}_y$ and $\mathrm{tp}(a|Mb)$ is a coheir of $\mathrm{tp}(a|M)$ over Mb, then $\mathrm{tp}(b|Ma)$ is a heir of $\mathrm{tp}(b|M)$.

Let $M \subseteq B \subseteq C$, and suppose $r \in \mathrm{St}_x(C)$ is a son of $q \in \mathrm{St}_x(B)$ and q is a son of p, and r is a coheir of p. Then it follows easily that q is a coheir of p, and also that r is a coheir of q when $B = N$ is an elementary submodel of $\mathbb{M}$.

Lemma 6.7. *Let $p \in \mathrm{St}_x(M)$, and $M \subseteq B \subseteq C$. Then*

(1) *p has a coheir over B;*
(2) *if q is a coheir of p over B, then q extends to a coheir of p over C;*

Proof. It is enough to prove (2), since (1) is a special case. We shall use a topological argument. Let q be a coheir of p over B. Let R be the closure in $\mathrm{St}_x(C)$ of $\{\mathrm{tp}(a|C) : a \in M_x\}$. For every $\phi(x) \in q$ the closed set $[\phi(x)] \in \mathrm{St}_x(C)$ intersects $\{\mathrm{tp}(a|C) : a \in M_x\}$, and thus R. Hence by compactness R has a point $r(x)$ that lies in every $[\phi(x)]$ with $\phi(x) \in q(x)$. Then r extends q and is a coheir of p. $\square$

Lemma 6.8. *Let $M \subseteq B$ and suppose every y-type over M is definable, for every finite y. Then every $p \in \mathrm{St}_x(M)$ has a unique coheir over B.*

Proof. Let $p \in \mathrm{St}_x(M)$, and suppose that $a, a' \in \mathbb{M}_x$ realize different coheirs q and q' of p over B. Take a finite y disjoint from x, $b \in B_y$, and an L_M-formula $\phi(x,y)$ such that $\models \phi(a,b) \wedge \neg\phi(a',b)$, and thus $\phi(a,y) \in$

$\text{tp}(b|Ma)$ and $\phi(a',y) \notin \text{tp}(b|Ma')$. But $\text{tp}(b|Ma)$ and $\text{tp}(b|Ma')$ are heirs of $\text{tp}(b|M)$, which gives a contradiction using (3) of Lemma 11.16 and $\text{tp}(a|M) = \text{tp}(a'|M)$. □

7. Indiscernibles and NIP

Let I be a chain, that is, a linearly ordered set. An I-sequence is just a family $(a_i)_{i\in I}$, and is said to be an I-sequence in the set X if all $a_i \in X$. For $I = \mathbb{N}$, ordered in the usual way, we just say "sequence" instead of $\mathbb{N}$-sequence. An I-sequence (a_i) in $\mathbb{M}_x$, is said to be *indiscernible over A* if for every L_A-formula $\phi(x_1,\dots,x_n)$, with $x_1,\dots,x_n$ similar to x, and all $i_1 < \dots < i_n$ and $j_1 < \dots, j_n$ in I we have

$$\models \phi(a_{i_1},\dots,a_{i_n}) \iff \phi(a_{j_1},\dots,a_{j_n}).$$

For $A = \emptyset$ we just say "indiscernible".

Definable types and coheirs generate indiscernible sequences:

Lemma 7.1. *Let $q \in \text{St}_y(M)$ be definable and $a, b \in \mathcal{M}_x$ with $\text{tp}(a|M) = \text{tp}(b|M)$. Suppose $c \in \mathcal{M}_y$ realizes $q{\upharpoonright}Ma$ and $d \in \mathcal{M}_y$ realizes $q{\upharpoonright}Mb$. Then $\text{tp}\big(a,c)|M\big) = \text{tp}\big((b,d)|M\big)$.*

Proof. Let $\phi(u,x,y)$ be an L-formula, $e \in M_u$. Then

$$\begin{aligned} \models \phi(e,a,c,) \;&\Leftrightarrow\; \phi(e,a,y) \in q \;\Leftrightarrow \models d_q\phi(e,a) \;\Leftrightarrow \models d_q\phi(e,b) \\ &\Leftrightarrow\; \phi(e,b,y) \in q \;\Leftrightarrow \models \phi(e,b,d). \end{aligned}$$

□

Corollary 7.2. *Let $q \in \text{St}_y(M)$ be definable, and take a sequence (a_n) in $\mathbb{M}_y$ such that a_0 realizes q and a_{n+1} realizes $q{\upharpoonright}Ma_0\dots a_n$ for each n. Then (a_n) is indiscernible over M, and for each m the type $\text{tp}\big((a_0,\dots,a_m)|M\big)$ is independent of the choice of the sequence (a_n).*

Proof. Let $i_0 < \dots < i_n < i_{n+1}$ in $\mathcal{N}$, and suppose inductively that

$$\text{tp}\big((a_0,\dots,a_n)|M\big) \;=\; \text{tp}\big((a_{i_0},\dots,a_{i_n})|M\big).$$

Since a_{n+1} realizes $q{\upharpoonright}Ma_0\dots a_n$ and $a_{i_{n+1}}$ realizes $q{\upharpoonright}Ma_{i_0}\dots a_{i_n}$, it follows from the previous lemma that

$$\text{tp}\big((a_0,\dots,a_n,a_{n+1})|M\big) \;=\; \text{tp}\big((a_{i_0},\dots,a_{i_n},a_{i_{n+1}})|M\big).$$

The second claim of the lemma is proved in the same way. □

A sequence (a_n) as in this lemma is called a *Morley sequence* of q over M.

For $M \subseteq B$, a type $q \in \text{St}_y(B)$ is said to be *finitely satisfiable* in M if every formula $\phi(y) \in q$ is realized by some element of M_y, in other words, q is a coheir of its restriction to M.

Let $M \subseteq N$, and assume in the next lemma and its corollary that N is κ-saturated with $\kappa > |M|$.

Lemma 7.3. *Let $q \in \mathrm{St}_y(B)$ be finitely satisfiable in M and let $a, b \in N_x$ be such that $\mathrm{tp}(a|M) = \mathrm{tp}(b|M)$. Suppose $c \in N_y$ realizes $q|Ma$ and $d \in N_y$ realizes $q|Mb$. Then $\mathrm{tp}\big((a,c)|M\big) = \mathrm{tp}\big((b,d)|M\big)$.*

Proof. Let $\phi(x,y)$ be an L_M-formula such that $\models \phi(a,c)$, so $\phi(a,y) \in q$; it suffices to show that then $\phi(b,y) \in q$. Towards a contradiction, assume that $\neg\phi(b,y) \in q$. Then $\phi(a,y) \wedge \neg\phi(b,y) \in q$, so we have $e \in M_y$ such that $\models \phi(a,e) \wedge \neg\phi(b,e)$, contradicting $\mathrm{tp}(a|M) = \mathrm{tp}(b|M)$. □

Corollary 7.4. *Suppose $q \in \mathrm{St}_y(B)$ is finitely satisfiable in M. Take a sequence $a_0, a_1, a_2, \dots$ in N_y such that a_0 realizes $q|M$ and a_{n+1} realizes $q|Ma_0 \dots a_n$ for all n. Then the sequence (a_n) is indiscernible over M, and for each n the type $\mathrm{tp}\big((a_0,\dots,a_n)|M\big)$ depends only on M and q, not on the choice of the sequence (a_n).*

Proof. Let $i_0 < i_1 < \cdots < i_n < i_{n+1}$ (in $\mathcal{N}$), and assume inductively that

$$\mathrm{tp}\big((a_0, a_1, \dots, a_n)|M\big) \;=\; \mathrm{tp}\big((a_{i_0}, a_{i_1}, \dots, a_{i_n})|M\big).$$

Now observe that a_{n+1} realizes $q|Ma_0 \dots a_n$ and $a_{i_{n+1}}$ realizes $q|Ma_{i_0} \dots a_{i_n}$. Then by the lemma above we have

$$\mathrm{tp}\big((a_0, a_1, \dots, a_n, a_{n+1})|M\big) \;=\; \mathrm{tp}\big((a_{i_0}, a_{i_1}, \dots, a_{i_n}, a_{i_{n+1}})|M\big).$$

The second part of the lemma follows in the same way. □

A sequence (a_n) as in this lemma is called a *coheir sequence* of q over M.

Dependence and independence. In this section, x and y are disjoint, and $R \subseteq \mathbb{M}_{x,y} = \mathbb{M}_x \times \mathbb{M}_y$ is a definable relation. We say that R, and any $L_{\mathbb{M}}$-formula $\rho(x,y)$ that defines R, has the *independence property* (or *IP*) if for every $I \subseteq \mathbb{N}$ there are $a_i \in \mathbb{M}_x$ for $i \in \mathbb{N}$ and a $b_I \in \mathbb{M}_y$ such that

$$R(a_i, b_I) \iff i \in I, \qquad \text{for all } i \in \mathbb{N}.$$

Equivalently, for each n and $I \subseteq \{1,\dots,n\}$ there are $a_i \in \mathbb{M}_x$ for $i = 1,\dots,n$ and a $b_I \in \mathbb{M}_y$, such that

$$R(a_i, b_I) \iff i \in I, \qquad \text{for all } i \in \{1,\dots,n\}.$$

The ordered index set $\mathbb{N}$ in this definition is just for convenience: using the finite version it is easy to see that in the definition of *IP* one can replace $\mathbb{N}$ by any small infinite chain.

Note also that in this definition we single out not only a particular product set $\mathbb{M}_x \times \mathbb{M}_y$ but also $\mathbb{M}_x$ as its *first* factor and $\mathbb{M}_y$ as its *second* factor. In this connection, recall that an $L_{\mathbb{M}}$-formula $\rho(x,y)$ is formally a triple ρ, x, y; let $\breve{\rho}(y,x)$ be the triple ρ, y, x, so if $\rho(x,y)$ defines R, then $\breve{\rho}(y,x)$ defines the reversed relation $\breve{R} \subseteq \mathbb{M}_y \times \mathbb{M}_x$ given by

$$R(a,b) \iff \breve{R}(b,a), \qquad \text{for } a \in \mathbb{M}_x, b \in \mathbb{M}_y.$$

Our main interest here is when R is *dependent* (that is, R does not have IP), in which case also any $L_{\mathbb{M}}$-formula $\rho(x,y)$ that defines it is said to be dependent.

Lemma 7.5. *Let $\rho(x,y)$ be a dependent $L_{\mathbb{M}}$-formula. Then*

1. *the formula $\breve{\rho}(y,x)$ is dependent;*
2. *the formula $\neg\rho(x,y)$ is dependent;*
3. *if the $L_{\mathbb{M}}$-formula $\rho'(x,y')$ is dependent with y' disjoint from x and y, then $(\rho \vee \rho')(x,(y,y'))$ and $(\rho \wedge \rho')(x,(y,y'))$ are dependent;*
4. *if $y=(u,v)$ and $c \in \mathbb{M}_v$, then $\rho_1(x,u) := \rho(x,u,c)$ is dependent.*

8. Stability

In this section, x and y are disjoint and finite, and $R \subseteq \mathbb{M}_{x,y} = \mathbb{M}_x \times \mathbb{M}_y$ is a definable relation. We say that R, and any $L_{\mathbb{M}}$-formula $\delta(x,y)$ that defines R, is *unstable* if there are $a_i \in \mathbb{M}_x$ and $b_i \in \mathbb{M}_y$ for $i \in \mathbb{N}$ such that

$$R(a_i,b_j) \Longleftrightarrow i<j, \qquad \text{for all } i,j \in \mathbb{N}.$$

Equivalently, for each n there are $a_i \in \mathbb{M}_x$ and $b_i \in \mathbb{M}_y$ for $i=1,\dots,n$, such that

$$R(a_i,b_j) \Longleftrightarrow i<j, \qquad \text{for all } i,j \in \{1,\dots,n\}.$$

The ordered index set $\mathbb{N}$ in this definition is just for convenience: using the finite version it is easy to see that in the definition of *unstable* one can replace $\mathbb{N}$ by any small infinite linearly ordered set.

Note also that in this definition we single out not only a particular product set $\mathbb{M}_x \times \mathbb{M}_y$ but also $\mathbb{M}_x$ as its *first* factor and $\mathbb{M}_y$ as its *second* factor. In this connection, recall that an $L_{\mathbb{M}}$-formula $\delta(x,y)$ is formally a triple δ, x, y; let $\breve{\delta}(y,x)$ be the triple δ, y, x, so if $\delta(x,y)$ defines R, then $\breve{\delta}(y,x)$ defines the reversed relation $\breve{R} \subseteq \mathbb{M}_y \times \mathbb{M}_x$ given by

$$R(a,b) \iff \breve{R}(b,a), \qquad \text{for } a \in \mathbb{M}_x, b \in \mathbb{M}_y.$$

As the terminology suggests, our main interest here is when R is *stable* (that is, not unstable), in which case also any $L_{\mathbb{M}}$-formula $\delta(x,y)$ that defines it is said to be stable.

Lemma 8.1. *Let $\delta(x,y)$ be a stable $L_{\mathbb{M}}$-formula. Then*

1. *the formula $\breve{\delta}(y,x)$ is stable;*
2. *the formula $\neg\delta(x,y)$ is stable;*
3. *if the $L_{\mathbb{M}}$-formula $\delta'(x,y')$ is stable with y' finite and disjoint from x and y, then $(\delta \vee \delta')(x,(y,y'))$ and $(\delta \wedge \delta')(x,(y,y'))$ are stable;*
4. *if $y=(u,v)$ and $c \in \mathbb{M}_v$, then $\delta_1(x,u) := \delta(x,u,c)$ is stable.*

To prove these items, consider the contrapositives. For example, in (3), suppose $(\delta \vee \delta')(x,(y,y'))$ is unstable. Take $a_i \in \mathbb{M}_x, b_i \in \mathbb{M}_y, b'_i \in \mathbb{M}_{y'}$ for $i \in \mathbb{N}$, such that for all $i,j \in \mathbb{N}$: $\models \delta(a_i,b_j) \vee \delta'(a_i,b'_j) \iff i<j$. Let

$$P := \{(i,j) \in \mathbb{N}^2 : i<j \text{ and } \models \delta(a_i,b_j)\},$$
$$P' := \{(i,j) \in \mathbb{N}^2 : i<j \text{ and } \models \delta'(a_i,b'_j)\}.$$

Then $P \cup P' = \{(i,j) \in \mathbb{N}^2 : i<j\}$, hence by Ramsey's theorem there is an infinite subset I of $\mathbb{N}$ such that either all pairs $(i,j) \in I^2$ with $i<j$ belong

to P, in which case $\delta(x,y)$ is unstable, or all pairs $(i,j) \in I^2$ with $i < j$ belong to P', in which case $\delta'(x,y')$ is unstable.

Exercise. Let $f : \mathbb{M}_{x'} \to \mathbb{M}_x$ and $g : \mathbb{M}_{y'} \to \mathbb{M}_y$ be definable, with disjoint finite x', y', and define the relation $R' \subseteq \mathbb{M}_{x'} \times \mathbb{M}_{y'}$ by

$$R'(a,b) \iff R(fa, gb).$$

Show that if R is stable, then R' is stable. Show that if f and g are surjective and R' is stable, then R is stable.

Let $\delta(x,y)$ be an L-formula, not just an $L_{\mathbb{M}}$-formula, and let $\mathcal{M} = (M; \cdots)$ be any model of T. Then $\delta(x,y)$ is unstable if and only if for each n there are $a_i \in M_x$ and $b_i \in M_y$ for $i = 1, \dots, n$, such that

$$\mathcal{M} \models \delta(a_i, b_j) \iff i < j, \quad \text{for all } i, j \in \{1, \dots, n\}.$$

It follows that the stability of $\delta(x,y)$ depends only on the theory T, not on the particular monster model $\mathbb{M}$ of T used in the definition.

The next results show how unstability relates to linear orderability.

Lemma 8.2. *Suppose R is unstable. Define $R^* \subseteq \mathbb{M}_{x,y} \times \mathbb{M}_{x,y}$ by*

$$R^*(a,b;a',b') \iff R(a,b'), \quad (a,a' \in \mathbb{M}_x, b, b' \in \mathbb{M}_y).$$

Then there are $c_i \in \mathbb{M}_{x,y}$ for $i \in \mathbb{N}$ such that for all $i,j \in \mathbb{N}$,

$$R^*(c_i, c_j) \iff i < j.$$

Lemma 8.3. *Let $\delta(x,y)$ be an unstable L-formula, and $\kappa \geq |L|$. Then there is a model $\mathcal{M} = \big(M; \cdots\big)$ of size κ such that $\mathrm{St}_x^{\mathcal{M}}(M)$ has size $> \kappa$.*

Proof. First construct a linearly ordered set I of size κ with more than κ upward closed subsets (see exercise below). With i and j ranging over I, we use compactness to obtain a model $\mathcal{M} = \big(M; \cdots\big)$ with elements $a_i \in M_x$ and b_i in M_y such that

$$\mathcal{M} \models \delta(a_i, b_j) \iff i < j, \quad \text{for all } i, j.$$

Using Löwenheim-Skolem, arrange that $\mathcal{M}$ has size κ. To each upward closed set $U \subseteq I$ with $U \neq I$ we associate the set of L_M-formulas

$$\Phi_U(x) := \{\delta(x, b_j) : j \in U\} \cup \{\neg\delta(x, b_j) : j \notin U\}.$$

One checks easily that $\Phi_U(x)$ is a partial x-type over M in $\mathcal{M}$. Let $p_U(x)$ be an x-type over M in $\mathcal{M}$ that extends $\Phi_U(x)$. Then $p_U(x) \neq p_{U'}(x)$ if $U \neq I$ and $U' \neq I$ are distinct upward closed subsets of I. $\square$

Exercise. Let κ be given, and let λ be the smallest cardinal such that $2^\lambda > \kappa$, so $\lambda \leq \kappa$. We order the set G of all functions $g : \lambda \to 2 = \{0,1\}$ lexicographically, that is, $g < h$ iff for some ordinal $\beta < \lambda$ we have

$$g(\beta) < h(\beta), \quad g(\alpha) = h(\alpha) \text{ for all ordinals } \alpha < \beta.$$

Observe that this makes G a linearly ordered set. Let I be the (ordered) subset of G consisting of those $i \in G$ that are eventually constant. For

$g \in G$, put $U(g) := \{i \in I : i \geq g\}$. Show that I has size $\leq \kappa$, and that if $g, h \in G$, $g \neq h$, then $U(g) \neq U(h)$.

If $\delta(x, y)$ is stable, then for some positive integer N there are no $a_i \in \mathbb{M}_x$ and $b_i \in \mathbb{M}_y$ for $i = 1, \dots, N$ such that

$$\models \delta(a_i, b_j) \iff i \leq j, \quad \text{for } i, j = 1, \dots, N,$$

and below $N(\delta)$ denotes a positive integer N with this property. (Here we abuse language by letting δ stand for $\delta(x, y)$.) The next result is crucial.

Lemma 8.4. *Let $\delta(x, y)$ be a stable L-formula, $a \in \mathbb{M}_x$, and M a model. Then there are positive integers I, J and elements $a^i_j \in M_x$ for $1 \leq i \leq I$ and $1 \leq j \leq J$, such that for each $b \in M_y$,*

$$\models \ \delta(a, b) \ \longleftrightarrow \ \bigvee_{i=1}^{I} \big(\bigwedge_{j=1}^{J} \delta(a^i_j, b) \big).$$

Proof. Put $J := N(\delta)$, and let $\psi(x) \in \operatorname{tp}(a|M)$.

Claim. There are $a_1, \dots, a_n \in M_x$ realizing $\psi(x)$, with $1 \leq n \leq J$, such that for all $b \in M_y$ we have $\models \big(\bigwedge_{i=1}^{n} \delta(a_i, b) \big) \longrightarrow \delta(a, b)$.

To prove this claim, define a ψ-sequence of length n to be a sequence

$$(a_1, \dots, a_n, b_1, \dots, b_n), \qquad (a_1, \dots, a_n \in M_x, \ b_1, \dots, b_n \in M_y)$$

such that for all $i, j \in \{1, \dots, n\}$,

$$\models \delta(a_i, b_j) \iff i \leq j, \qquad \models \psi(a_i) \wedge \neg\delta(a, b_i).$$

Note that we have a ψ-sequence of length 0. Let $(a_1, \dots, a_n, b_1, \dots, b_n)$ be a ψ-sequence, so $n < J$. The L_M-formula $\psi(x) \wedge \bigwedge_i \neg\delta(x, b_i)$ is realized by a, and thus realized by an element $a_{n+1} \in M_x$. If there is a $b \in M_y$ such that $\models \bigwedge_{i=1}^{n+1} \delta(a_i, b) \wedge \neg\delta(a, b)$, then we let b_{n+1} be such a b, and we have extended our sequence to a longer ψ-sequence $(a_1, \dots, a_{n+1}, b_1, \dots, b_{n+1})$. If there is no such b, then the claim holds for $a_1, \dots, a_{n+1} \in M_x$. Since ψ-sequences have length $< J$, the extension process must come to a halt. Thus the claim is established. Note that in this claim we can arrange $n = J$. For use in the second part of the proof, we say that a tuple $\vec{a} = (a_1, \dots, a_J) \in M_x^J$ *satisfies the claim* if for all $b \in M_y$ we have

$$\models \big(\bigwedge_{j=1}^{J} \delta(a_j, b) \big) \ \longrightarrow \ \delta(a, b).$$

Now the second part of the proof. Introduce a multivariable $\vec{x} = (x_1, \dots, x_J)$ where the x_j are multivariables similar to x, pairwise disjoint, and disjoint from y. Consider the L-formula

$$\delta_J(\vec{x}, y) \ := \ \delta(x_1, y) \wedge \dots \wedge \delta(x_J, y).$$

Then $\delta_J(\vec{x}, y)$ is stable, and so is $\neg\delta_J(\vec{x}, y)$. Put $I := N(\neg\delta_J)$. Suppose the tuples $\vec{a}^1, \vec{a}^2, \dots, \vec{a}^n \in M_x^J = M_{\vec{x}}$ satisfy the claim, and $b_1, \dots, b_n \in M_y$ are such that for all $i, j \in \{1, \dots, n\}$

$$\models \delta_J(\vec{a}^i, b_j) \iff i > j, \qquad \models \delta(a, b_j).$$

Let us express this by saying that $(\vec{a}^1, \vec{a}^2, \dots, \vec{a}^n, b_1, \dots, b_n)$ is a δ-sequence. Note that then $n < I$, and that for $n = 0$ there exists a δ-sequence. We now try to extend our δ-sequence. Applying the claim with $\psi(x) := \bigwedge_{j=1}^n \delta(x, b_j)$ yields a tuple $\vec{a}^{n+1} \in M_{\vec{x}}$ such that $\models \delta_J(\vec{a}^{n+1}, b_j)$ for $j = 1, \dots, n$, and for each $b \in M_y$, $\models \delta_J(\vec{a}^{n+1}, b) \to \delta(a, b)$. If there is a $b \in M_y$ such that $\models \delta(a, b)$ and $\models \neg\delta_J(\vec{a}^i, b)$ for $i = 1, \dots, n+1$, then we let b_{n+1} be such a b, and we have constructed a longer δ-sequence

$$(\vec{a}^1, \vec{a}^2, \dots, \vec{a}^{n+1}, b_1, \dots, b_{n+1}).$$

If there is no such b, then for all $b \in M_y$ we have $\models \delta(a, b) \to \bigvee_{i=1}^{n+1} \delta_J(\vec{a}^i, b)$, and thus

$$\models\ \delta(a, b) \ \longleftrightarrow\ \bigvee_{i=1}^{n+1} \delta_J(\vec{a}^i, b).$$

Then the lemma holds for the above values of I and J by taking the a_j^i such that $\vec{a}^i = (a_1^i, \dots, a_J^i)$ for $i = 1, \dots, n+1$, and $a_j^i = a_j^1$ for $n + 1 < i \le I$. The process of extending a δ-sequence must come to a halt. □

Remarks. The proof shows that we can take I and J to depend only on $\delta(x, y)$, not on a or M. In addition, if $\psi(x) \in \mathrm{tp}(a|M)$, we can take the a_j^i to realize $\psi(x)$. Also, if the model M is κ-saturated and $A \subseteq M$ has size $< \kappa$, then we can choose the a_j^i to realize $\mathrm{tp}(a|A)$.

A striking consequence of the lemma is that the subset $\delta(a, M_y)$ of M_y is definable in the model M, although a might not be an M-tuple!

Corollary 8.5. *Given the L-formula $\delta(x, y)$, the following are equivalent:*

(1) *$\delta(x, y)$ is stable;*
(2) *every δ-type over any M is definable;*
(3) *$|\,\mathrm{St}_\delta(M)| \le \kappa$, for each $\kappa \ge |L|$ and each M of size κ.*
(4) *for some small $\kappa \ge |L|$ we have $|\,\mathrm{St}_\delta(M)| \le \kappa$ for each M of size κ.*

Proof. For (1) ⇒ (2), assume (1) and let $p(x)$ be a δ-type over M. Take $a \in \mathbb{M}_x$ realizing $p(x)$. Then Lemma 8.4 yields an L_M-formula $\psi(y)$ such that $\delta(a, M_y) = \psi(M_y)$, that is, $\psi(y)$ defines $p(x)$.

The implications (2) ⇒ (3) and (3) ⇒ (4) are clear. The implication (4) ⇒ (1) is clear from the proof of Lemma 8.3. □

We say that T is *stable* if each L-formula $\delta(x, y)$ is stable. To prove that stability of T is equivalent to other conditions on T we use that a type is determined by the δ-types contained in it for relevant δ. It follows that if T is stable, then each x-type over M has a defining scheme, hence $|\,\mathrm{St}_x(M)| \le$

$\kappa^{|L|}$ for M of size $\kappa \geq |L|$. Here is a useful consequence of stability for global types.

Corollary 8.6. *If T is stable and has* EI*, then every global x-type has a canonical base.*

Proof. Stability of T implies that global types are definable, since we may view $\mathbb{M}$ as a "small" elementary submodel of an even bigger monster. Now apply Lemma 6.5. $\square$

From now on we shall also assume (as we may) that

$$\kappa(\mathbb{M}) > 2^{|L|}.$$

This will be used in proving some of the implications below.

Corollary 8.7. *The following are equivalent:*

(1) *T is stable;*
(2) *no 0-definable binary relation $R \subseteq \mathbb{M}_x \times \mathbb{M}_x$ with finite x linearly orders any infinite subset of $\mathbb{M}_x$ when restricted to that subset;*
(3) *whenever $Y \subseteq \mathbb{M}_y$ with finite y is definable, then the subset $Y \cap M_y$ of M_y is definable in the model M;*
(4) *$|\operatorname{St}_x(M)| \leq \kappa^{|L|}$, for each finite x, $\kappa \geq |L|$ and M of size κ;*
(5) *for some small $\kappa \geq |L|$ we have $|\operatorname{St}_x(M)| \leq \kappa$ for each finite x and M of size κ;*
(6) *for some small $\kappa \geq |L|$ we have $|\operatorname{St}_v(M)| \leq \kappa$ for each variable v and M of size κ;*
(7) *each L-formula $\phi(v, y)$ where v is a variable and y is finite, is stable.*

Proof. The equivalence of (1) and (2) is clear from the definition of "stable" and Lemma 8.2. The equivalence of (1) and (3) follows from the equivalence of (1) and (2) in Corollary 8.5. The implication (1) $\Rightarrow$ (4) was already noted. From (4) we obtain (5) with $\kappa = 2^{|L|}$, since this κ satisfies $\kappa^{|L|} = \kappa$. The implication (5) $\Rightarrow$ (6) is obvious. The implication (6) $\Rightarrow$ (5) follows by induction on the number of variables in x from Lemma 5.7, and (5) $\Rightarrow$ (1) follows from (3) $\Rightarrow$ (1) of Corollary 8.5. It remains to show that (7) is equivalent to (6). The implication (6) $\Rightarrow$ (7) follows from (4) $\Rightarrow$ (1) of Corollary 8.5. Assume (7), and let v be a variable. Then each v-type over any M has a defining scheme, hence $|\operatorname{St}_v(M)| \leq \kappa^{|L|}$ for M of size $\kappa \geq |L|$, and thus (6) holds with $\kappa = 2^{|L|}$. $\square$

Remark. The reader may find it annoying that some of the above conditions refer implicitly or explicitly to the monster model. Since $\mathbb{M}$ and $\kappa(\mathbb{M})$ can be chosen as large as we want, it is easy to see that if T is stable, then (3) actually holds with M and $\mathbb{M}$ replaced by any $\mathcal{M}$ and $\mathcal{N}$ such that $\mathcal{M} \preceq \mathcal{N} \models T$, and that (4) holds with M replaced by arbitrary models $\mathcal{M}$ of T. Likewise, (6) $\Rightarrow$ (1) holds without κ restricted to being small and with M replaced by arbitrary models $\mathcal{M}$ of T.

Exercise. Suppose T is stable. Show that for each A the L_A-theory T_A of $\mathbb{M}$ is stable. Show that T^{eq} is stable.

Suppose now that $\delta(x,y)$ is a stable L-formula and $p(x)$ is a δ-type over M defined by $\psi(y)$. Lemma 8.4 shows that $\psi(y)$ is a $\breve{\delta}$-formula over M, in particular, either $\psi(y) \in q$ or $\neg\psi(y) \in q$ whenever q is a $\breve{\delta}$-type over M. This remark is relevant in connection with the following symmetry result.

Lemma 8.8. *Let $\delta(x,y)$ be a stable L-formula and $A \subseteq M$. Let $p(x)$ be a δ-type over M defined by the L_A-formula $\psi(y)$, and let $q(y)$ be a $\breve{\delta}$-type over M defined by the L_A-formula $\chi(x)$. Then*

$$\psi(y) \in q(y) \iff \chi(x) \in p(x).$$

Proof. Suppose towards a contradiction that $\psi(y) \in q(y)$ but $\chi(x) \notin p(x)$. Let $\kappa := \max\{|A|, |L|\}^+$. By working in a big enough monster model we can pass to a κ-saturated $N \supseteq M$ and replace p and q by $p \upharpoonright N$ and $q \upharpoonright N$ we reduce to the situation that M is κ-saturated. Take $a_0 \in M_x$ realizing $p(x) \upharpoonright A$, so $\models \neg\chi(a_0)$, hence $\neg\breve{\delta}(y, a_0) \in q(y)$, that is, $\neg\delta(a_0, y) \in q(y)$. Next, take $b_0 \in M_y$ realizing $q(y) \upharpoonright Aa_0$, so $\models \neg\delta(a_0, b_0)$ and $\models \psi(b_0)$, hence $\delta(x, b_0) \in p(x)$. Now, take $a_1 \in M_x$ realizing $p(x) \upharpoonright Ab_0$, so $\models \delta(a_1, b_0)$ and $\models \neg\chi(a_1)$, and thus $\neg\delta(a_1, y) \in q(y)$. Continuing this way, we construct infinite sequences $a_0, a_1, a_2, \ldots$ in M_x and $b_0, b_1, b_2, \ldots$ in M_y, with a_{n+1} realizing $p(x) \upharpoonright Ab_1 \ldots b_n$ and b_n realizing $q(y) \upharpoonright Aa_1 \ldots a_n$, for all n. One shows by induction that $\models \neg\chi(a_n) \wedge \psi(b_n)$, and thus, for all m, n:

$$\models \delta(a_m, b_n) \iff m > n,$$

contradicting the stability of $\delta(x,y)$. $\square$

Indiscernibility. Let $(a_i)_{i \in \mathbb{N}}$ be a sequence in $\mathbb{M}_x$. We say that (a_i) is *indiscernible over A* if

$$\mathrm{tp}((a_{i_1}, \ldots, a_{i_n})|A) = \mathrm{tp}((a_{j_1}, \ldots, a_{j_n})|A)$$

for all n and all n-element sets $\{i_1, \ldots, i_n\}, \{j_1, \ldots, j_n\} \subseteq \mathbb{N}$. We say that (a_i) is *order-indiscernible over A* if the displayed identity holds for all n and $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$ in $\mathbb{N}$. In the presence of stability these two notions coincide:

Lemma 8.9. *Suppose T is stable and the sequence $(a_i)_{i \in \mathbb{N}}$ in $\mathbb{M}_x$ is order-indiscernible over A. Then (a_i) is indiscernible over A.*

Proof. To keep notation simple, assume $A = 0$. We first prove the special case $\mathrm{tp}(a_0, a_1, a_2) = \mathrm{tp}(a_1, a_0, a_2)$. Suppose $\mathrm{tp}(a_0, a_1, a_2) \neq \mathrm{tp}(a_1, a_0, a_2)$. Then we have an L-formula $\phi(x, y, z)$ with y and z similar to x such that

$$\models \phi(a_0, a_1, a_2) \quad \text{and} \ \models \neg\phi(a_1, a_0, a_2), \ \text{hence}$$
$$\models \phi(a_i, a_j, a_k) \quad \text{and} \ \models \neg\phi(a_j, a_i, a_k) \ \text{ for all } i < j < k.$$

By saturation we can take $a \in \mathbb{M}_x$ such that $\models \phi(a_i, a_j, a)$ and $\models \neg\phi(a_j, a_i, a)$ for all $i < j$, contradicting stability. Having now established the special case, it follows that $\mathrm{tp}(a_i, a_j, a_k) = \mathrm{tp}(a_j, a_i, a_k)$ for all $i < j < k$ in $\mathbb{N}$.

For the general case, proceed likewise, and use that the permutations of $\{1, \dots, n\}$ are generated by the transpositions $(i, i+1)$ with $1 \le i < n$. $\square$

9. Rank and degree in model theory

From now on x, y, z are *finite* multivariables, unless specified otherwise.

Let $\mathcal{M}$ be an L-structure and Y a definable set in $\mathcal{M}$, say $Y \subseteq M_y$. Then we define the *Cantor rank* of Y (in $\mathcal{M}$) to be the Cantor rank of Y as an element of the boolean algebra $\mathrm{Def}_y(\mathcal{M})$ of definable subsets of M_y, equivalently, it is the Cantor rank of Y as an element of the boolean algebra of definable subsets of Y. We denote it by $\mathrm{CR}(Y)$, and by $\mathrm{CR}^{\mathcal{M}}(Y)$ if we wish to indicate the ambient structure $\mathcal{M}$. If $\mathrm{CR}(Y)$ is an ordinal, then we call Y ranked, and define $\mathrm{CD}(Y)$ (the *Cantor degree* of Y) to be the Cantor degree of Y as an element of the boolean algebra of definable subsets of M_y. (If we wish to indicate the ambient structure $\mathcal{M}$ we write this degree as $\mathrm{CD}^{\mathcal{M}}(Y)$.) Suppose now that X is also a definable set in $\mathcal{M}$, say $X \subseteq M_x$, and that the map $f : X \to Y$ is definable. Then

(i) if $f(X) = Y$, then $\mathrm{CR}(X) \ge \mathrm{CR}(Y)$, and in case of equality with X ranked we have $\mathrm{CD}(X) \ge \mathrm{CD}(Y)$.
(ii) if f is injective, then $\mathrm{CR}(X) \le \mathrm{CR}(Y)$, and in case of equality with X ranked we have $\mathrm{CD}(X) \le \mathrm{CD}(Y)$.

To see this, note that $U \mapsto f^{-1}(U)$ is a morphism from the boolean algebra of definable subsets of Y into the boolean algebra of definable subsets of X, and that this morphism sends Y to X. If $f(X) = Y$, then this morphism is injective, so we can apply part (1) of Proposition 2.3. If f is injective, then this morphism is surjective, so we can apply part (2) of that proposition.

Unfortunately, the Cantor rank of a definable set may change in passing to an elementary extension of $\mathcal{M}$. Take for example $\mathcal{M} = (\mathbb{N}, <)$ and let $\mathcal{N} = (N, <)$ be an $\aleph_0$-saturated elementary extension of $\mathcal{M}$. The definable subsets of $\mathbb{N}$ in $\mathcal{M}$ are exactly the finite and cofinite subsets of $\mathbb{N}$ (exercise), so $\mathbb{N}$ has Cantor rank 1 as a definable set in $\mathcal{M}$. But N as a definable set in $\mathcal{N}$ has Cantor rank > 1, as is easily verified.

More generally, let $\mathcal{M} \preceq \mathcal{N}$, so we have a boolean algebra embedding

$$\iota : \mathrm{Def}_y(\mathcal{M}) \to \mathrm{Def}_y(\mathcal{N}), \qquad \phi(a, M_y) \mapsto \phi(a, N_y),$$

where $\phi(x, y)$ ranges over L-formulas of the indicated form with $a \in M_x$. By Proposition 2.3 we have $\mathrm{CR}(Y) \le \mathrm{CR}(\iota Y)$ for $Y \in \mathrm{Def}_y(\mathcal{M})$. With a mild saturation assumption we have equality:

Lemma 9.1. *Let $\mathcal{M}$ be $\aleph_0$-saturated. Then for all $Y \in \mathrm{Def}_y(\mathcal{M})$,*

$$\mathrm{CR}(Y) = \mathrm{CR}(\iota Y), \qquad \mathrm{CD}(Y) = \mathrm{CD}(\iota Y) \text{ in case } Y \text{ is ranked.}$$

Proof. By increasing $\mathcal{N}$ we can assume $\mathcal{N}$ is strongly $\aleph_1$-homogeneous. Let $Y \in \mathrm{Def}_y(\mathcal{M})$. The rank equality follows from

$$\mathrm{CR}(\iota Y) > \lambda \implies \mathrm{CR}(Y) > \lambda, \quad \text{for all } \lambda,$$

which we prove by induction on λ. Suppose $\mathrm{CR}(\iota Y) > \lambda$. Then we have disjoint definable sets $Y_n \subseteq \iota Y$ of Cantor rank $\geq \lambda$ in $\mathcal{N}$. Take a finite tuple a in $\mathcal{M}$ such that Y is a-definable in $\mathcal{M}$. Take a tuple b in $\mathcal{N}$ of countable size such that all Y_n are b-definable in $\mathcal{N}$. By the saturation and homogeneity assumptions on $\mathcal{M}$ and $\mathcal{N}$ we have an automorphism $f \in \mathrm{Aut}(\mathcal{N}|a)$ such that fb is an M-tuple. Then $f(\iota Y) = \iota Y$ and $f(Y_n)$ is definable over M, so $f(Y_n) = \iota X_n$ with $X_n \subseteq Y$ definable in $\mathcal{M}$. Since

$$\mathrm{CR}(\iota X_n) = \mathrm{CR}\, f(Y_n) = \mathrm{CR}(Y_n) \geq \lambda,$$

we can assume inductively that $\mathrm{CR}(X_n) \geq \lambda$. The X_n being disjoint subsets of Y, we obtain $\mathrm{CR}(Y) > \lambda$. The degree equality follows in a similar way. $\square$

Let $Y \in \mathrm{Def}_y(\mathcal{M})$ be given. Choose an $\aleph_0$-saturated elementary extension $\mathcal{N}$ of $\mathcal{M}$ and let ι be as above. Then by the lemma $\mathrm{CR}(\iota Y)$ is independent of the choice of $\mathcal{N}$, so we can now define the *Morley rank* of Y by

$$\mathrm{MR}(Y) = \mathrm{MR}^{\mathcal{M}}(Y) := \mathrm{CR}(\iota Y)$$

In case this rank is an ordinal, the *Morley degree* of Y is defined likewise by

$$\mathrm{MD}(Y) = \mathrm{MD}^{\mathcal{M}}(Y) := \mathrm{CD}(\iota Y).$$

Often $\mathcal{M}$ is already $\aleph_0$-saturated, and then $\mathrm{MR}(Y) = \mathrm{CR}(Y)$, and in case $\mathrm{MR}(Y)$ is an ordinal, $\mathrm{MD}(Y) = \mathrm{CD}(Y)$; this applies in particular to the definable sets in our monster model $\mathbb{M}$.

Exercise. Suppose $X \subseteq \mathbb{M}_x$ is M-definable (so the set $X(M) := X \cap M_x$ is definable in the model M). Prove that

$$\mathrm{CR}^M\big(X(M)\big) \leq \mathrm{MR}(X).$$

Totally transcendental theories. We define T to be *totally transcendental* if every nonempty definable set in $\mathbb{M}$ is ranked. By Lemma 9.1 this is indeed a property of T, that is, it does not depend on the particular monster model $\mathbb{M}$ of T.

Lemma 9.2. *Suppose T is totally transcendental. Then*

(1) *the theory T_A of the L_A-structure $\mathbb{M}$ is totally transcendental;*
(2) *T^{eq} is totally transcendental;*
(3) $|\,\mathrm{St}_x(M)| \leq |M|$ *for* $|M| \geq |L|$;
(4) *T is stable.*

Proof. Item (1) is obvious, and (2) follows from the fact that every definable set in $\mathbb{M}^{\mathrm{eq}}$ is the image of a definable set in $\mathbb{M}$ under a definable map. Item (3) follows from the implication (iii)⇒(iv) of Theorem 2.5 applied to the subalgebras $B = \mathrm{Def}_x(\mathbb{M}|M)$ of $\mathrm{Def}_x(\mathbb{M})$ for $|M| \geq |L|$, since $|B| \leq |M|$ for such M. Item (4) follows from (3) and Corollary 8.7. □

When the language is countable we can say a bit more:

Corollary 9.3. *Let L be countable. The following are equivalent:*

(1) *T is totally transcendental;*
(2) *$|\mathrm{St}_x(M)| \leq |M|$ for each x and infinite M;*
(3) *$|\mathrm{St}_x(M)| \leq \aleph_0$ for each x and countable M;*
(4) *$|\mathrm{St}_v(M)| \leq \aleph_0$ for each variable v and countable M;*
(5) *each nonempty definable subset of $\mathbb{M}_v$ is ranked, for each variable v.*

Proof. The implication (1) ⇒ (2) is part of the previous lemma. The implications (2) ⇒ (3) and (3) ⇒ (4) are obvious. The equivalence (4) ⇔ (5) follows from the equivalence (3) ⇔ (4) of Theorem 2.5 by noting that each countable subalgebra of $\mathrm{Def}_v(\mathbb{M})$ is contained in a subalgebra of the form $\mathrm{Def}_v(\mathbb{M}|M)$ with countable M. The implication (4) ⇒ (3) follows by induction on the size of x from Lemma 5.7, and (3) ⇒ (1) is again an easy consequence of Theorem 2.5. □

By Theorem 2.5, all nonzero elements in a boolean algebra are ranked if each countable subalgebra has this property. Thus:

Corollary 9.4. *The following are equivalent:*

(1) *T is totally transcendental;*
(2) *for each countable sublanguage L_0 of L the theory of the L_0-reduct of $\mathbb{M}$ is totally transcendental;*
(3) *each nonempty definable subset of $\mathbb{M}_v$ is ranked, for each variable v.*

Proof. The equivalence of (1) and (2) follows easily from Theorem 2.5. We already know that (1) implies (3), and the implication (3) ⇒ (2) follows from the previous corollary. □

For countable L, a totally transcendental T is also said to be *omega-stable* (or *ω-stable*).

Stability and rank. Let $\delta(x, y)$ be an L-formula, and consider the boolean subalgebra $\mathrm{Def}_\delta(\mathbb{M})$ of $\mathrm{Def}_x(\mathbb{M})$ that is generated by the sets $\delta(\mathbb{M}_x, b)$ with $b \in \mathbb{M}_y$. For $X \in \mathrm{Def}_\delta(\mathbb{M})$ we define the δ-rank $\mathrm{R}_\delta(X)$ of X to be its Cantor rank in this boolean algebra, and if this rank is an ordinal, we let $\mathrm{D}_\delta(X)$ be the Cantor degree of X in this boolean algebra.

Proposition 9.5. *The formula $\delta(x, y)$ is stable if and only if each nonempty $X \in \mathrm{Def}_\delta(\mathbb{M})$ is ranked.*

Proof. We prove the contrapositives. Suppose $\delta(x, y)$ is unstable. Then we have a model M of size $\kappa \geq |L|$ such that $\mathrm{St}_\delta(M) > \kappa$. Then by Theorem 2.5, the rank of $\mathbb{M}_x$ as an element of the boolean subalgebra $\mathrm{Def}_\delta(M)$ of $\mathrm{Def}_\delta(\mathbb{M})$ is $+\infty$, and thus $\mathrm{R}_\delta(\mathbb{M}_x) = +\infty$.

For the converse, suppose $\mathrm{R}_\delta(\mathbb{M}_x) = +\infty$. Then we take an M of size $\kappa \geq |L|$ such that $\mathbb{M}_x$ as an element of the boolean algebra $\mathrm{Def}_\delta(M)$ has rank $+\infty$. For each n, Corollary 2.8 yields elements $b_1, \dots, b_n \in M_y$ and types $p_0(x), \dots, p_n(x) \in \mathrm{St}_x(M)$ such that

$$\delta(x, b_i) \in p_j(x) \iff i \leq j, \text{ for } i = 1, \dots, n, \ j = 0, \dots, n.$$

With such b_i and p_j we take $a_j \in \mathbb{M}_x$ realizing p_j and obtain

$$\models \delta(a_j, b_i) \iff i \leq j, \text{ for } i = 1, \dots, n, \ j = 0, \dots, n.$$

Hence $\delta(x, y)$ is unstable. □

10. Basic facts on Morley rank

From now on we restrict our attention mainly to Morley rank and totally transcendental theories, although much will go through, by similar arguments, for δ-rank and stable theories. We begin with improving the result that $\mathrm{MR}(X) \leq \mathrm{MR}(Y)$ if there is an injective definable map from X to Y.

Lemma 10.1. *Let the sets X, Y and the map $f : X \to Y$ be definable in $\mathbb{M}$ such that $f^{-1}(b)$ is finite for all $b \in Y$. Then $\mathrm{MR}(X) \leq \mathrm{MR}(Y)$.*

Proof. By induction on $\mathrm{MR}(Y)$. The case that $\mathrm{MR}(Y) \in \{-\infty, 0, +\infty\}$ is obvious, so we can assume $\mathrm{MR}(Y) = \lambda \geq 1$, and make a further reduction to $\mathrm{MD}(Y) = 1$. Saturation gives an n such that $|f^{-1}(b)| \leq n$ for all $b \in Y$. Suppose that $\mathrm{MR}(X) > \lambda$. Take disjoint definable subsets $X_1, \dots, X_{n+1}$ of X with $\mathrm{MR}(X_i) \geq \lambda$ for all i. If $\mathrm{MR}\, f(X_i) < \lambda$ for some i, then by the induction hypothesis,

$$\mathrm{MR}(X_i) \leq \mathrm{MR}\, f(X_i) < \lambda$$

for such i, a contradiction. Hence $\mathrm{MR}\, f(X_i) = \lambda$ for all i. As $\mathrm{MD}(Y) = 1$, we have a $b \in f(X_1) \cap \cdots \cap f(X_{n+1})$, contradicting $|f^{-1}(b)| \leq n$. □

Corollary 10.2. *Let the sets X, Y, and $R \subseteq X \times Y$ be definable in $\mathbb{M}$ such that for each $a \in X$ the section $R(a)$ is finite and for each $b \in Y$ there is an $a \in X$ with $(a, b) \in R$. Then $\mathrm{MR}(X) \geq \mathrm{MR}(R) \geq \mathrm{MR}(Y)$.*

Proof. All fibers of the map $(a, b) \mapsto a : R \to X$ are finite, so $\mathrm{MR}(X) \geq \mathrm{MR}(R)$ by the previous lemma. The map $(a, b) \mapsto b : R \to Y$ is surjective, so $\mathrm{MR}(R) \geq \mathrm{MR}(Y)$. □

Exercise. Show that if $X \subseteq \mathbb{M}_x$, $Y \subseteq \mathbb{M}_y$ and $f : X \to Y$ are definable, and $\mathrm{MR}\left(f^{-1}(y)\right) \leq 1$ for all $y \in Y$, then $\mathrm{MR}(X) \leq \mathrm{MR}(Y) + 1$.

For definable $X, Y \subseteq \mathbb{M}_x$ we set

$$\begin{aligned} X =_\lambda Y \ &:\Longleftrightarrow\ \mathrm{MR}(X \triangle Y) < \lambda, \\ X \subseteq_\lambda Y \ &:\Longleftrightarrow\ \mathrm{MR}(X \setminus Y) < \lambda, \\ -X \ &:= \mathbb{M}_x \setminus X, \end{aligned}$$

or in terms of the boolean algebra $\mathrm{Def}_x(\mathbb{M})$ and its ideal $I := I_{<\lambda}$:

$$X =_\lambda Y \iff X =_I Y, \qquad X \subseteq_\lambda Y \iff X/I \leq Y/I.$$

Let X be a definable set in $\mathbb{M}$. Clearly, $\mathrm{MR}(X) = 0$ iff X is finite and nonempty. If $\mathrm{MR}(X) = 0$, then $\mathrm{MD}(X) = |X|$. The next possible "size" is when $\mathrm{MR}(X) = 1$, $\mathrm{MD}(X) = 1$, and such X is called *strongly minimal* (in $\mathbb{M}$), and this is equivalent to X being infinite such that each definable subset of X is finite or cofinite in X. We say that X is *A-irreducible* if X is A-definable and ranked, and X is not the union of two disjoint A-definable subsets of the same Morley rank. For example, if X is A-definable and ranked with $\mathrm{MD}(X) = 1$, then X is A-irreducible.

Exercise. Show that if X, Y and $R \subseteq X \times Y$ are definable sets in $\mathbb{M}$ and Y is strongly minimal, then there is an m such that for all $a \in X$, either $|R(a)| \leq m$ or $|Y \setminus R(a)| \leq m$.

Types over A and A-irreducible sets. Let $p(x)$ be an x-type over A. We shall identify $p(x)$ with the corresponding ultrafilter $\{\phi(\mathbb{M}_x) : \phi(x) \in p(x)\}$ of the boolean algebra $\mathrm{Def}_x(\mathbb{M}|A)$, so for each A-definable $X \subseteq \mathbb{M}_x$, either $X \in p$ or $-X \in p$.

If p does not contain any ranked set we put $\mathrm{MR}(p) = +\infty$. If p contains a ranked set, we choose $X \in p$ of minimal Morley rank λ and of minimal Morley degree d among the sets in p of Morley rank λ, and put $\mathrm{MR}(p) = \lambda$, $\mathrm{MD}(p) = d$. This set X is then A-irreducible, and

$$p = \{Y \in \mathrm{Def}_x(\mathbb{M}|A) : X \subseteq_\lambda Y\}.$$

Lemma 10.3. *Suppose $X \subseteq \mathbb{M}_x$ is ranked with $\mathrm{MR}(X) = \lambda$. Then*

1. *If X is A-irreducible, then $p := \{Y \in \mathrm{Def}_x(\mathbb{M}|A) : X \subseteq_\lambda Y\}$ is an x-type over A with $\mathrm{MR}(p) = \mathrm{MR}(X)$, $\mathrm{MD}(p) = \mathrm{MD}(X)$.*
2. *$\mathrm{MR}(X) = \max\{\mathrm{MR}(p) : p \in \mathrm{St}_x(A), X \in p\}$.*
3. *The set $\{p : p \in \mathrm{St}_x(A), X \in p, \mathrm{MR}(p) = \lambda\}$ is finite and*

$\mathrm{MD}(X) = \sum_p \mathrm{MD}(p)$ *where the sum is over the p in this finite set.*

Proof. For (1), use that if X is A-irreducible, then for each A-definable $Y \subseteq \mathbb{M}_y$, either $\mathrm{MR}(X \cap Y) < \lambda$ (in which case $X \subseteq_\lambda -Y$), or $\mathrm{MR}(X \setminus Y) < \lambda$ (in which case $X \subseteq_\lambda Y$). To obtain (2) and (3), decompose X as $X_1 \cup \cdots \cup X_n$ with $n \geq 1$ and disjoint A-irreducible $X_1, \ldots, X_n$ of Morley rank λ. Now apply (1) to each X_i. □

We say that an A-irreducible set $X \subseteq \mathbb{M}_x$ *determines* the type p defined in (1) of Lemma 10.3. Let $p(x) \in \mathrm{St}_x(A)$; then

$$\mathrm{MR}(p) = 0 \iff p(\mathbb{M}_x) \text{ is finite;}$$

we say that p is *algebraic* if $\mathrm{MR}(p) = 0$; in that case, $\mathrm{MD}(p) = |p(\mathbb{M}_x)|$.

For $a \in \mathbb{M}_x$ we define

$$\mathrm{MR}(a|A) := \mathrm{MR}\left(\mathrm{tp}(a|A)\right).$$

Thus $\mathrm{MR}(a|A) = 0 \iff a$ is A-algebraic. Note that if $\mathrm{MR}(a|A) = \lambda$, then there is an A-irreducible $X \subseteq \mathbb{M}_x$ such that $a \in X$ and $\mathrm{MR}(X) = \lambda$ and there is no A-definable $X \subseteq \mathbb{M}_x$ with $a \in X$ and $\mathrm{MR}(X) < \lambda$. Conversely, if $X \subseteq \mathbb{M}_x$ is A-irreducible, then there is an $a \in X$ such that $\mathrm{MR}(a|A) = \mathrm{MR}(X)$, and such an a is called a *generic element* of X over A (or just a *generic* of X over A).

From now on we let $a, a', a_1, a_2, b, b', c$ denote finite tuples in $\mathbb{M}$.

Lemma 10.4. *Suppose b is a-algebraic over A. Then* $\mathrm{MR}(b|A) \leq \mathrm{MR}(a|A)$.

Proof. Let $a \in \mathbb{M}_x$ and $b \in \mathbb{M}_y$, and take an A-definable $X \subseteq \mathbb{M}_x$ with $a \in X$ and $\mathrm{MR}(X) = \mathrm{MR}(a|A)$. Take an A-definable $R \subseteq \mathbb{M}_x \times \mathbb{M}_y$ such that $b \in R(a)$ and $R(a)$ is finite. By shrinking R we can arrange that $R \subseteq X \times \mathbb{M}_y$ and $|R(a')| \leq |R(a)|$ for all $a' \in X$. Let

$$Y := \{b' \in \mathbb{M}_y : R(a', b') \text{ for some } a' \in X\}.$$

Then Y is A-definable and $R \subseteq X \times Y$ satisfies the conditions of Corollary 10.2, so $\mathrm{MR}(b|A) \leq \mathrm{MR}(Y) \leq \mathrm{MR}(X) = \mathrm{MR}(a|A)$. $\square$

In particular, if a and b are interalgebraic over A, then $\mathrm{MR}(a|A) = \mathrm{MR}(b|A)$.

Definable types again. Following Ziegler we shall give another proof of definability of types in the case of totally transcendental theories.

Lemma 10.5. *Suppose the model M is $\aleph_0$-saturated, and $X \subseteq \mathbb{M}_x$ is definable and ranked, and contained in an M-definable subset of $\mathbb{M}_x$ of the same Morley rank. Then $X \cap M_x \neq \emptyset$.*

Proof. Let $Y \subseteq \mathbb{M}_x$ be M-definable with $\mathrm{MR}(X) = \mathrm{MR}(Y) = \lambda$. We proceed by induction on λ. If $\lambda = 0$, then Y is finite, so $X \subseteq Y \subseteq M_x$. Let $\lambda > 0$ and decompose Y as a disjoint union $Y_1 \cup \cdots \cup Y_d$ of M-definable sets of Morley rank λ and Morley degree 1. (This is possible by Lemma 9.1.) Then $\mathrm{MR}(X \cap Y_i) = \lambda$ for some i, so by replacing X and Y by $X \cap Y_i$ and Y_i for suitable i we have made a reduction to the case that $\mathrm{MD}(Y) = 1$. We can also assume that $X \neq Y$. So $\mathrm{MR}(Y \setminus X) = \alpha < \lambda$. Take disjoint M-definable sets $Y_0, Y_1, Y_2, \dots \subseteq Y$ of Morley rank $\geq \alpha$ and $< \lambda$. Not all Y_n can be contained in $Y \setminus X$, so we can take n such that $X \cap Y_n \neq \emptyset$. Since $\mathrm{MR}(Y_n) < \lambda$ we can assume inductively that $X \cap Y_n \cap M_x \neq \emptyset$. $\square$

Given M and an M-definable $X \subseteq \mathbb{M}_x$ we put

$$X(M) := X \cap M_x \quad \text{(the set of M-points of X).}$$

Theorem 10.6. *Let $\phi(x, y)$ be a stable L-formula and let $X \subseteq \mathbb{M}_x$ be definable with $\mathrm{MR}(X) = \lambda$. Then the subset*

$$\{b \in \mathbb{M}_y : X \subseteq_\lambda \phi(\mathbb{M}_x, b)\}$$

of $\mathbb{M}_y$ is definable.

Proof. By decomposing X into irreducible sets of Morley rank λ we can assume that $\mathrm{MD}(X) = 1$. Take an $\aleph_0$-saturated model M such that X is M-definable.

Claim 1. $X(M)$ has a finite subset Δ such that for all $b \in \mathbb{M}_y$,

$$\Delta \subseteq \phi(\mathbb{M}_x, b) \implies X \subseteq_\lambda \phi(\mathbb{M}_x, b).$$

To prove this claim, suppose there is no such Δ. Then we obtain a contradiction by constructing infinite sequences $a_1, a_2, \dots \in X(M)$ and $b_1, b_2, \dots \in \mathbb{M}_y$ such that for all i, j

$$\models \phi(a_i, b_j) \iff i \leq j.$$

We obtain such sequences as follows: Assume $a_1, \dots, a_n \in X(M)$ and $b_1, \dots, b_n \in \mathbb{M}_y$ are such that this equivalence holds for $i, j = 1, \dots, n$ and such that $X \not\subseteq_\lambda \phi(\mathbb{M}_x, b_i)$ for $i = 1, \dots, n$. Then

$$X \not\subseteq_\lambda \phi(\mathbb{M}_x, b_1) \cup \dots \cup \phi(\mathbb{M}_x, b_n),$$

so by the previous lemma we have an $a_{n+1} \in X(M)$ outside each $\phi(\mathbb{M}_x, b_i)$. Since $\{a_1, \dots, a_{n+1}\}$ cannot serve as Δ, we obtain $b_{n+1} \in \mathbb{M}_y$ such that $\{a_1, \dots, a_{n+1}\} \subseteq \phi(\mathbb{M}_x, b_{n+1})$ and $X \not\subseteq_\lambda \phi(\mathbb{M}_x, b_{n+1})$, so the displayed equivalence holds for $i, j = 1, \dots, n + 1$.

Claim 2. Suppose that $X \subseteq_\lambda \phi(\mathbb{M}_x, c)$, $c \in \mathbb{M}_y$. Then there exists a Δ as in Claim 1 such that $\Delta \subseteq \phi(\mathbb{M}_x, c)$.

The proof is a variant of that of Claim 1: take the $a_i \in \phi(\mathbb{M}_x, c)$.

For each Δ as in Claim 1, let $Y_\Delta := \{b \in \mathbb{M}_y : \Delta \subseteq \phi(\mathbb{M}_x, b)\}$, so $Y_\Delta \subseteq \mathbb{M}_y$ is definable. Claims 1 and 2 yield that

$$\{b \in \mathbb{M}_y : X \subseteq_\lambda \phi(\mathbb{M}_x, b)\} = \bigcup_\Delta Y_\Delta$$

where the union is over the small set of Δ's as in Claim 1. Since X is irreducible, the complement in $\mathbb{M}_y$ of the set on the left is

$$\{b \in \mathbb{M}_y : X \subseteq_\lambda \neg\phi(\mathbb{M}_x, b)\},$$

which is likewise a union of a small number of definable sets. Now use an exercise from Section 2. $\square$

In Section 5 we proved that if T is stable, then types over a model M are definable over M. Here we extend this to types over parameter sets, for totally transcendental T.

Corollary 10.7. *Assume T is totally transcendental, and let $p(x) \in \mathrm{St}_x(A)$. Then p is* definable *over A: for each L-formula $\phi(x,y)$ there is an L_A-formula $\psi(y)$ such that for all $b \in A_y$:*

$$\phi(x,b) \in p(x) \iff \models \psi(b).$$

Proof. Take an A-irreducible set $X \subseteq \mathbb{M}_x$ that determines p, so $\mathrm{MR}(X) = \mathrm{MR}(p) = \lambda$, say. Then for all $b \in A_y$,

$$\phi(x,b) \in p(x) \iff \phi(\mathbb{M}_x, b) \in p \iff X \subseteq_\lambda \phi(\mathbb{M}_x, b).$$

Now apply Theorem 10.6. □

Morley rank of global types. Let $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$ be a global type, viewed as an ultrafilter of $\mathrm{Def}_x(\mathbb{M})$. If $\mathbf{p}$ does not contain any ranked set, we put $\mathrm{MR}(\mathbf{p}) = +\infty$. Suppose that $\mathbf{p}$ contains a ranked set. Then we choose $X \in \mathbf{p}$ of minimal Morley rank λ, and put $\mathrm{MR}(\mathbf{p}) = \lambda$. Such an X can be chosen to have Morley degree 1, and then determines $\mathbf{p}$ in the sense that

$$\mathbf{p} = \{Y \in \mathrm{Def}_x(\mathbb{M}) : X \subseteq_\lambda Y\}.$$

Conversely, if $X \subseteq \mathbb{M}_x$ is definable, $\mathrm{MR}(X) = \lambda$, $\mathrm{MD}(X) = 1$, then X determines the global type $\mathbf{p} = \{Y \in \mathrm{Def}_x(\mathbb{M}) : X \subseteq_\lambda Y\}$ of Morley rank λ.

11. Forking and Independence

In this section we assume that T is totally transcendental.

Let $p \in \mathrm{St}_x(A)$, $A \subseteq B$ and suppose $q \in \mathrm{St}_x(B)$ extends p, that is, $p \subseteq q$, equivalently, $q \restriction A = p$. Then $\mathrm{MR}(p) \geq \mathrm{MR}(q)$, and in case of equality we say that q is a *nonforking extension* of p (to B). We also say that q *does not fork over* A if it is a nonforking extension of p; this terminology makes sense, since q determines p by $p = q \restriction A$.

For example, if p is algebraic, then q is necessarily a nonforking extension of p, and if $\mathrm{MR}(p) = 1$, then q is a nonforking extension of p iff q is not algebraic. We also say that a global type $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$ is a nonforking extension of p if $p \subseteq \mathbf{p}$ and $\mathrm{MR}(p) = \mathrm{MR}(\mathbf{p})$.

Lemma 11.1. *Let $p \in \mathrm{St}_x(A)$, $A \subseteq B$. Then p has a nonforking extension to B, and has at most $\mathrm{MD}(p)$ such extensions:*

$$\mathrm{MD}(p) = \sum_q \mathrm{MD}(q),$$

where the sum is over the nonforking extensions q of p to B.

Proof. Take an A-irreducible set X that determines p, so $\mathrm{MR}(X) = \mathrm{MR}(p)$ and $\mathrm{MD}(X) = \mathrm{MD}(p)$. Take disjoint B-irreducible sets $X_1, \dots, X_n$ of the same Morley rank as X such that $X = X_1 \cup \dots \cup X_n$, $n \geq 1$. Then each X_i determines a nonforking extension $p_i \in \mathrm{St}_x(B)$ of p, with $\mathrm{MD}(X_i) = \mathrm{MD}(p_i)$. Also each nonforking extension of p to B contains an X_i, and thus equals a p_i. It remains to note that $\mathrm{MD}(X) = \sum_i \mathrm{MD}(X_i)$. □

Almost the same proof, with "irreducible" in place of "B-irreducible", yields a similar result for *global* nonforking extensions:

Lemma 11.2. *Each $p \in \mathrm{St}_x(A)$ has exactly $\mathrm{MD}(p)$ global nonforking extensions in $\mathrm{St}_x(\mathbb{M})$.*

Corollary 11.3. *Let $p \in \mathrm{St}_x(A)$. Then $\mathrm{MD}(p) = 1$ if and only if p has exactly one nonforking extension to every $B \supseteq A$.*

Proof. The forward direction is immediate from Lemma 11.1. For the converse, suppose $n := \mathrm{MD}(p) > 1$. Take an A-irreducible set X that determines p, so $\mathrm{MD}(X) = n$. Then X is a disjoint union of definable subsets $X_1, \dots, X_n$ of the same Morley rank. Take a $B \supseteq A$ such that $X_1, \dots, X_n$ are B-definable, hence B-irreducible. Then $X_1, \dots, X_n$ determine distinct types $q_1, \dots, q_n \in \mathrm{St}_x(B)$, and all are nonforking extensions of p. □

A type $p \in \mathrm{St}_x(A)$ with $\mathrm{MD}(p) = 1$ is also said to be *stationary*, because it has the property expressed in this corollary. Note that if $p \in \mathrm{St}_x(A)$ is stationary, then p has also a unique global nonforking extension $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$.

It is obvious that nonforking is *monotone* and *transitive*:

Lemma 11.4. *Let $A \subseteq B \subseteq C$ and $p \in \mathrm{St}_x(A), q \in \mathrm{St}_x(B), r \in \mathrm{St}_x(C)$ be such that $p \subseteq q \subseteq r$. Then r is a nonforking extension of p if and only if r is a nonforking extension of q and q is a nonforking extension of p.*

Nonforking is also *continuous*:

Lemma 11.5. *Let $q \in \mathrm{St}_x(B)$. Then*

(1) *there is a finite $B_0 \subseteq B$ such that q does not fork over B_0 and q is the only nonforking extension of $q \restriction B_0$ to B;*
(2) *if $A \subseteq B$ and q forks over A, then there is a finite $B_0 \subseteq B$ such that $q \restriction (A \cup B_0)$ forks over A.*

Proof. Take a set $X \in q$ with $\mathrm{MR}(X) = \mathrm{MR}(q)$, and $\mathrm{MD}(X) = \mathrm{MD}(q)$. Take a finite $B_0 \subseteq B$ such that X is B_0-definable. Then B_0 witnesses (1) by Lemma 11.1. For (2), assume $A \subseteq B$ and q forks over A, and put $p := q \restriction A$. Then we can take $X \in q$ such that $\mathrm{MR}(X) = \mathrm{MR}(q) < \mathrm{MR}(p)$. Take a finite $B_0 \subseteq B$ such that X is $(A \cup B_0)$-definable. Then B_0 witnesses the conclusion of (2). □

We say that a is *independent* from B over A, notation:

$$a \underset{A}{\downarrow} B,$$

if $\mathrm{MR}(a|AB) = \mathrm{MR}(a|A)$, in other words, $\mathrm{tp}(a|AB)$ does not fork over A. In this notation, Lemma 10.4 and the lemmas above in this section yield:

$$a \text{ is } A\text{-algebraic} \implies a \underset{A}{\downarrow} B,$$

$$\text{if } a \text{ and } b \text{ are interalgebraic over } A, \text{ then: } \; a \underset{A}{\downarrow} B \;\Leftrightarrow\; b \underset{A}{\downarrow} B,$$

$$p \in \mathrm{St}_x(A) \implies p \text{ has a realization } a \text{ such that } a \underset{A}{\perp} B,$$

$$a \underset{A}{\perp} B \text{ and } B' \subseteq B \implies a \underset{A}{\perp} B',$$

$$\big(a \underset{A}{\perp} B \text{ and } a \underset{AB}{\perp} C\big) \iff a \underset{A}{\perp} BC,$$

$$a \underset{A}{\perp} B \iff a \underset{A}{\perp} b \text{ for each finite } B\text{-tuple } b.$$

We shall use these rules freely from now on. Note the special case

$$a \underset{A}{\perp} BC \implies a \underset{AB}{\perp} C.$$

Symmetry. In the next symmetry lemma, a and b also stand for certain parameter sets, according to the convention in Section 2.

Lemma 11.6. $a \underset{A}{\perp} b \iff b \underset{A}{\perp} a.$

Proof. We first assume that $A = M$ is an $\aleph_0$-saturated model. Set $\alpha := \mathrm{MR}(a|M)$ and $\beta := \mathrm{MR}(b|M)$. Let $a \in \mathbb{M}_x$ and $b \in \mathbb{M}_y$ with disjoint x, y, and take M-definable $X \subseteq \mathbb{M}_x$ and $Y \subseteq \mathbb{M}_y$ that determine $\mathrm{tp}(a|M)$ and $\mathrm{tp}(b|M)$. Suppose that b is not independent from a over M. Then

$$\mathrm{MR}(b|Ma) < \mathrm{MR}(b|M) = \beta,$$

so we have an M-definable $Z \subseteq \mathbb{M}_x \times \mathbb{M}_y$ such that

$$(a, b) \in Z, \qquad \mathrm{MR}\big(Z(a)\big) < \beta.$$

We can assume that $Z \subseteq X \times Y$. We have

$$\{a' \in X : \mathrm{MR}\big(Z(a')\big) < \beta\} = \{a' \in X : Y \not\subseteq_\beta Z(a')\},$$

so by Theorem 10.6, the set $X' := \{a' \in X : \mathrm{MR}\big(Z(a')\big) < \beta\}$ is M-definable, and contains a. After replacing Z by $Z \cap (X' \times Y)$ we can assume that $\mathrm{MR}\big(Z(a')\big) < \beta$ for all $a' \in X$. Hence $\breve{Z}(b) \cap M_x = \emptyset$, so by Lemma 10.1 we have $\mathrm{MR}\big(\breve{Z}(b)\big) < \mathrm{MR}(X) = \alpha$. So a is not independent from b over A.

We now consider any A, and take an $\aleph_0$-saturated model $M \supseteq A$. Take a $b' \in \mathbb{M}_y$ such that

$$b' \text{ realizes } \mathrm{tp}(b|A), \qquad b' \underset{A}{\perp} M.$$

Take $f \in \mathrm{Aut}(\mathbb{M}|A)$ with $fb = b'$, and take an $a' \in \mathbb{M}_x$ such that

$$a' \text{ realizes } \mathrm{tp}(fa|Ab'), \qquad a' \underset{Ab'}{\perp} Mb',$$

and take $g \in \mathrm{Aut}(\mathbb{M}|Ab')$ with $g(fa) = a'$. Then $h := gf \in \mathrm{Aut}(\mathbb{M}|A)$ satisfies $h(a) = a'$ and $h(b) = b'$, so

$$a \underset{A}{\perp} b \iff a' \underset{A}{\perp} b', \qquad b \underset{A}{\perp} a \iff b' \underset{A}{\perp} a',$$

Suppose now that a is independent from b over A. Then $a' \underset{A}{\downarrow} b'$, so $a' \underset{A}{\downarrow} Mb'$ (since $a' \underset{Ab'}{\downarrow} Mb'$), so $a' \underset{M}{\downarrow} b'$. By the first part of the proof, this gives $b' \underset{M}{\downarrow} a'$, which in view of $b' \underset{A}{\downarrow} M$ yields $b' \underset{A}{\downarrow} a'$, so b is independent from a over A. □

Corollary 11.7. $a \underset{A}{\downarrow} \operatorname{acl}(A)$.

Proof. For A-algebraic b we have $b \underset{A}{\downarrow} a$, hence $a \underset{A}{\downarrow} b$. □

The next result is needed in dealing with one-basedness in Section 13.

Corollary 11.8. *Suppose* $a \underset{A}{\downarrow} b$. *Then* $\operatorname{acl}(Aa) \cap \operatorname{acl}(Ab) = \operatorname{acl}(A)$.

Proof. Let c be algebraic over Aa and over Ab. Then

$$\operatorname{MR}(a|Abc) = \operatorname{MR}(a|Ab) = \operatorname{MR}(a|A),$$

and thus $\operatorname{MR}(a|Ac) = \operatorname{MR}(a|A)$. Hence $a \underset{A}{\downarrow} c$, so $c \underset{A}{\downarrow} a$, that is, $\operatorname{MR}(c|A) = \operatorname{MR}(c|Aa) = 0$, so c is algebraic over A. □

Lemma 11.9. *Let* $p \in \operatorname{St}_x(A)$. *Then*

(1) *each* $q \supseteq p$ *in* $\operatorname{St}_x(\operatorname{acl}(A))$ *is a nonforking extension of* p*;*
(2) *any two extensions of* p *in* $\operatorname{St}_x(\operatorname{acl}(A))$ *are conjugate over* A.

Proof. Item (1) is immediate from Corollary 11.7 by taking a realization $a \in \mathbb{M}_x$ of q. For (2) we take an A-irreducible X that determines p, and take an $\operatorname{acl}(A)$-irreducible $Y \subseteq X$ with $\operatorname{MR}(X) = \operatorname{MR}(Y)$. Take an A-algebraic b such that Y is b-definable. Then clearly $\{f(b) : f \in \operatorname{Aut}(\mathbb{M}|A)\}$ is finite, and whenever $f(b) = g(b)$, then $f(Y) = g(Y)$, for $f, g \in \operatorname{Aut}(\mathbb{M}|A)$, so Y has only finitely many conjugates over A. Let $Y_1, \dots, Y_n$ be the distinct conjugates of Y over A. Then $Y_1 \cup \dots \cup Y_n \subseteq X$ is A-definable by Corollary 5.2, so $X =_\lambda Y_1 \cup \dots \cup Y_n$ where $\lambda = \operatorname{MR}(X)$. The sets $Y_1, \dots, Y_n$ determine extensions $q_1, \dots, q_n \in \operatorname{St}_x(\operatorname{acl}(A))$ of p, and $q_1, \dots, q_n$ are conjugate over A the same way $Y_1, \dots, Y_n$ are.

Suppose that $q \in \operatorname{St}_x(\operatorname{acl}(A))$ extends p. Since $Y_1 \cup \dots \cup Y_n \in p$ we have an $i \in \{1, \dots, n\}$ with $Y_i \in q$. Since $\operatorname{MR}(q) = \lambda$ by (1), this Y_i determines q, so $q = q_i$. □

Theorem 11.10. *Suppose* T *has* EI *and* A *is algebraically closed. Let* $A \subseteq M$*, and let* $p, p_1, p_2 \in \operatorname{St}_x(M)$. *Then*

(1) *if* p *does not fork over* A*, then* p *is definable over* A*;*
(2) *if* p_1, p_2 *are definable over* A*, and* $p_1 \restriction A = p_2 \restriction A$*, then* $p_1 = p_2$*;*
(3) *each* x*-type over* A *is stationary.*

Proof. For (1), assume p does not fork over A. Take a global nonforking extension $\mathbf{p} \in \operatorname{St}_x(\mathbb{M})$ of p. The conjugates $f(\mathbf{p})$ with $f \in \operatorname{Aut}(\mathbb{M}|A)$ are all nonforking extensions of $p \restriction A$, so there are only finitely many such conjugates. Let $\phi(x,y)$ be an L-formula, and $Y := d_{\mathbf{p}}x\phi(x,\mathbb{M}_y)$. Then by the above, Y has only finitely many conjugates over A, so Y is A-algebraic by Corollary 5.2, hence Y is A-definable by Corollary 5.3.

For (2), let p_1, p_2 be definable over A, and $p_1 \restriction A = p_2 \restriction A$. Let $\phi(x,y)$ be an L-formula, and let $\psi_1(y)$ and $\psi_2(y)$ be L_A-formulas that define $p_1 \restriction \phi$ and $p_2 \restriction \phi$. It is enough to show that then $\psi_1(M_y) = \psi_2(M_y)$. Let $b \in \psi_1(M_y)$, and let $q \in \mathrm{St}_y(M)$ be a nonforking extension of $\mathrm{tp}(b|A)$. Then q is definable over A by (1), so we have an L_A-formula $\chi(x)$ defining $q \restriction \breve{\phi}$. Since $\psi_1(y) \in q(y)$ we can apply Lemma 8.8 to $p_1 \restriction \phi$ and $q \restriction \breve{\phi}$ to get $\chi(x) \in p_1(x)$, hence $\chi(x) \in p_2(x)$. Now apply the same lemma to $p_2 \restriction \phi$ and $q \restriction \breve{\phi}$ to get $\psi_2(y) \in q$, that is, $b \in \psi_2(M_y)$.

Item (3) follows from (1) and (2). $\square$

In the next three corollaries we do not assume that T has EI.

Corollary 11.11. *Every x-type over a model M is stationary.*

Proof. Each $p \in \mathrm{St}_x(M)$ extends uniquely to a type $p^{\mathrm{eq}} \in \mathrm{St}_x(M^{\mathrm{eq}})$ in $\mathbb{M}^{\mathrm{eq}}$. Since T^{eq} has EI, and M^{eq} is algebraically closed in $\mathbb{M}^{\mathrm{eq}}$, the type p^{eq} is stationary, that is, $\mathrm{MD}(p^{\mathrm{eq}}) = 1$. But $\mathrm{Def}_x(\mathbb{M}|M)$ and $\mathrm{Def}_x(\mathbb{M}^{\mathrm{eq}}|M^{\mathrm{eq}})$ are the same boolean algebra, and, viewed as ultrafilters on these boolean algebras, p and p^{eq} are also the same, so $\mathrm{MD}(p) = 1$, so p is stationary. $\square$

Exercises. Show that if $X \subseteq \mathbb{M}_x$ is M-definable, $\mathrm{MR}(X) = \lambda$, $\mathrm{MD}(X) = d$, then there are M-definable disjoint $X_1, \dots, X_d \subseteq \mathbb{M}_x$ such that

$$X = X_1 \cup \dots \cup X_k, \text{ and } \mathrm{MR}(X_i) = \lambda, \ \mathrm{MD}(X_i) = 1 \text{ for } i = 1, \dots, d.$$

Show that for $a, b \in \mathbb{M}_x$, $\mathrm{stp}(a|A) = \mathrm{stp}(b|A)$ if and only if there is a model $M \supseteq A$ such that $\mathrm{tp}(a|M) = \mathrm{tp}(b|M)$.

Elaborating on the proof of Corollary 11.11 we note that each x-type over A in $\mathbb{M}$, as an ultrafilter of the boolean algebra $\mathrm{Def}_x(\mathbb{M}|A)$ is also an x-type over A in $\mathbb{M}^{\mathrm{eq}}$ since $\mathrm{Def}_x(\mathbb{M}|A)$ and $\mathrm{Def}_x(\mathbb{M}^{\mathrm{eq}}|A)$ are the same boolean algebra. In view of $\mathrm{Aut}(\mathbb{M}|A) = \mathrm{Aut}(\mathbb{M}^{\mathrm{eq}}|A)$ (see Section 3), this yields

Corollary 11.12. *Any two global nonforking extensions of $p \in \mathrm{St}_x(A)$ are conjugate over A.*

Proof. By the observation preceding the statement of this corollary it suffices to prove this for types in $\mathbb{M}^{\mathrm{eq}}$ instead of $\mathbb{M}$, that is, we can assume that T has EI. Let $\mathbf{p}_1$ and $\mathbf{p}_2$ be nonforking extensions of $p \in \mathrm{St}_x(A)$, and let $p_i := \mathbf{p}_i \restriction \mathrm{acl}(A)$ for $i = 1, 2$. By Lemma 11.9, (2), we have an $f \in \mathrm{Aut}(\mathbb{M}|A)$ such that $f(p_1) = p_2$. Then by Theorem 11.10, (3), we have $f(\mathbf{p}_1) = \mathbf{p}_2$. $\square$

Corollary 11.13. *Suppose $M \subseteq B$ and $q \in \mathrm{St}_x(B)$ does not fork over M. Then q is definable over M.*

Proof. As subalgebras of the boolean algebra $\mathrm{Def}_x(\mathbb{M}) = \mathrm{Def}_x(\mathbb{M}^{\mathrm{eq}})$ we have

$$\mathrm{Def}_x(\mathbb{M}|B) = \mathrm{Def}_x(\mathbb{M}^{\mathrm{eq}}|B), \quad \mathrm{Def}_x(\mathbb{M}|M) = \mathrm{Def}_x(\mathbb{M}^{\mathrm{eq}}|M^{\mathrm{eq}}).$$

Viewing in this way q as an x-type over B in $\mathbb{M}^{\text{eq}}$, Theorem 11.10, (1), yields that q is definable over M^{eq}. It follows that q is definable over M. $\square$

Corollary 11.14. *Suppose T has* EI*, and $\mathbf{p} \in \text{St}_x(\mathbb{M})$. Then*

(1) *$\mathbf{p}$ does not fork over A if and only if $\text{cb}(\mathbf{p}) \subseteq \text{acl}(A)$;*
(2) *$\mathbf{p}$ does not fork over A and $\mathbf{p} \restriction A$ is stationary if and only if $\text{cb}(\mathbf{p}) \subseteq \text{dcl}(A)$.*

Proof. Assume $\text{cb}(\mathbf{p}) \subseteq \text{acl}(A)$. Then $\mathbf{p}$ is defined over $\text{acl}(A)$. Take a nonforking extension $\mathbf{q} \in \text{St}_x(\mathbb{M})$ of $\mathbf{p} \restriction \text{acl}(A)$. Then $\mathbf{q}$ is also defined over $\text{acl}(A)$ by Theorem 11.10, (1), so $\mathbf{p} = \mathbf{q}$ by part (2) of that theorem, so $\mathbf{p}$ does not fork over $\text{acl}(A)$, and therefore does not fork over A by Lemma 11.9. Conversely, if $\mathbf{p}$ does not fork over A, then $\mathbf{p}$ does not fork over $\text{acl}(A)$, hence is definable over $\text{acl}(A)$ by Theorem 11.10, and thus $\text{cb}(\mathbf{p}) \subseteq \text{acl}(A)$. This proves (1).

For (2), if $\text{cb}(\mathbf{p}) \subseteq \text{dcl}(A)$, then $\mathbf{p}$ does not fork over A by (1), and has no conjugates over A different from itself, hence $\mathbf{p} \restriction A$ is stationary. Conversely, if $\mathbf{p}$ does not fork over A and $\mathbf{p} \restriction A$ is stationary, then $f(\mathbf{p}) = \mathbf{p}$ for all $f \in \text{Aut}(\mathbb{M}|A)$, and thus $\text{cb}(\mathbf{p}) \subseteq \text{dcl}(A)$. $\square$

Corollary 11.15. *Suppose T has* EI*, and $\mathbf{p} \in \text{St}_x(\mathbb{M})$. Then $\text{cb}(\mathbf{p}) = \text{dcl}(a)$ for some a.*

Proof. Let $A := \text{cb}(\mathbf{p})$. Then by (2) of the previous corollary, $\mathbf{p}$ is the unique nonforking extension of $p := \mathbf{p} \restriction A$, and $\text{MD}(p) = 1$. Take an irreducible A-definable $X \subseteq \mathbb{M}_x$ that determines p, and let a be a code of X. Then X is a-definable, so $X \in \mathbf{p} \restriction a$, so $\mathbf{p} \restriction a$ is stationary, and $\mathbf{p}$ is a nonforking extension of $\mathbf{p} \restriction a$. Hence by (2) of the previous corollary we have $A \subseteq \text{dcl}(a)$. But a is A-definable, so $\text{dcl}(a) = A$. $\square$

Suppose T has EI, and $p \in \text{St}_x(A)$ is stationary. Then we put

$$\text{cb}(p) := \text{cb}(\mathbf{p}), \text{where } \mathbf{p} \in \text{St}_x(\mathbb{M}) \text{ is the nonforking extension of } p.$$

Also, given $a \in \mathbb{M}_x$ such that $\text{tp}(a|A)$ is stationary, we put

$$\text{cb}(a|A) := \text{cb}(\text{tp}(a|A)).$$

Theorem 11.16. *Let $p \in \text{St}_x(M)$, $M \subseteq B$ and $p \subseteq q \in \text{St}_x(B)$. Then*

$$q \text{ is a nonforking extension of } p \iff q \text{ is a heir of } p.$$

Proof. Let q be the nonforking extension of p to B. Then q is definable over M by Corollary 11.13. Let $\phi(x,y)$ be an L_M-formula, $b \in B_y$, and $\phi(x,b) \in q$. Take an L_M-formula $\psi(y)$ that defines $q \restriction \phi$. Then $\models \psi(b)$, which in view of $M \preceq \mathbb{M}$ gives $a \in M_y$ such that $\models \psi(a)$, so $\phi(x,a) \in q$, and thus $\phi(x,a) \in p$. This proves $\Longrightarrow$. For $\Longleftarrow$, note that p is definable, so p has exactly one heir over B. $\square$

Morley sequences. This material will be needed in Section 13. Given a stationary $p \in \text{St}_x(A)$, a Morley sequence in p is a sequence $(a_i)_{i \in \mathbb{N}}$ in $\mathbb{M}_x$

such that a_i realizes the unique nonforking extension of p to $Aa_0 \dots a_{i-1}$ for each i. (In particular, all a_i realize p.) Note that there exists a Morley sequence in p.

Lemma 11.17. *Let $p \in \mathrm{St}_x(A)$ be stationary, and let (a_i) be a Morley sequence in p. Then*

(1) $a_i \underset{A}{\downarrow} (a_0, \dots, a_{i-1})$ *for all i;*
(2) *the sequence (a_i) is indiscernible over A;*
(3) *if (b_i) is also a Morley sequence in p, then there is an $f \in \mathrm{Aut}(\mathbb{M}|A)$ such that $f(a_i) = b_i$ for all i;*
(4) *suppose T has* EI*; then* $\mathrm{cb}(p) \subseteq \mathrm{dcl}(\{a_i : i \in \mathbb{N}\})$*;*

Proof. Item (1) is evident from the definition of Morley sequence. For (2) it is enough, by Lemma 8.9, to show that for all $i_0 < \dots < i_n$,

$$\mathrm{tp}((a_0, \dots, a_n)|A) = \mathrm{tp}((a_{i_0}, \dots, a_{i_n})|A).$$

We prove this by induction on n. The case $n = 0$ is clear. Let $i_0 < \dots < i_n < i_{n+1}$. Now a_{n+1} realizes $p \upharpoonright Aa_0 \dots a_n$ and $a_{i_{n+1}}$ realizes $p \upharpoonright Aa_0 \dots a_N$ where $N = i_{n+1} - 1$, hence $a_{i_{n+1}}$ realizes $p \upharpoonright Aa_{i_0} \dots a_{i_n}$. Assuming inductively the displayed identity above, it follows that

$$\mathrm{tp}((a_0, \dots, a_n, a_{n+1})|A) = \mathrm{tp}((a_{i_0}, \dots, a_{i_n}, a_{i_{n+1}})|A).$$

For (3), let (b_i) be a Morley sequence in p. Induction on n as in the proof of (2) yields $\mathrm{tp}((a_0, \dots, a_n)|A) = \mathrm{tp}((b_0, \dots, b_n)|A)$, and (3) follows. To prove (4), let $\mathbf{p}$ be the global nonforking extension of p. Consider an L-formula $\delta(x, y)$. The proof of Lemma 8.4 shows that $\mathbf{p} \upharpoonright \delta$ is defined over (b_i) for some Morley sequence (b_i) in p. But $\mathbf{p}$ and thus $\mathbf{p} \upharpoonright \delta$ is also defined over A. Then by (3) $\mathbf{p} \upharpoonright \delta$ is defined over (a_i). Since δ is arbitrary, it follows that $\mathbf{p}$, and thus p is defined over $\{a_i : i \in \mathbb{N}\}$. ∎

There is a problem with this proof since we only defined $p \upharpoonright B$ for types p over models. For stationary p it should be defined in general.

Independence of parameter sets. For use in Section 13 we extend our notion of independence from finite tuples to parameter sets: *A is independent from B over C*, notation: $A \underset{C}{\downarrow} B$, if for each A-tuple a we have $a \underset{C}{\downarrow} B$. It is easily checked that if A and B are the parameter sets corresponding to tuples a and b, then $A \underset{C}{\downarrow} B$ is equivalent to $a \underset{C}{\downarrow} b$ as defined previously. Thus our notation for independence of parameter sets agrees with the convention of letting a tuple in $\mathbb{M}$ stand for the corresponding parameter set when it suits us.

12. Combinatorial Geometries and Strongly Minimal Sets

Strongly minimal sets are basic building blocks of totally transcendental structures. As we shall see, some very robust features of a strongly minimal set are in turn controlled by something more primitive: a pregeometry. We begin with discussing pregeometries. As the name suggests, a pregeometry can be turned into a geometry.

Pregeometries. A *pregeometry* is a set Ω with an operation

$$\mathrm{cl}: \ \mathcal{P}(\Omega) \longrightarrow \mathcal{P}(\Omega)$$

such that for all $E \subseteq \Omega$ and $a, b \in \Omega$:

(1) $E \subseteq \mathrm{cl}\, E$;
(2) $\mathrm{cl}\, E = \bigcup\{\mathrm{cl}\, F : \ F \text{ a finite subset of } E\}$;
(3) $\mathrm{cl}(\mathrm{cl}\, E) = \mathrm{cl}\, E$;
(4) $a \in \mathrm{cl}(E \cup \{b\}), a \notin \mathrm{cl}\, E \implies b \in \mathrm{cl}(E \cup \{a\})$.

In most cases (1) and (2) are trivially satisfied, and (3) and (4) may require a little work. Condition (4) is called the *Steinitz Exchange Axiom* and is the most significant of the four conditions.

Examples.

(1) Let Ω be any set. For $E \subseteq \Omega$, let $\mathrm{cl}\, E := E$. This makes Ω into a (rather trivial) pregeometry.
(2) Let V be a (left) vector space over a division ring $\mathbf{k}$. For $E \subseteq V$, put $\mathrm{cl}\, E :=$ $\mathbf{k}$-linear span of E. This makes V into a pregeometry.
(3) Let K be a field. For $E \subseteq K$, let $\mathrm{cl}\, E$ be the set of all $a \in K$ that are algebraic over the subfield of K generated by E. This makes K into an pregeometry.

Let Ω be a pregeometry as above, and $E \subseteq \Omega$. Condition (2) yields

$$E \subseteq E' \subseteq \Omega \implies \mathrm{cl}\, E \subseteq \mathrm{cl}\, E'.$$

We call E *closed* if $\mathrm{cl}\, E = E$. The intersection (inside Ω) of any collection of closed subsets of Ω is also closed, and $\mathrm{cl}\, E$ is the smallest closed subset of Ω that contains E.

We say that E is *independent* if $a \notin \mathrm{cl}(E \setminus \{a\})$ for all $a \in E$, that E *generates* or *spans* Ω if $\mathrm{cl}\, E = \Omega$, and that E is a *basis* of Ω if E is both independent and generates Ω. It is easy to show that the following three conditions are equivalent:

(1) E is a maximal independent subset of Ω;
(2) E is a minimal generating set of Ω;
(3) E is a basis of Ω.

For example, to show that (2) $\implies$ (3), prove first that if E spans Ω and $F \subseteq E$ is independent, then Ω has a basis B such that $F \subseteq B \subseteq E$. (Use Zorn.) Special cases: if F is independent, then it is a subset of a basis of Ω, and if E spans Ω, then E has a subset that is a basis of Ω. In particular, Ω has a basis. We leave the proof of the following exchange lemma as an exercise.

Lemma 12.1. *Let E and F be bases of Ω and $e \in E \setminus F$. Then there is an $f \in F$ such that $(E \setminus \{e\}) \cup \{f\}$ is also a basis of Ω.*

Now a key result:

Proposition 12.2. *All bases of Ω have the same size.*

Proof. Let E and F be bases of Ω. Consider first the case that F is finite. We claim that then E is finite. To see why, note that each $f \in F$ lies in $\mathrm{cl}(E_f)$ for some finite $E_f \subseteq E$, so $F \subseteq \mathrm{cl}(E')$ with finite $E' \subseteq E$, so $E \subseteq \Omega = \mathrm{cl}(F) \subseteq \mathrm{cl}(E')$. Since E is independent, this gives $E = E'$, so E is finite. With E and F both finite, one uses the exchange lemma above to obtain $|E| = |F|$. It remains to consider the case that E and F are both infinite. Then we argue as before: for each $f \in F$ we take a finite $E_f \subseteq E$ such that $f \in \mathrm{cl}(E_f)$. Then

$$F \subseteq \bigcup_{f\in F} \mathrm{cl}\, E_f \subseteq \mathrm{cl}\,\Big(\bigcup_{f\in F} E_f\Big) = \mathrm{cl}\, E', \qquad E' := \bigcup_{f\in F} E_f,$$

so $E \subseteq \Omega = \mathrm{cl}\, F \subseteq \mathrm{cl}\, E'$ with $E' \subseteq E$, so $E' = E$. Hence $|E| = |E'| \le |F| \cdot \aleph_0 = |F|$. Likewise, $|F| \le |E|$. $\square$

The *rank* of the pregeometry, denoted by $\mathrm{rk}\,\Omega$, is the size of any basis of Ω. Some people call it the *dimension* of the pregeometry, but in some cases, like projective geometry, this conflicts with more natural notions of dimension, so the neutral term "rank" is to be preferred.

Note that in example (1) the set Ω is itself a basis, so $\mathrm{rk}\,\Omega = |\Omega|$, and in example (2), the notions of *spanning set*, *independent set*, and *basis* are the familiar ones in vector spaces, and thus $\mathrm{rk}\,\Omega = \dim_{\mathbf{k}} \Omega$. In example (3), *independent* means *algebraically independent*, so $\mathrm{rk}\,\Omega$ is the transcendence degree of the field Ω over its prime field.

Let $X \subseteq \Omega$. Then we consider X as a pregeometry with respect to the closure operation $E \mapsto \mathrm{cl}(E) \cap X : \mathcal{P}(X) \to \mathcal{P}(X)$. Note that a set $E \subseteq X$ is independent in the pregeometry X if and only if E is independent in Ω, so, with harmless ambiguity:

$$\mathrm{rk}\, X \;=\; \text{size of any maximal independent subset of } X.$$

We also have the pregeometry $\Omega|X$ ("Omega over X") which has Ω as its underlying set, and closure operation

$$E \mapsto \mathrm{cl}(E \cup X) : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega).$$

A set $F \subseteq \Omega$ that is independent in the pregeometry $\Omega|X$ is also said to be *independent over* X or *X-independent*. If E is a basis of X and F a basis of $\Omega|X$, then $E \cap F = \emptyset$, and $E \cup F$ is a basis of Ω. Thus

$$\mathrm{rk}\,\Omega \;=\; \mathrm{rk}\, X + \mathrm{rk}\,\Omega|X \qquad \text{(additivity of rank)}$$

We also define $\mathrm{rk}(E|X)$ to be the rank of the pregeometry $(E \cup X)|X$. Note that any maximal X-independent subset of E is a basis of $(E \cup X)|X$, so $\mathrm{rk}(E|X)$ is the size of any maximal X-independent subset of E. The pregeometry $\Omega|X$ is also referred to as the *localization of* Ω *at* X.

From a pregeometry to a geometry. A (combinatorial) geometry is a pregeometry Ω such that $\mathrm{cl}\,\emptyset = \emptyset$ and $\mathrm{cl}\{a\} = \{a\}$ for each $a \in \Omega$. The pregeometry on a set Ω with $\mathrm{cl}\,E = E$ for each $E \subseteq \Omega$ is a geometry. But the pregeometry of a vector space over a division ring is *not* a geometry, since it has $\mathrm{cl}\,\emptyset = \{0\}$.

Let Ω be a pregeometry. By a *line of* Ω we mean a set $\mathrm{cl}\{a\}$ with $a \in \Omega \setminus \mathrm{cl}\,\emptyset$. Note that any two distinct lines of Ω have the trivial intersection $\mathrm{cl}\,\emptyset$. The sets $E, F \subseteq \Omega$ are said to *intersect nontrivially* if their intersection contains an element outside $\mathrm{cl}\,\emptyset$.

To the pregeometry Ω we associate a geometry Ω' with closure operation cl' as follows: The points (elements) of Ω' are the lines of Ω; for $X \subseteq \Omega'$, the union $\cup X$ is the set of all elements of Ω that lie on some line $p \in X$, and we define $\mathrm{cl}'\,X \subseteq \Omega'$ to be the set of all lines of Ω that are contained in $\mathrm{cl}\cup X$.

Let $E \subseteq \Omega$, and put $E' := \{\mathrm{cl}\{a\} : a \in E \setminus \mathrm{cl}\,\emptyset\}$, that is, $E' \subseteq \Omega'$ is the set of lines of Ω that intersect E nontrivially. It is easy to see that $\mathrm{cl}'\,E'$ is the set of lines of Ω that intersect $\mathrm{cl}\,E$ nontrivially. Note that we have a bijection $E \mapsto E'$ from the set of closed subsets E of Ω onto the set of closed subsets of Ω'. Note also that if E is independent, then E' is independent. In particular, $\mathrm{rk}\,E = \mathrm{rk}\,E'$ for all E.

Let Ω be the pregeometry of a vector space over a division ring $\mathbf{k}$. Then the points of Ω' are the lines $\mathbf{k}a$ ($a \in \Omega$, $a \neq 0$), that is $\Omega' = \mathbb{P}(\Omega)$, the projective space associated to the vector space Ω. Note also that if $E \subseteq \Omega$ is closed, that is, E is a $\mathbf{k}$-linear subspace of Ω, then $E' = \mathbb{P}(E) := \{\mathbf{k}a : 0 \neq a \in E\}$. Thus our construction of a geometry from a pregeometry generalizes the construction of the projective space associated to a vector space. We call Ω' the *projective geometry associated to the vector space* Ω. Note that if Ω as a vector space over $\mathbf{k}$ has dimension n, then $\mathrm{rk}\,\Omega = n = \mathrm{rk}\,\Omega'$, but the projective space $\mathbb{P}(\Omega)$ is regarded as a space of *dimension* $n-1$ for $n > 0$. So our combinatorially defined rank does not always agree with more geometrically inspired notions of dimension.

We also associate to a vector space V over a division ring $\mathbf{k}$ another combinatorial geometry, namely its *affine geometry*: define a *flat* to be either a translate $a + E$ of a linear subspace E or the empty subset of V. Then the flats are the closed sets of the affine geometry of V, which has V as its underlying set. Localizing this geometry at $\{0\}$ we get back the pregeometry of the vector space V.

The pregeometry on a strongly minimal set. We now return to the setting of our monster model $\mathbb{M}$ of T, but we do not assume that T is totally transcendental unless we say so.

Let $\Omega \subseteq \mathbb{M}_x$ be infinite and type-definable over A, such that for each definable $X \subseteq \mathbb{M}_x$, either $\Omega \cap X$ is finite, or $\Omega \setminus X$ is finite. (For A-definable $\Omega \subseteq \mathbb{M}_x$ this just says that Ω is strongly minimal, but there are situations of interest where the weaker assumption is relevant.)

Then we have the following exchange lemma:

Lemma 12.3. *Let $a \in \Omega$, $b \in \mathbb{M}_y$, and suppose b is Aa-algebraic, but not A-algebraic. Then a is Ab-algebraic.*

Proof. Take an A-definable relation $R \subseteq \mathbb{M}_x \times \mathbb{M}_y$ such that $R(a)$ is finite, say $|R(a)| \le n$, and $b \in R(a)$. By shrinking R we can arrange that $|R(a')| \le n$ for all $a' \in \mathbb{M}_x$. If the set $\breve{R}(b) \cap \Omega$ is finite, then it is Ab-definable, which in view of $a \in \breve{R}(b) \cap \Omega$ yields that a is Ab-algebraic, and we are done. Assume $\breve{R}(b) \cap \Omega$ is infinite, so $|\Omega \setminus \breve{R}(b)| = m$, say. Representing Ω as an intersection of A-definable subsets of $\mathbb{M}_x$, one of these A-definable sets, call it X, satisfies

$$X \subseteq \mathbb{M}_x, \quad X \supseteq \Omega, \quad |X \setminus \breve{R}(b)| = m.$$

By shrinking R further we can arrange that for all $b' \in \mathbb{M}_y$,

$$\text{either } |X \setminus \breve{R}(b')| \le m, \text{ or } \breve{R}(b') = \emptyset.$$

The set $Y := \{b' \in \mathbb{M}_y : |X \setminus \breve{R}(b')| \le m\}$ is A-definable and contains b, so Y is infinite. Take distinct $b_1, \dots, b_{n+1} \in Y$. Then the sets

$$\breve{R}(b_1) \cap \Omega, \dots, \breve{R}(b_{n+1}) \cap \Omega$$

are cofinite in Ω, so have a common element a', hence $|R(a')| \ge n+1 > n$, a contradiction. □

This allows us to introduce a pregeometry on Ω, but to do this we need to view a set $E \subseteq \mathbb{M}_x$ as a *parameter set.* More precisely, we assign to such E the parameter set $[E]$ consisting of the components of the elements of E; that is, if $x = (x_i)_{i \in I}$ with each variable x_i of sort s_i, then $[E]$ is the parameter set given by

$$[E]_s := \bigcup_{\{i \in I : s_i = s\}} \{e_i : e \in E\}.$$

(Note that if E is not small, then $[E]$ is not small.) When using terminology like "E-definable" and "E-algebraic" for $E \subseteq \mathbb{M}_x$ we regard E as standing for the parameter set $[E]$. Likewise, AE is the parameter set $A[E]$.

For the rest of this section E ranges just over subsets of Ω, and we define the closure operation $\mathrm{cl}_A : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ by

$$\mathrm{cl}_A(E) := \{a \in \Omega : a \text{ is algebraic over } AE\}.$$

Here we have fixed the (small) parameter set A over which Ω is type-definable. Of course, Ω is then also type-definable over any $B \supseteq A$, and this yields likewise a closure operation

$$\mathrm{cl}_B : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega).$$

Theorem 12.4. *The set Ω with the closure operation cl_A is a pregeometry, denoted by Ω_A. For any $a, b \in \Omega$ outside $\mathrm{cl}_A(E)$ there is an $f \in \mathrm{Aut}(\mathbb{M}|AE)$ such that $f(a) = b$.*

Proof. The first three axioms defining pregeometries are obviously satisfied, and the Exchange Axiom is satisfied because of Lemma 12.3 applied to AE instead of A. Let $a, b \in \Omega$ be outside $\mathrm{cl}_A(E)$. It clearly suffices to show:

Claim. $\mathrm{tp}(a|AE) = \mathrm{tp}(b|AE)$. Suppose this claim fails. Then we have an AE-definable $X \subseteq \mathbb{M}_x$ such that $a \in X$ and $b \notin X$. We can assume that $\Omega \cap X$ is finite. (Otherwise, interchange the roles of a and b, replacing X by its complement in $\mathbb{M}_x$.) Then $\Omega \cap X$ is AE-definable, so $a \in \mathrm{cl}_A(E)$, a contradiction. $\square$

We let $\mathrm{rk}_A E$ denote the rank of E in the pregeometry Ω_A, that is, $\mathrm{rk}_A E$ is the size of any maximal independent subset of E ("independent" in the sense of the pregeometry Ω_A).

We have a global type $\mathbf{p} \in \mathrm{St}_x(\mathbb{M})$ associated to Ω:

$$\mathbf{p} := \{X \in \mathrm{Def}_x(\mathbb{M}) : \Omega \setminus X \text{ is finite }\}.$$

If Ω is strongly minimal, then $\mathbf{p}$ is the global type determined by Ω, and $\mathrm{MR}(\mathbf{p}) = 1$. By the result above and its proof, an element $a \in \Omega$ realizes $\mathbf{p} \restriction AE$ iff $a \notin \mathrm{cl}_A(E)$. An n-tuple $(a_1, \ldots, a_n) \in \Omega^n$ is said to be *E-independent* if $a_1, \ldots, a_n$ are distinct and $\{a_1, \ldots, a_n\} \subseteq \Omega$ is E-independent in the pregeometry Ω_A that is, independent in the pregeometry $\Omega_A|E$. For $E = \emptyset$ we write "independent" instead of "$\emptyset$-independent".

Lemma 12.5. *Let $n \geq 1$, and suppose $(a_1, \ldots, a_n) \in \Omega^n$ is E-independent and $(b_1, \ldots, b_n) \in \Omega^n$ is E-independent. Then there is an $f \in \mathrm{Aut}(\mathbb{M}|AE)$ such that $f(a_i) = b_i$ for $i = 1, \ldots, n$.*

Proof. By the theorem above we can take a $g \in \mathrm{Aut}(\mathbb{M}|AE)$ such that $g(a_1) = b_1$. By replacing $(a_1, \ldots, a_n)$ by its image under g (and renaming) we reduce to the case that $a_1 = b_1$. If $n > 1$, use that $(a_2, \ldots, a_n)$ and $(b_2, \ldots, b_n)$ are $E \cup \{a_1\}$-independent, and proceed inductively. $\square$

Note that for each n there are $a_1, \ldots, a_n \in \Omega$ such that $(a_1, \ldots, a_n)$ is independent: just take $a_i \in \Omega$ to be outside $\mathrm{cl}_A(\{a_1, \ldots, a_{i-1}\})$ for $i = 1, \ldots, n$. Thus the pregeometry Ω_A has infinite rank.

We now continue with the more restrictive strongly minimal case, that is, in the rest of this section we assume:

$$\Omega \text{ is } A\text{-definable}, \quad \mathrm{MR}(\Omega) = 1, \quad \mathrm{MD}(\Omega) = 1.$$

(But we do not assume that T is totally transcendental.) It follows by induction on n and an exercise in Section 2.3 that $\mathrm{MR}(\Omega^n) \leq n$.

Lemma 12.6. *Let E be small. Then for all $(a_1, \ldots, a_n) \in \Omega^n$,*

$$\mathrm{MR}((a_1, \ldots, a_n)|AE) = n \iff (a_1, \ldots, a_n) \text{ is } E\text{-independent.}$$

Proof. By induction on n. The case $n = 0$ is trivial. Assume the lemma holds for a certain n, and let $a = (a_1, \dots, a_n, a_{n+1}) \in \Omega^{n+1}$.

To obtain the direction $\Leftarrow$, let a be E-independent, and let $X \subseteq \Omega^{n+1}$ be AE-definable with $a \in X$; it remains to show that $\mathrm{MR}(X) = n + 1$. By identifying Ω^{n+1} with $\Omega \times \Omega^n$ we make X into a binary relation:

$$X \subseteq \Omega \times \Omega^n.$$

The set $X(a_1) \subseteq \Omega^n$ is definable over AEa_1 and contains the n-tuple $(a_2, \dots, a_{n+1})$, which is $(E \cup \{a_1\})$-independent, so $\mathrm{MR}\big(X(a_1)\big) = n$ by the inductive assumption. Hence $\mathrm{MR}\big(X(b)\big) = n$ for $b \in \Omega$ outside $\mathrm{cl}_A(E)$, and there are infinitely many such b. These b produce infinitely many disjoint definable subsets $\{b\} \times X(b)$ of X of Morley rank n. Hence $\mathrm{MR}(X) = n+1$.

For the direction $\Rightarrow$, suppose a is not E-independent. By permutating the coordinates of the $(n + 1)$-tuple a we arrange that a_{n+1} is algebraic over AEa' where $a' := (a_1, \dots, a_n)$, so we have an AE-definable relation $Y \subseteq \Omega^n \times \Omega = \Omega^{n+1}$ such that $a \in Y$ and $Y(a')$ is finite, say $|Y(a')| = m$. We can shrink Y to arrange that $|Y(b')| \le m$ for all $b' \in \Omega^n$. It follows that $\mathrm{MR}(Y) \le \mathrm{MR}(\Omega^n) \le n$, and thus $\mathrm{MR}(a|AE) \le n$. □

Corollary 12.7. *Given any $a_1, \dots, a_n \in \Omega$, we have*

$$\mathrm{MR}\big((a_1, \dots, a_n)|A\big) = \mathrm{rk}_A\{a_1, \dots, a_n\}.$$

Proof. Suppose $\mathrm{rk}_A\{a_1, \dots, a_n\} = k$. By permutating the coordinates of the tuple $(a_1, \dots, a_n)$ we can arrange that $(a_1, \dots, a_k)$ is independent, and that $a_{k+1}, \dots, a_n$ are algebraic over $A(a_1, \dots, a_k)$. Then $\mathrm{MR}\big((a_1, \dots, a_n)|A\big) = \mathrm{MR}\big((a_1, \dots, a_k)|A\big) = k$ by Lemma 10.4 and Lemma 12.6. □

Corollary 12.8. *Let $a = (a_1, \dots, a_m) \in \Omega^m$ and $b = (b_1, \dots, b_n) \in \Omega^n$, so $(a, b) := (a_1, \dots, a_m, b_1, \dots, b_n) \in \Omega^{m+n}$. Then*

$$\mathrm{MR}((a, b)|A) = \mathrm{MR}(a|bA) + \mathrm{MR}(b|A).$$

Proof. Let $E = \{a_1, \dots, a_m\}$ and $F = \{b_1, \dots, b_n\}$. Then $\mathrm{MR}((a, b)|A) = \mathrm{rk}_A(E \cup F)$, $\mathrm{MR}(a|bA) = \mathrm{rk}_{bA}(E)$, and $\mathrm{MR}(b|A) = \mathrm{rk}_A(F)$. Now use that

$$\mathrm{rk}_A(E \cup F) = \mathrm{rk}_{bA}(E) + \mathrm{rk}_A(F).$$

□

For definable $Y \subseteq \Omega^n$ we can determine $\mathrm{MR}(Y)$ inductively. To see how, let $n > 0$ and use the identification $\Omega^n = \Omega^{n-1} \times \Omega$ to view Y as a binary relation: $Y \subseteq \Omega^{n-1} \times \Omega$. By an exercise in Section 10 we have an m such that for all $b \in \Omega^{n-1}$, either $|Y(b)| \le m$ or $|\Omega \setminus Y(b)| \le m$. Put

$$Y' := \{b \in \Omega^{n-1} : 1 \le |Y(b)| \le m\},$$
$$Y'' := \{b \in \Omega^{n-1} : |\Omega \setminus Y(b)| \le m\}.$$

Then Y' and Y'' are definable, and Y is the disjoint union of $(Y' \times \Omega) \cap Y$ and $(Y'' \times \Omega) \cap Y$, whose Morley ranks are $\mathrm{MR}(Y')$ and $\mathrm{MR}(Y'')+1$, respectively.

(Why?) Thus

$$\mathrm{MR}(Y) = \max\{\mathrm{MR}(Y'),\ \mathrm{MR}(Y'') + 1\}.$$

This argument can be done "with parameters" to obtain the definability of Morley rank within definable families of subsets of Ω^n:

Corollary 12.9. *Let X be a definable set in $\mathbb{M}$ and $R \subseteq X \times \Omega^n$ a definable relation. Then $\{a \in X : \mathrm{MR}\big(R(a)\big) \geq d\}$ is definable for $d = 0, \ldots, n$.*

Proof. We proceed by induction on n. The case $n = 0$ being obvious, assume $n > 0$. Identify Ω^n with $\Omega^{n-1} \times \Omega$ and take m such that for all $a \in X$ and $b \in \Omega^{n-1}$, either $|X(a)(b)| \leq m$ or $|\Omega \setminus X(a)(b)| \leq m$. Put

$$R' := \{(a, b) \in X \times \Omega^{n-1} : 1 \leq |X(a)(b)| \leq m\} \subseteq X \times \Omega^{n-1},$$
$$R'' := \{(a, b) \in X \times \Omega^{n-1} : |\Omega \setminus X(a)(b)| \leq m\} \subseteq X \times \Omega^{n-1}.$$

By the argument preceding this corollary we have for each $a \in X$:

$$\mathrm{MR}\big(R(a)\big) = \max\{\mathrm{MR}\big(R'(a)\big),\ \mathrm{MR}\big(R''(a)\big) + 1\}.$$

The desired result now follows by applying the inductive assumption to the definable relations R' and R''. $\square$

Extension to algebraic elements. In this subsection we assume

$$T \text{ is totally transcendental.}$$

An element $h \in \mathbb{M}_y$ is said to be *A-algebraic over Ω* if h is algebraic over AE for some (finite) E. We extend some of our results on tuples in Ω to elements that are A-algebraic over Ω.

Lemma 12.10. *Suppose that $h \in \mathbb{M}_y$ is A-algebraic over Ω. Then there are $a = (a_1, \ldots, a_m) \in \Omega^m$ and $b = (b_1, \ldots, b_n) \in \Omega^n$ such that the tuple*

$$(a, b) := (a_1, \ldots, a_m, b_1, \ldots, b_n) \in \Omega^{m+n}$$

is independent in the pregeometry Ω_A, h is independent from a over A, and h and b are interalgebraic over Aa (so $\mathrm{MR}(h|A) = \mathrm{MR}(h|Aa) = n$).

Proof. Take a finite E of minimal size such that h is algebraic over AE. Note that then E is independent in Ω_A. Let $a_1, \ldots, a_m \in E$ be distinct such that $\{a_1, \ldots, a_m\}$ is a basis of E in the pregeometry Ω_{Ah}, and let $E = \{a_1, \ldots, a_m, b_1, \ldots, b_n\}$ with $|E| = m + n$. Then each $b_i \in \mathrm{cl}_{Ah}\{a_1, \ldots, a_m\}$, so $(b_1, \ldots, b_n)$ is algebraic over $Aa_1 \ldots a_m h$. Hence h and $(b_1, \ldots, b_n)$ are interalgebraic over $Aa_1 \ldots a_m$. Also,

$$\mathrm{MR}\big((a_1, \ldots, a_m)|Ah\big) = m = \mathrm{MR}\big((a_1, \ldots, a_m)|A\big),$$

hence $(a_1, \ldots, a_m) \mathop{\downarrow}\limits_A h$, and thus $h \mathop{\downarrow}\limits_A (a_1, \ldots, a_m)$, as desired. $\square$

In this lemma, with $n = \mathrm{MR}(h|A)$, we can take for $(b_1, \ldots, b_n)$ any element of Ω^n that is independent in the pregeometry Ω_{Ah}, since all such n-tuples are conjugate under $\mathrm{Aut}(\mathbb{M}|Ah)$.

Lemma 12.11. *Let $g \in \mathbb{M}_y$ and $h \in \mathbb{M}_z$ be A-algebraic over Ω. Then*

$$\mathrm{MR}\big((g,h)|A\big) = \mathrm{MR}(g|hA) + \mathrm{MR}(h|A).$$

Proof. By the previous lemma we can take $a \in \Omega^p$ and $b \in \Omega^m$ such that

$$g \underset{A}{\downarrow} a \text{ and } g \text{ is interalgebraic with } b \text{ over } Aa,$$

and we can take $c \in \Omega^q$ and $d \in \Omega^n$ such that

$$h \underset{A}{\downarrow} c \text{ and } h \text{ is interalgebraic with } d \text{ over } Ac.$$

We arrange that for $(g,h) \in \mathbb{M}_y \times \mathbb{M}_z$ and $(a,c) \in \Omega^{p+q}$ we have

$$(g,h) \underset{A}{\downarrow} (a,c).$$

This is done as follows: Take an automorphism of $\mathbb{M}$ over Ag that sends a to a realization of a nonforking extension of $\mathrm{tp}(a|Ag)$ to Agh, and replace a and b by their images under this automorphism; this achieves

$$\mathrm{MR}(a|Agh) = \mathrm{MR}(a|Ag) = \mathrm{MR}(a|A).$$

Next, take an automorphism of $\mathbb{M}$ over Ah that sends c to a realization of a nonforking extension of $\mathrm{tp}(c|Ah)$ to $Agha$ and replace c and d by their images under this automorphism; this guarantees

$$\mathrm{MR}(c|Agha) = \mathrm{MR}(c|Ah) = \mathrm{MR}(c|A).$$

Hence

$$\begin{aligned}\mathrm{MR}((a,c)|A) &= \mathrm{MR}(c|Aa) + \mathrm{MR}(a|A) = \mathrm{MR}(c|Agha) + \mathrm{MR}(a|Agh)\\ &= \mathrm{MR}((a,c)|Agh),\end{aligned}$$

so $(a,c) \underset{A}{\downarrow} (g,h)$, and thus $(g,h) \underset{A}{\downarrow} (a,c)$ by symmetry. Then

$$\begin{aligned}\mathrm{MR}\big((g,h)|A\big) &= \mathrm{MR}\big((g,h)|Aac\big) = \mathrm{MR}((b,d)|Aac)\\ &= \mathrm{MR}(b|Aacd) + \mathrm{MR}(d|Aac)\\ &= \mathrm{MR}(b|Aacd) + \mathrm{MR}(h|Aac).\end{aligned}$$

From $(g,h) \underset{A}{\downarrow} (a,c)$ we obtain $g \underset{Ah}{\downarrow} (a,c)$, so $\mathrm{MR}(g|Ah) = \mathrm{MR}(g|Aach) = \mathrm{MR}(b|Aach) = \mathrm{MR}(b|Aacd)$. Likewise, using $h \underset{A}{\downarrow} (a,c)$ we get

$$\mathrm{MR}(h|A) = \mathrm{MR}(h|Aac).$$

□

A set $Z \subseteq \mathbb{M}_z$ is said to be *almost strongly minimal with respect to* Ω, A if Z is definable and each element of Z is A-algebraic over Ω.

An easy saturation argument yields:

Lemma 12.12. *A definable set $Z \subseteq \mathbb{M}_z$ is almost strongly minimal with respect to Ω, A if and only if there are $k, N \in \mathbb{N}$ and A-definable relations $R_1 \subseteq \Omega^{n_1} \times \mathbb{M}_z, \dots, R_k \subseteq \Omega^{n_k} \times \mathbb{M}_z$ whose sections are all finite of size at most N, such that each element of Z lies in a section of some R_i.*

With Lemmas 12.11 and 12.12 some properties of almost strongly minimal sets follow easily from similar properties of definable subsets of Ω^n:

Exercises. Let $Z \subseteq \mathbb{M}_z$ be almost strongly minimal with respect to Ω, A. Show that $\mathrm{MR}(Z) < \omega$. Let also $Y \subseteq \mathbb{M}_y$ and $R \subseteq Y \times Z$ be definable. Prove:

(1) $\{g \in Y : \mathrm{MR}\big(R(g)\big) \geq d\}$ is definable;
(2) if $\mathrm{MR}\big(R(g)\big) = d$ for all $g \in Y$, then $\mathrm{MR}(R) = \mathrm{MR}(Y) + d$.

13. Modularity

In this section we view the collection of closed sets in a pregeometry Ω as a *lattice* with respect to inclusion. The meet and join operations of this lattice are given by

$$E \wedge F := E \cap F, \qquad E \vee F := \mathrm{cl}(E \cup F).$$

We begin with a short excursion on (modular) lattices in general.

Modular Lattices. A poset (partially ordered set) is said to be a *lattice* if any elements a, b in it have a least upperbound $a \vee b$ (their join) and a greatest lowerbound $a \wedge b$ (their meet). For example, a boolean algebra viewed as a poset is a lattice.

In this subsection P is a lattice, and a, b, c range over P. The partial ordering of P can be recovered from its join operation as well as from its meet operation:

$$a \leq b \iff a \vee b = b \iff a \wedge b = b.$$

The join and meet operations of P are idempotent ($a \vee a = a \wedge a = a$), commutative, associative, and satisfy the absorption identities:

$$a \wedge (a \vee b) = a \vee (a \wedge b) = a.$$

A *sublattice* of P is just a subset of P closed under $\vee$ and $\wedge$, and is thus a lattice with respect to the induced partial ordering. For example, given a, c we have the sublattice $[a, c] := \{b : a \leq b \leq c\}$ of P. Note that for all a, b, c:

$$a \geq c \implies a \wedge (b \vee c) \geq (a \wedge b) \vee c.$$

The lattice P is said to be *modular* if for all a, b, c

$$a \geq c \implies a \wedge (b \vee c) = (a \wedge b) \vee c.$$

This identity implies its own dual: if P is modular, then for all a, b, c

$$a \leq c \implies a \vee (b \wedge c) = (a \vee b) \wedge c.$$

Of course, a sublattice of a modular lattice is modular as well.

Example. Let G be an abelian (additively written) group and consider the lattice of all subgroups of G with respect to inclusion. Its join and meet operations are given by $A \vee B = A + B, \quad A \wedge B = A \cap B$. It is easy to check that this lattice is modular. Thus the lattice of linear subspaces of a vector space over a division ring is modular.

There are many results about modular lattices, but we develop here only what is needed for our purpose. Let $\mathcal{N}_5$ be the lattice with exactly five distinct elements o, p, q, r, i with o as least element, i as the greatest element, and $p < q,\ p \vee r = i,\ q \wedge r = o$. This pentagon shaped lattice $\mathcal{N}_5$ is *not* modular: $q \geq p$, but $q \wedge (r \vee p) = q$ and $(q \wedge r) \vee p = p$. The significance of $\mathcal{N}_5$ is that modularity of P is equivalent to P not containing a copy of $\mathcal{N}_5$:

Proposition 13.1. *The following conditions on P are equivalent:*

(1) *P is modular;*
(2) *$a \wedge (b \vee c) = a \wedge \big((b \wedge (a \vee c)) \vee c\big)$ for all a, b, c;*
(3) *P has no sublattice isomorphic to $\mathcal{N}_5$.*

Proof. For (1) $\Rightarrow$ (2), assume P is modular. Since $a \vee c \geq c$, this yields

$$(b \wedge (a \vee c)) \vee c = (b \vee c) \wedge (a \vee c),$$

hence $a \wedge \big((b \wedge (a \vee c)) \vee c\big) = a \wedge \big((b \vee c) \wedge (a \vee c)\big) = a \wedge (b \vee c)$.

For (2) $\Rightarrow$ (3), note that the identity of (2) fails in $\mathcal{N}_5$:

$$q \wedge (r \vee p) = q \wedge i = q \text{ and } q \wedge \big((r \wedge (q \vee p)) \vee p = q \wedge ((r \wedge q) \vee p) = q \wedge p = p.$$

For (3) $\Rightarrow$ (1), suppose P is not modular. Take a, b, c such that $a \geq c$ but $a \wedge (b \vee c) \neq (a \wedge b) \vee c$. We leave it as an exercise to show that

$$a \wedge b,\ (a \wedge b) \vee c,\ a \wedge (b \vee c),\ b \vee c,\ b$$

are distinct, and are the elements of a sublattice of P isomorphic to $\mathcal{N}_5$. $\square$

Given a second lattice Q, we make the cartesian product set $P \times Q$ into a lattice by defining (for $d, e \in Q$):

$$(a, d) \leq (b, e) \iff a \leq b \text{ and } d \leq e,$$

so that $(a, d) \vee (b, e) = (a \vee b, d \vee e)$, $(a, d) \wedge (b, e) = (a \wedge b, d \wedge e)$. The following isomorphisms are very useful, and easily verified:

Proposition 13.2. *Let P be modular. Then we have an isomorphism*

$$s \mapsto s \vee b \ :\ [a \wedge b, a] \to [b, a \vee b],$$

of sublattices, with inverse isomorphism

$$t \mapsto t \wedge a \ :\ [b, a \vee b] \to [a \wedge b, a].$$

Moreover, the map

$$(s, s') \mapsto s \vee s' \ : [a \wedge b, a] \times [a \wedge b, b] \to [a \wedge b, a \vee b]$$

is an isomorphism of the product lattice $[a \wedge b, a] \times [a \wedge b, b]$ onto a sublattice of $[a \wedge b, a \vee b]$, with inverse given by $c \mapsto (c \wedge a, c \wedge b)$.

We now return to the setting of a pregeometry with its lattice of closed sets.

Modular pairs. In this subsection E and F are closed sets of finite rank in a pregeometry Ω. Consider the inclusion diagram

....

Take bases B, C, D of $E \wedge F$, $E|E \wedge F$ and $F|E \wedge F$, respectively. Then B, C, D are finite, pairwise disjoint, $B \cup C$ is a basis of E, $B \cup D$ is a basis of F, and $B \cup C \cup D$ generates $E \vee F$. Hence

$$\operatorname{rk} E \vee F \leq \operatorname{rk}(E \wedge F) + \operatorname{rk}(E|E \wedge F) + \operatorname{rk}(F|E \wedge F).$$

In view of

$$\operatorname{rk}(E|E \wedge F) = \operatorname{rk}(E) - \operatorname{rk}(E \wedge F) \text{ and } \operatorname{rk}(F|E \wedge F) = \operatorname{rk}(F) - \operatorname{rk}(E \wedge F),$$

this inequality becomes

$$\operatorname{rk} E + \operatorname{rk} F \geq \operatorname{rk}(E \vee F) + \operatorname{rk}(E \wedge F),$$

in other words, $\operatorname{rk}(E|E \wedge F) \geq \operatorname{rk}(E \vee F|F)$. This argument also shows:

$$\begin{aligned} \operatorname{rk} E + \operatorname{rk} F = \operatorname{rk}(E \vee F) + \operatorname{rk}(E \wedge F) &\iff C \text{ is independent in } \Omega|F \\ &\iff \operatorname{rk}(E|E \wedge F) = \operatorname{rk}(E \vee F|F). \end{aligned}$$

Let us call the pair E, F *modular* if

$$\operatorname{rk} E + \operatorname{rk} F = \operatorname{rk}(E \vee F) + \operatorname{rk}(E \wedge F),$$

equivalently, $\operatorname{rk}(E|E \wedge F) = \operatorname{rk}(E \vee F|F)$. Note that if $\operatorname{rk}(E|E \wedge F) \leq 1$, then E, F is modular. (This is because $C \cap F = \emptyset$.)

Next, let $E \wedge F \subseteq D \subseteq E$ with D closed. Then the inclusion diagram above consists of two subdiagrams:

....

Take a basis C_1 of $E|D$, and a basis C_2 of $D|E \wedge F$, and let $C = C_1 \cup C_2$ (a disjoint union) be the corresponding basis of $E|E \wedge F$. Note that then C_1 generates $E \vee F|D \vee F$ and C_2 generates $D \vee F|F$. Hence E, F is modular if and only if C_1 is a basis of $E \vee F|D \vee F$ and C_2 is a basis of $D \vee F|F$, that is, if and only if $\operatorname{rk}(E|D) = \operatorname{rk}(E \vee F|D \vee F)$ and D, F is modular.

Claim. $\operatorname{rk}(E|D) = \operatorname{rk}(E \vee F|D \vee F)$ if and only if $E \wedge (D \vee F) = D$ and $E, D \vee F$ is modular.

The "if" direction is clear, and the "only if" direction follows from

$$\operatorname{rk}(E|D) \geq \operatorname{rk}(E|E \wedge (D \vee F)) \geq \operatorname{rk}(E \vee F|D \vee F).$$

Corollary 13.3. *Suppose E, F is modular. Consider the maps*

$$\begin{aligned} \text{up} &: [E \wedge F, E] \to [F, E \vee F], \quad \text{up}(D) := D \vee F \\ \text{down} &: [F, E \vee F] \to [E \wedge F, E], \quad \text{down}(G) := E \wedge G. \end{aligned}$$

Then $\text{down} \circ \text{up} =$ *the identity on $[E \wedge F, E]$, in particular,* up *is injective. If in addition E, G is modular for each $G \in [F, E \vee F]$, then both maps are isomorphisms between sublattices of the lattice of closed sets in Ω.*

Proof. The arguments above proves the first assertion. Let $G \in [F, E \vee F]$ and put $D = E \wedge G$. Then $D \vee F \subseteq G$, and D, F is modular, so if E, G is also modular, then necessarily $G = D \vee F$, as is easily seen in an inclusion diagram. □

Modular pregeometries. Let Ω be a pregeometry; we let E, F range over subsets of Ω. Modularity is a very strong condition, and appears in model theory in the company of other properties that a pregeometry Ω may or may not have:

(1) *Triviality*: $\operatorname{cl} E = \bigcup_{a \in E} \operatorname{cl}\{a\}$, for all finite E;
(2) *Modularity*: E, F is modular for all closed E and F of finite rank;
(3) *Local modularity*: E, F is modular for all closed E and F of finite rank with $E \wedge F \neq \operatorname{cl} \emptyset$;
(4) *Local finiteness*: $\operatorname{cl} E$ is finite for all finite E;
(5) *Homogeneity*: for any finite E and $a, b \in \Omega \setminus \operatorname{cl} E$ the pregeometry has an automorphism fixing E pointwise and sending a to b.

In the first example of the previous section, where $\operatorname{cl} E = E$ for all E, the pregeometry is trivial, modular, locally finite, and homogeneous. The second example, where Ω is the pregeometry of a vector space over a division ring $\mathbf{k}$, is modular and homogeneous; it is also locally finite if $\mathbf{k}$ is a finite field. If in the third example the field K is algebraically closed, then the pregeometry is homogeneous. Note also that the pregeometry Ω_A of Theorem 12.4 is homogeneous. If Ω is the affine geometry of a vector space over a division ring $\mathbf{k}$, then Ω is homogeneous, and locally finite if in addition $\mathbf{k}$ is finite. Note that if Ω is a geometry, then triviality of Ω means that $\operatorname{cl} E = E$ for all E.

Combinatorial geometries became important in model theory when Zilber conjectured the following theorem, subsequently proved by him and independently by Cherlin and by Mills:

Theorem 13.4. *Suppose Ω is a homogeneous locally finite geometry of infinite rank. Then one of the following happens:*

(1) *Ω is trivial;*
(2) *Ω is isomorphic to the projective geometry associated to a vector space of infinite dimension over a finite field;*
(3) *Ω is isomorphic to the affine geometry of a vector space of infinite dimension over a finite field.*

We do not prove this here, and just mention it to give some perspective on what we are engaged in. The next result has a routine proof.

Lemma 13.5. *Each of the five properties listed above is inherited from Ω by the geometry Ω' associated to Ω, and by each localization $\Omega|X$ with $X \subseteq \Omega$.*

Lemma 13.6. *Ω is modular if and only if its lattice of closed sets is modular.*

Proof. The lattice of closed sets is modular iff its sublattice of closed sets of finite rank is modular. This follows easily from the definitions and the fact that a closed set is the directed union of its closed subsets of finite rank.

Suppose now that Ω is modular. To show that the lattice of closed sets of finite rank is modular, assume towards a contradiction that this lattice has a sublattice isomorphic to $\mathcal{N}_5$. Then we have closed sets E, F of finite rank, and a closed set D strictly between $E \wedge F$ and E such that $D \vee F = E \vee F$, but this contradicts Corollary 13.3.

For the converse, assume the lattice of closed sets is modular, and let E and F be closed of finite rank. Let $n := \mathrm{rk}(E|F)$, so we have a strictly increasing sequence

$$E \wedge F = D_0 \subset D_1 \subset \cdots \subset D_n = E$$

of closed sets. By Proposition 13.2 this yields a strictly increasing sequence

$$F = D_0 \vee F \subset D_1 \vee F \subset \cdots \subset D_n \vee F = E \vee F$$

of closed sets, hence $\mathrm{rk}(E \vee F|F) \geq n = \mathrm{rk}(E|E \wedge F)$. Also

$$\mathrm{rk}(E \vee F|F) \leq \mathrm{rk}(E|E \wedge F),$$

and thus E, F is modular. $\square$

The next lemma has a routine proof.

Lemma 13.7. *Ω is locally modular if and only if each localization $\Omega|\{a\}$ with $a \in \Omega \setminus \mathrm{cl}\,\emptyset$ is modular. If Ω is homogeneous and $\Omega|\{a\}$ is modular for some $a \in \Omega \setminus \mathrm{cl}\,\emptyset$, then Ω is locally modular.*

Lemma 13.8. *If E, F is modular for all closed E, F of finite rank with $\mathrm{rk}(E|E \wedge F) = 2$, then Ω is modular.*

Proof. We prove the contrapositive. Assume Ω is not modular, so E, F is not modular for certain closed E and F of finite rank; take such E and F for which $n := \mathrm{rk}(E|E \wedge F)$ is minimal, so $n \geq 2$ and $\mathrm{rk}(E \vee F|F) < n$. Let $C = \{c_1, \ldots, c_n\}$ be a basis of $E|E \vee F$, and put

$$E_1 := \mathrm{cl}((E \wedge F) \cup \{c_1, \ldots, c_{n-1}\},$$

so $\mathrm{rk}(E_1|E \wedge F) = n - 1$. Also E_1, F is modular by the minimality of n, hence $\mathrm{rk}(E_1 \vee F|F) = n - 1$, and thus $E_1 \vee F = E \wedge F$. Next, let

$$E_2 := \mathrm{cl}((E \wedge F) \cup \{c_1, \ldots, c_{n-2}\},$$

so $\mathrm{rk}(E_1|E_2) = 1$, and thus for $G := E_2 \vee F$ we have $\mathrm{rk}(E \vee F|G) = 1$ and $E_1 \wedge G = E_2$, by Corollary 13.3.

Claim. $E \wedge G = E_2$.

To prove the claim, note that $E_2 \subseteq E \wedge G$. Suppose E_2 is properly contained in $E \wedge G$. Then $\mathrm{rk}(E \wedge G | E \wedge F) = n-1$, so $E \wedge G, F$ is modular by minimality of n. Hence $\mathrm{rk}(G|F) \geq \mathrm{rk}((E \wedge G) \vee F | F) = \mathrm{rk}(E \wedge G | E \wedge F) = n-1$, contradicting $\mathrm{rk}(G|F) = n-2$. This proves the claim.

It follows that E, G is not modular, since $\mathrm{rk}(E|E_2) = 2$ and $\mathrm{rk}(E \vee G|G) = \mathrm{rk}(E \vee F|G) = 1$. Thus E, G has the desired property. $\square$

We can go further along these lines:

Lemma 13.9. *If E, F is modular for all closed E, F of finite rank with $\mathrm{rk}(E|E \wedge F) = 2 = \mathrm{rk}(F|E \wedge F)$, then Ω is modular.*

Proof. Again, we prove the contrapositive, and assume Ω is not modular. By the previous lemma we can take closed E, F of finite rank such that E, F is not modular and $\mathrm{rk}(E|E \wedge F) = 2$. Take such E, F with minimal $m := \mathrm{rk}(F|E \wedge F)$, so $m \geq 2$. Let $\{b_1, \ldots, b_m\}$ be a basis of $F|E \wedge F$, let $H := \mathrm{cl}(E \cup \{b_1\})$. Then

$$\begin{aligned} \mathrm{rk}(H|E \wedge F) &= \mathrm{rk}(H|E) + \mathrm{rk}(E|E \wedge F) = 3 \\ &= \mathrm{rk}(H|H \wedge F) + \mathrm{rk}(H \wedge F|E \wedge F), \end{aligned}$$

so $\mathrm{rk}(H \wedge F|E \wedge F)$ equals 1 or 2. It it equals 1, then $\mathrm{rk}(H|H \wedge F) = 2$ and $\mathrm{rk}(F|H \wedge F) = m-1$, and H, F is nonmodular, contradicting the minimality of m. So $\mathrm{rk}(H \wedge F|E \wedge F) = 2$, and thus $E, H \wedge F$ is nonmodular of the desired form. $\square$

Lemma 13.10. *The following are equivalent:*

(1) *Ω is modular;*
(2) *for all closed E and F such that $\mathrm{rk}\, E = 2$, F has finite rank, and $\mathrm{rk}(E \vee F|F) = 1$ we have $E \wedge F \neq \mathrm{cl}\, \emptyset$;*
(3) *for each nonempty closed E and $b \in \Omega$ we have*

$$\mathrm{cl}(E \cup \{b\}) = \bigcup_{a \in E} \mathrm{cl}\{a, b\};$$

(4) *for all nonempty closed E, F we have*

$$E \vee F = \bigcup_{a \in E, b \in F} \mathrm{cl}\{a, b\}.$$

Proof. The direction (1) $\Longrightarrow$ (2) is clear. For the converse, suppose Ω is not modular. Then Lemma 13.8 gives closed E, F of finite rank with

$$\mathrm{rk}(E|E \wedge F) = 2, \qquad \mathrm{rk}(E \vee F|F) = 1.$$

Let $\{a, b\}$ be a basis of $E|E \wedge F$ and put $E^* := \mathrm{cl}\{a, b\}$. Then $\mathrm{rk}\, E^* = 2$ and $\mathrm{rk}(E^*|E^* \wedge F) = 2$, so $E^* \wedge F = \mathrm{cl}\, \emptyset$. Also $E^* \vee F = E \vee F$, so $\mathrm{rk}(E^* \vee F|F) = 1$. This contradicts (2), and proves (the contrapositive of) (2) $\Longrightarrow$ (1). For (1) $\Longrightarrow$ (3), assume (1), let E be closed and nonempty,

$b \in \Omega$ and $c \in \mathrm{cl}(E \cup \{b\})$. To find $a \in E$ such that $c \in \mathrm{cl}\{a, b\}$ we can assume E has finite rank and $c \notin \mathrm{cl}\{b\}$. By modularity,

$$\mathrm{rk}(E \cup \{b, c\}) = \mathrm{rk}\, E + \mathrm{rk}\{b, c\} - \mathrm{rk}(E \wedge \mathrm{cl}\{b, c\}), \text{ and}$$
$$\mathrm{rk}(E \cup \{b, c\}) = \mathrm{rk}(E \cup \{b\}) = \mathrm{rk}\, E + \mathrm{rk}\{b\} - \mathrm{rk}(E \wedge \mathrm{cl}\{b\}).$$

Also, $\mathrm{rk}\{b, c\} = \mathrm{rk}\{b\} + 1$, so $E \wedge \mathrm{cl}\{b\}$ is properly contained in $E \wedge \mathrm{cl}\{b, c\}$. Take $a \in E \wedge \mathrm{cl}\{b, c\}$ such that $a \notin \mathrm{cl}\{b\}$. Then $c \in \mathrm{cl}\{a, b\}$ by exchange.

To prove (3) $\Longrightarrow$ (4), assume (3), and let E, F be closed, nonempty, and, without loss of generality, of finite rank. Let $c \in E \vee F$. We have to find $a \in E$ and $b \in F$ such that $c \in \mathrm{cl}\{a, b\}$. We proceed by induction on $\mathrm{rk}\, F$. If $\mathrm{rk}\, F = 0$, we take any $b \in F$, and apply (2). Next, let $F = \mathrm{cl}(F_0 \cup \{b_1\})$ with closed F_0 and $\mathrm{rk}\, F = \mathrm{rk}\, F_0 + 1$. Then by (3) we have an $a_0 \in E \vee F_0$ such that $c \in \mathrm{cl}\{a_0, b_1\}$. The inductive assumption gives $a \in E$ and $b_0 \in F_0$ such that $a_0 \in \mathrm{cl}\{a, b_0\}$. Then $c \in \mathrm{cl}\{a, b_0, b_1\} \subseteq \mathrm{cl}(F \cup \{a\})$, so by (3) there is $b \in F$ such that $c \in \mathrm{cl}\{a, b\}$.

We leave (4) $\Longrightarrow$ (1) to the reader. □

14. Modularity, One-basedness, and Linearity

We now return to the setting of a totally transcendental theory T with monster model $\mathbb{M}$ and a strongly minimal A-definable set $\Omega \subseteq \mathbb{M}_z$. This gives the pregeometry Ω_A with closure operation cl_A. We shall prove in this section that local modularity of Ω_A is equivalent to various other conditions.

Given a tuple $a = (a_1, \ldots, a_n) \in \Omega^n$ and $E \subseteq \Omega$, we put

$$\mathrm{cl}_A(a) := \mathrm{cl}_A\{a_1, \ldots, a_n\}, \text{ a subset of } \Omega,$$
$$\mathrm{rk}_A(a|E) := \mathrm{rk}_A(\{a_1, \ldots, a_n\}|E).$$

Lemma 14.1. *The following are equivalent:*

(1) Ω_A *is modular;*
(2) *for all* $a \in \Omega^m$ *and* $b \in \Omega^n$*:* $a \underset{B}{\downarrow} b$*, where* $B := A\big(\mathrm{cl}_A(a) \cap \mathrm{cl}_A(b)\big)$.

Proof. Recall first that by our notational conventions

$$\mathrm{rk}_A(a|E) = \mathrm{rk}_A\left(\mathrm{cl}_A(a)|E\right), \quad (a \in \Omega^n, E \subseteq \Omega).$$

Suppose Ω_A is modular, and let $a \in \Omega^m$ and $b \in \Omega^n$. Then

$$\mathrm{rk}_A\left(a|\,\mathrm{cl}_A(a) \cap \mathrm{cl}_A(b)\right) = \mathrm{rk}_A\left((a, b)|\,\mathrm{cl}_A(b)\right) = \mathrm{rk}_A\left(a|\,\mathrm{cl}_A(b)\right),$$

hence $\mathrm{MR}\left(a|A(\mathrm{cl}_A(a) \cap \mathrm{cl}_A(b))\right) = \mathrm{MR}(a|A\,\mathrm{cl}_A(b))$ by Corollary 12.7, that is, $a \underset{B}{\downarrow} b$, with $B = A\big(\mathrm{cl}_A(a) \cap \mathrm{cl}_A(b)\big)$.

Suppose next that Ω_A is not modular. Then by Lemma 13.8 there are $a \in \Omega^2$ and $b \in \Omega^n$ such that $\mathrm{rk}_A\left(a|\,\mathrm{cl}_A(a) \cap \mathrm{cl}_A(b)\right) = 2$ but $\mathrm{rk}_A\left(a|\,\mathrm{cl}_A(b)\right) < 2$. Hence

$$\mathrm{MR}\left(a|A(\mathrm{cl}_A(a) \cap \mathrm{cl}_A(b))\right) = 2, \qquad \mathrm{MR}\left(a|A\,\mathrm{cl}_A(b)\right) < 2,$$

so $a \mathop{\smile\!\!\!\!\!\!\,|}_{B} b$ fails, with B as above. $\square$

One-basedness. In the rest of this section T has EI, in addition to being totally transcendental. This assumption is just for the convenience of having simpler statements of some definitions and results. (To avoid assuming that T has EI one would have to work in $\mathbb{M}^{\mathrm{eq}}$ instead of $\mathbb{M}$.)

Consider a parameter set P in $\mathbb{M}$ that is A-invariant, that is, $f(P) = P$ for all $f \in \mathrm{Aut}(\mathbb{M}|A)$. (We do not assume P is small.) We say that P is *one-based over* A if for each P-tuple a and each $B \supseteq A$ such that $\mathrm{tp}(a|B)$ is stationary we have $\mathrm{cb}(a|B) \subseteq \mathrm{acl}(Aa)$. Clearly, if P is one-based over A and $A \subseteq A'$, then P is one-based over A'.

Lemma 14.2. *The following are equivalent:*

(1) *P is one-based over A;*
(2) *for each P-tuple a and each $B \supseteq A$ we have*

$$a \mathop{\smile\!\!\!\!\!\!\,|}_{C} B, \qquad \textit{where } C := \mathrm{acl}(Aa) \cap \mathrm{acl}(B).$$

Proof. Assume (1), let a be a P-tuple, and let $B \supseteq A$. Since $\mathrm{tp}(a|\,\mathrm{acl}(B))$ is stationary, we have $\mathrm{cb}(a|\,\mathrm{acl}(B)) \subseteq C := \mathrm{acl}(Aa) \cap \mathrm{acl}(B)$. Hence $\mathrm{tp}(a|\,\mathrm{acl}(B))$ does not fork over C, that is, $a \mathop{\smile\!\!\!\!\!\!\,|}_{C} \mathrm{acl}(B)$, and thus $a \mathop{\smile\!\!\!\!\!\!\,|}_{C} B$.

The converse follows by similar reasoning. $\square$

Lemma 14.3. *If P is one-based over A, then* $\mathrm{acl}(P)$ *is also one-based over A. If $A' \supseteq A$ and P is one-based over A', then P is one-based over A.*

Proof. Assume P is one-based over A. Let a be an $\mathrm{acl}(P)$-tuple such that $\mathrm{tp}(a|B)$ is stationary, where $B \supseteq A$. To obtain $\mathrm{cb}(a|B) \subseteq \mathrm{acl}(Aa)$ we can assume B is algebraically closed, replacing B by $\mathrm{acl}(B)$ if necessary. Take a P-tuple b such that a is b-algebraic. We can arrange that $b \mathop{\smile\!\!\!\!\!\!\,|}_{Aa} B$, by replacing b by a realization of a nonforking extension of $\mathrm{tp}(b|Aa)$ to AaB. Then $C := \mathrm{acl}(Ab) \cap B \subseteq \mathrm{acl}(Aa)$ by Lemma 11.8, so $C = \mathrm{acl}(Aa) \cap B$. By the lemma above, $b \mathop{\smile\!\!\!\!\!\!\,|}_{C} B$, hence $a \mathop{\smile\!\!\!\!\!\!\,|}_{C} B$, so $\mathrm{cb}(a|B) \subseteq C \subseteq \mathrm{acl}(Aa)$.

Next, assume P is one-based over A', where $A' \supseteq A$. Let a be a P-tuple and let $B \supseteq A$ be such that $\mathrm{tp}(a|B)$ is stationary. We can arrange that B is algebraically closed, and that $A' \mathop{\smile\!\!\!\!\!\!\,|}_{A} Ba$, so $A' \mathop{\smile\!\!\!\!\!\!\,|}_{AB} a$, that is, $A' \mathop{\smile\!\!\!\!\!\!\,|}_{B} a$, so $a \mathop{\smile\!\!\!\!\!\!\,|}_{B} A'$, that is, $\mathrm{tp}(a|A'B)$ does not fork over B. So $\mathrm{cb}(a|B) = \mathrm{cb}(a|A'B)$. By assumption, $\mathrm{cb}(a|A'B) \subseteq \mathrm{acl}(A'a)$, hence $\mathrm{cb}(a|B) \subseteq \mathrm{acl}(A'a) \cap \mathrm{acl}(B)$. Since $B \mathop{\smile\!\!\!\!\!\!\,|}_{Aa} A'a$, it follows that $\mathrm{cb}(a|B) \subseteq \mathrm{acl}(Aa)$. $\square$

Lemma 14.4. *Suppose Ω_A is modular. Then Ω is one-based over A.*

Proof. Let $a \in \Omega^n$ and let $B \supseteq A$ be such that $\mathrm{tp}(a|B)$ is stationary. Take b such that $\mathrm{cb}(a|B) = \mathrm{dcl}(b)$. By Lemma 11.17 the tuple b is definable over a tuple $c \in \Omega^m$, and we arrange in the usual way that $c \underset{Ab}{\perp} a$, so $a \underset{Ab}{\perp} c$. Note that then

$$\mathrm{cb}(a|B) = \mathrm{cb}(a|Ab) = \mathrm{cb}(a|Abc) = \mathrm{cb}(a|Ac).$$

By the modularity assumption and Lemma 14.1 we have

$$a \underset{C}{\perp} c, \qquad \text{where } C := A(\mathrm{cl}_A(a) \cap \mathrm{cl}_A(c)).$$

Note that $Ac \subseteq Cc \subseteq \mathrm{acl}(Ac)$, so

$$\mathrm{cb}(a|Ac) = \mathrm{cb}(a|Cc) = \mathrm{cb}(a|C) \subseteq \mathrm{dcl}(C) \subseteq \mathrm{acl}(Aa),$$

hence $\mathrm{cb}(a|B) \subseteq \mathrm{acl}(Aa)$. $\square$

Linearity. We say that Ω is A-*linear* if for all $\xi, \eta \in \Omega$ and each algebraically closed $B \supseteq A$ such that $\mathrm{MR}((\xi,\eta)|B) = 1$ there is a c such that $\mathrm{MR}(c|A) \le 1$ and $\mathrm{tp}((\xi,\eta)|B)$ is defined over c.

Note that if Ω is A-linear, $\xi, \eta \in \Omega$, $B \supseteq A$ is algebraically closed, $\mathrm{MR}((\xi,\eta)|B) = 1$, and $\mathrm{cb}((\xi,\eta)|B) = \mathrm{dcl}(c)$, then necessarily $\mathrm{MR}(c|A) \le 1$. Note also that if Ω is A-linear and $A \subseteq A'$, then Ω is A'-linear.

Lemma 14.5. *The following are equivalent:*

(1) *Ω is A-linear;*
(2) *for all $\xi, \eta \in \Omega$ and each algebraically closed $B \supseteq A$ there is a c such that $\mathrm{tp}((\xi,\eta)|B)$ is definable over c and c is algebraic over $A\xi\eta$.*

Proof. Assume Ω is A-linear, let $\xi, \eta \in \Omega$, and let $B \supseteq A$ be algebraically closed. The type $\mathrm{tp}((\xi,\eta)|B)$ is stationary, so it has a unique global nonforking extension $\mathbf{p}$. Take c such that $\mathrm{dcl}(c) = \mathrm{cb}(\mathbf{p})$. Thus $\mathbf{p}$ and its restriction $\mathrm{tp}((\xi,\eta)|B)$ are definable over c. We shall prove that c is algebraic over $A\xi\eta$.

If $\mathrm{MR}((\xi,\eta)|B) = 2$, then $\mathrm{MR}((\xi,\eta)|A) = 2$, so $\mathbf{p}$ does not fork over A, hence c is algebraic over A by Corollary 10.13(1). If $\mathrm{MR}((\xi,\eta)|B) = 0$, then (ξ,η) is the only realization of $\mathrm{tp}((\xi,\eta)|B)$, so c is interdefinable with (ξ,η). Suppose $\mathrm{MR}((\xi,\eta)|B) = 1$, the remaining case. Then by the A-linearity of Ω we have $\mathrm{MR}(c|A) \le 1$.

If $\mathrm{MR}((\xi,\eta)|A) = 1$, then again c is algebraic over A. So we can assume $\mathrm{MR}((\xi,\eta)|A) = 2$. Then $(\xi,\eta) \underset{A}{\not\perp} B$, hence $(\xi,\eta) \underset{A}{\not\perp} c$: otherwise, $\mathrm{MR}((\xi,\eta)|c) \ge 2$, but $\mathbf{p}$ is a nonforking extension of $\mathrm{tp}((\xi,\eta)|c)$, so $\mathrm{MR}(\mathbf{p}) = 2$, hence $\mathrm{MR}((\xi,\eta)|B) = 2$, a contradiction. Therefore, c is not algebraic over A, and thus $\mathrm{MR}(c|A) = 1$. Also $c \underset{A}{\not\perp} (\xi,\eta)$ by symmetry, so $\mathrm{MR}(c|A) > \mathrm{MR}(c|A\xi\eta$, so $\mathrm{MR}(c|A\xi\eta) = 0$, hence c is algebraic over $A\xi\eta$.

For the converse, assume (2), let $\xi, \eta \in \Omega$, and let $B \supseteq A$ be a parameter set in $\mathbb{M}$ such that $\mathrm{MR}((\xi,\eta)|B) = 1$. The type $\mathrm{tp}((\xi,\eta)|B)$ is stationary, so it has a unique global nonforking extension $\mathbf{p}$. Take c such that $\mathrm{dcl}(c) = \mathrm{cb}(\mathbf{p})$. Then $\mathbf{p}$ and its restriction $\mathrm{tp}((\xi,\eta)|B)$ are definable over c, so c is algebraic over B. By (2) the tuple c is also algebraic over $A\xi\eta$, so $\mathrm{MR}(c|A) \le \mathrm{MR}((\xi,\eta)|A) \le 2$. We shall prove that $\mathrm{MR}(c|A) \le 1$. Assume towards a contradiction that $\mathrm{MR}(c|A) = \mathrm{MR}((\xi,\eta)|A) = 2$. By Lemma 11.11,

$$\mathrm{MR}((\xi,\eta,c)|A) = \mathrm{MR}((\xi,\eta)|Ac) + \mathrm{MR}(c|A) = \mathrm{MR}(c|A\xi\eta) + \mathrm{MR}((\xi,\eta)|A),$$

so $\mathrm{MR}((\xi,\eta)|Ac) = \mathrm{MR}(c|A\xi\eta) = 0$. But c is algebraic over B, so

$$\mathrm{MR}((\xi,\eta)|Ac) \ge \mathrm{MR}((\xi,\eta)|B) \ge 1,$$

a contradiction. □

Lemma 14.6. *Suppose Ω is A-linear. Then Ω_A is locally modular.*

Proof. Take $a_0 \in \Omega \setminus \mathrm{cl}_A\,\emptyset$. We claim that $\Omega_A|\{a_0\}$ is modular. By 12.10 is is enough to show:

Let $a_1, a_2 \in \Omega$ and $\mathrm{rk}_A(\{a_1,a_2\}|\{a_0\}) = 2$, and let $F \subseteq \Omega$ be closed in Ω_A of finite rank with $a_0 \in F$ and $\mathrm{rk}_A(\{a_1,a_2\}|F) = 1$. Then

$$\mathrm{cl}_A\{a_0,a_1,a_2\} \cap F \ne \mathrm{cl}_A\{a_0\}.$$

By Lemma 11.17 we can take c such that $\mathrm{dcl}(c) = \mathrm{cb}((a_1,a_2)|\,\mathrm{acl}(AF))$, so c is algebraic over AF. By A-linearity and the previous lemma, c is also algebraic over Aa_1a_2. Note that $\mathrm{MR}((a_1,a_2)|Ac) = 1$ and a_i is not algebraic over Ac for $i = 1,2$. So a_2 is algebraic over Aca_1. Also, a_0 is not algebraic over Aa_1a_2, so a_0 is not algebraic over Ac. Hence a_0 and a_1 are conjugate over Ac, so (a_0,c) and (a_1,c) are conjugate over A. Take $a_2' \in \Omega$ such that (a_0,a_2',c) and (a_1,a_2,c) are conjugate over A. Then a_2' is algebraic over Aca_0, so $a_2' \in F$ and $a_2' \notin \mathrm{cl}_A\{a_0\}$. Since $a_2' \in \mathrm{cl}_A\{a_0,a_1,a_2\}$, this completes the proof. □

We can summarize most of the above as follows.

Theorem 14.7. *The following conditions on Ω, A are equivalent:*

1. *there is $B \supseteq A$ such that Ω_B is modular;*
2. *Ω is one-based over A;*
3. *Ω is A-linear;*
4. *Ω_A is locally modular.*

Proof. The direction (1) $\Longrightarrow$ (2) follows from Lemmas 14.3 and 14.4. The direction (2) $\Longrightarrow$ (3) follows from Lemma 14.5. The direction (3) $\Longrightarrow$ (4) is Lemma 14.6. For (4) $\Longrightarrow$ (1), assume Ω_A is locally modular. Take $a \in \Omega$ such that $\Omega_A|\{a\}$ is modular. Let $B := Aa$, and note that the closure operations of Ω_B and $\Omega_A|\{a\}$ coincide. Thus $\Omega_B = \Omega_A|\{a\}$ is modular. □

Condition (1) in this theorem raises the question whether Ω_B is always of the form $\Omega_A|E$ for a suitable $E \subseteq \Omega$. An obvious candidate to try is $E = \mathrm{cl}_B\,\emptyset$.

Linearity and plane curves. A *family of curves in Ω^2* is a pair (X, C) where $X \subseteq \mathbb{M}_x$ and $C \subseteq X \times \Omega^2$ are definable such that $C(a)$ is strongly minimal for all $a \in X$.

Let (X, C) be a family of curves in Ω^2. We think of Ω^2 as the plane, and of each section $C(a) \subseteq \Omega^2$ as a *plane curve*. Note that for all $a, b \in X$, either $C(a) \cap C(b)$ is finite, or $C(a) \triangle C(b)$ is finite. Thus by saturation, there is a natural number N such that for all $a, b \in X$, either $|C(a) \cap C(b)| \leq N$, or $|C(a) \triangle C(b)| \leq N$. In the former case we wish to consider $C(a)$ and $C(b)$ as essentially different curves, and in the latter case as essentially the same curve. How many essentially different curves are there in this family? To make sense of this question we introduce the (definable) equivalence relation $\sim$ on X by:

$$a \sim b \iff C(a) \triangle C(b) \text{ is finite.}$$

Since T has EI we have a definable surjective map $f : X \to X'$ with X' a definable set in $\mathbb{M}$, such that $\sim$ is the kernel of f. Thus $\mathrm{MR}(X')$ does not depend on the choice of f, X'. We call $\mathrm{MR}(X')$ the *essential Morley rank* of (X, C) and think of it as a rough measure of how many essentially different curves there are in the family.

Examples. In these examples we just consider one-sorted $\mathbb{M}$.

(1) Let $\mathbb{M} = \Omega$ be an infinite set (no further structure), and let $C \subseteq \Omega \times \Omega^2$ be given by $C = \{(b, \xi, \eta) \in \Omega^3 : \eta = b\}$, so each section $C(b)$ is a "horizontal" line. Then (Ω, C) is a family of curves in Ω^2, the equivalence relation $\sim$ on the parameter space Ω is just equality, so the essential Morley rank of the family is 1.

(2) Let $\mathbb{M}$ be an infinite vector space over a division ring $\mathbf{k}$ with underlying set Ω, fix a scalar $\lambda \in \mathbf{k}^\times$, and let $C \subseteq \Omega \times \Omega^2$ be given by $C = \{(b, \xi, \eta) \in \Omega^3 : \eta = \lambda\xi + b\}$. Then (Ω, C) is a family of curves in Ω^2, the equivalence relation $\sim$ on the parameter space Ω is just equality, so the essential Morley rank of the family is 1.

(3) Let $\mathbb{M}$ be an algebraically closed field with underlying set Ω (so $\mathbb{M}$, being large, is actually of infinite transcendence degree). Then

$$C := \{(a, b, \xi, \eta) \in \Omega^2 \times \Omega^2 : \eta = a\xi + b\}.$$

Then (Ω^2, C) is a family of curves in Ω^2, the equivalence relation $\sim$ on the parameter space Ω^2 is just equality, so the essential Morley rank of the family is 2.

The significance of A-linearity of Ω is that there are few curves in Ω^2. The following result makes one direction of this precise. (There is also a sort of converse.)

Proposition 14.8. *Suppose Ω is A-linear. Then every family of curves in Ω^2 has essential Morley rank ≤ 1.*

Proof. Let (X, C) be a family of curves in Ω^2, $X \subseteq \mathbb{M}_x$. To keep notations simple we assume that Ω, X and C are all 0-definable. (Why is this no loss of generality?) Let $f : X \to X'$ be a 0-definable map with kernel $\sim$ as above. Let $a \in X$, and let (ξ, η) be a generic point of $C(a)$ over a, so $\mathrm{MR}((\xi, \eta)|a) = \mathrm{MR}((\xi, \eta)|\operatorname{acl}(a)) = 1$. Let $\mathrm{cb}((\xi, \eta)|\operatorname{acl}(a)) = \mathrm{dcl}(c)$. Then

$\mathrm{MR}(c|0) \leq 1$ by linearity. Since $\mathrm{tp}((\xi, \eta)|\,\mathrm{acl}(a))$ is definable over c, there is an L_c-formula $\phi(x)$ such that for all $b \in X$ we have:

$$C(b) \in \mathrm{tp}((\xi, \eta)|\,\mathrm{acl}(a)) \iff \models \phi(b).$$

The set $C(a)$ determines the type $\mathrm{tp}((\xi, \eta)|\,\mathrm{acl}(a))$, in the sense of Section 9, so for all $b \in X$ we have:

$$C(b) \in \mathrm{tp}((\xi, \eta)|\,\mathrm{acl}(a)) \iff b \sim a \iff f(b) = f(a).$$

So $f(a)$ is the unique element of X' equal to $f(b)$ for some b such that $\models \phi(b)$. Hence $f(a)$ is c-definable, so $\mathrm{MR}(f(a)|0) \leq \mathrm{MR}(c|0) \leq 1$. Since $a \in X$ was arbitrary, this yields $\mathrm{MR}(X') \leq 1$. $\square$