Topics in Stability Theory

Anand Pillay and friends

Fall 2018

Contents

1 Introduction

These are notes from the course ”Topics in Stability Theory” given at Notre Dame by Anand Pillay in the fall of 2018. They are the result of the collaborative work of all students enrolled in the class. We will assume familiarity with basic model theory, for which [TZ12] is a good reference.

The course first covered, in the preliminaries section, the tools required to develop stability theory, such as indiscernibles, imaginaries... A brief incursion into continuous logic was also on the program. After this, in the typical modern fashion, local stability was developed, and consequences for stable theories were then discussed.

The course then became more expository, starting with an survey of classification theory, and different strengthening of stability. Followed an exploration of Zilber's conjecture and related algebraic issues. Finally, after a brief introduction to Keisler measures, the stable regularity lemma was proven, using the tools of stability theory.

2 Preliminaries

The notations for this class follow Pillay's model theory notes [Pil02]. Although most of them are standard, we will start these notes by reminding them to the reader :

The letter T will always denote a complete 1-sorted first order theory, in a language L. Models of said theory will be $M, N \models T$, and their subset will be $A, B, \cdots$. Finite tuples of some models are denoted $\bar{a}, \bar{b} \cdots$. If $A \subset M$, then by $\mathrm{Th}(M, A)$ we mean the complete theory of M with a constant symbol for each element of A .

As is the usual in model theory, we will fix a large cardinal $\bar{\kappa}$, and a $\bar{\kappa}$-saturated and strongly $\bar{\kappa}$-homogeneous model $\overline{M}$ of T. We will sometimes refer to it as the monster model. All sets of parameters considered will be of size strictly smaller than $\bar{\kappa}$.

Note. With the above notation, any model of cardinality $\leq \bar{\kappa}$ will be isomorphic to an elementary substructure of $\bar{M}$.

If $\Sigma(\bar{y})$ denotes a set of L_A-formulas, and $\varphi(\bar{y})$ denotes an L_A-formula, then $\Sigma(\bar{y}) \models \varphi(\bar{y})$ means $\forall \bar{b} \in \bar{M}$, if $\bar{M} \models \Sigma(\bar{b})$, then $\bar{M} \models \varphi(\bar{b})$.

Remark 2.1. If M is $\bar{\kappa}$-saturated and $\Sigma(\bar{y}) \models \varphi(\bar{y})$, then there exists a finite set of formula $\Sigma'(\bar{y}) \subset \Sigma(\bar{y})$ such that $\Sigma'(\bar{y}) \models \varphi(\bar{y})$.

Proof. Otherwise every finite subset of $\Sigma(\bar{y}) \cup \{\neg\varphi(\bar{y})\}$ will be consistent. By compactness, this set of formulas would then be consistent. Saturation of M yields that the set of formulas is realized in M, contradicting $\Sigma(\bar{y}) \models \varphi(\bar{y})$. □

A subset X of $\bar{M}$ is said to be A-definable in $\bar{M}$ if $\exists\varphi(\bar{x}) \in L_A$ such that $X = \{\bar{b} \in \bar{M} : \bar{M} \models \varphi(\bar{b})\}$. Finally, If $A \subseteq M$, and $\bar{b} \in M$, then the type of $\bar{b}$ in M in L_A is the set of formulas $\{\varphi(\bar{x}), \bar{M} \models \varphi(\bar{b})\}$, and is denoted $\mathrm{tp}_{\bar{M}}(\bar{b}/A)$.

2.1 Indiscernibles

Definition 2.2 (Indiscernible)**.** Let $(I, <)$ be a totally ordered set and $(\bar{a}_i : i \in I)$ a sequence of finite tuples from M of the same length. We say $(\bar{a}_i : i \in I)$ is *indiscernible* in M if for each $i_1 < \ldots < i_n$ and $j_1 < \ldots < j_n \in I$,

$$\mathrm{tp}_{\bar{M}}(\bar{a}_{i_1}, \ldots, \bar{a}_{i_n}) = \mathrm{tp}_{\bar{M}}(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}).$$

Equivalently, $\forall\varphi(\bar{x}_1, \ldots, \bar{x}_n) \in L$,

$$\bar{M} \models \varphi(\bar{a}_{i_1}, \ldots, \bar{a}_{i_n}) \quad \text{iff} \quad \bar{M} \models \varphi(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}).$$

Note. The above definition can be extended to make sense for any model M and also for formulas over $A \subseteq M$.

Example. Let $M = (\mathbb{Q}, <)$, and $I = (\mathbb{Q}, <)$. Then $(q, q \in \mathbb{Q})$ is indiscernible by quantifier elimination for dense linear orders.

Theorem 2.3. *Compactness lets us "stretch" indiscernibles. Formally, let $(a_i : i \in \omega)$ be indiscernible in $\bar{M}$, and $(I, <)$ an ordering of cardinality smaller than $\bar{\kappa}$. Then there exists an indiscernible $(b_i : i \in I)$ in $\bar{M}$ such that $\forall i_1 < \ldots < i_n \in I$,*

$$\mathrm{tp}_{\bar{M}}(\bar{a}_1, \ldots, \bar{a}_n) = \mathrm{tp}_{\bar{M}}(\bar{b}_{i_1}, \ldots, \bar{b}_{i_n}).$$

Proof. Introduce new constant symbols $c_i, i \in I$ and let

$$\Sigma = \{\varphi(c_1, \ldots, c_n) : \bar{M} \models \varphi(a_1, \ldots, a_n)\},$$

$$\Sigma' = \{\varphi(c_1, \ldots, c_n) \leftrightarrow \varphi(c_{i_1}, \ldots, c_{i_n}) : i_1 < \ldots < i_n \in I\}.$$

Then $\Sigma \cup \Sigma' \cup T$ is consistent by compactness, i.e. has a model. Take a such model, of cardinality $\leq \bar{\kappa}$. By $\bar{\kappa}$-saturation and strong $\bar{\kappa}$-homogeneity, such a model is isomorphic to an elementary substructure of $\bar{M}$. Then

$$\begin{aligned} \bar{M} \models \varphi(\bar{a}_1, \ldots, \bar{a}_n) &\leftrightarrow \bar{M} \models \varphi(\bar{c}_1, \ldots, \bar{c}_n) && \text{(by } \Sigma'\text{)} \\ &\leftrightarrow \bar{M} \models \varphi(\bar{c}_{i_1}, \ldots, \bar{c}_{i_n}). && \text{(by } \Sigma\text{)} \end{aligned}$$

$\square$

Indiscernible sequences are a fundamental tool of model theory, and there are many ways to obtain them. In what follows, we will discuss three such methods : Ramsey's theorem, coheirs, and the Erdös-Rado theorem.

Fact (Ramsey, extended)**.** *Let $n_1, \ldots, n_r < \omega$. For each $i = 1, \ldots, r$, let $X_{i,1}$, $X_{i,2}$ be a partition of $\omega^{[n_i]}$, the set of n_i-element subsets of $\mathbb{N}$. Then there is an infinite subset $Y \subseteq \omega$ which is homogeneous, i.e. $\forall i = 1, \ldots, r$, either $Y^{[n_i]} \subseteq X_{i,1}$ or $Y^{[n_i]} \subseteq X_{i,2}$.*

We can apply Ramsey's theorem to obtain indiscernible sequences.

Definition 2.4. Let $A \subseteq \bar{M}$ be small and $\Sigma(\bar{x})$ be a collection of L_A-formulas in the variable $\bar{x}$. We say "$\Sigma(\bar{x})$ is consistent" if $\Sigma(\bar{x}) \cup \mathrm{Th}(\bar{M}, A)$ is consistent.

Fact 2.5. $\Sigma(\bar{x})$ *is consistent iff every finite subset* $\Sigma'(\bar{x}) \subseteq \Sigma(\bar{x})$ *is realized in* $\bar{M}$.

Proof. If $\Sigma(\bar{x})$ is consistent, then by saturation it is realized in $\bar{M}$, and in particular every finite subset $\Sigma'(\bar{x}) \subseteq \Sigma(\bar{x})$ will also be realized in $\bar{M}$. For the converse, if every finite subset is realized, every finite subset is consistent. Then from compactness Σ is consistent. □

Proposition 2.6. *For each* $n \in \omega$*, let* $\Sigma_n(x_1, \ldots, x_n)$ *be a collection of* L*-formulas in variables* $x_1, \ldots, x_n$*. Suppose that there are* $a_1, a_2, \dots \in \bar{M}$ *such that*

$$\bar{M} \models \Sigma_n(a_{i_1}, \ldots, a_{i_n}), \quad \forall i_1 < \ldots < i_n < \omega.$$

Then there exists an indiscernible $(b_i : i \in \omega)$ *in* $\bar{M}$ *such that*

$$\bar{M} \models \Sigma_n(b_{i_1}, \ldots, b_{i_n}), \quad \forall i_1 < \ldots < i_n < \omega.$$

Example. Suppose $\Sigma_2 = \{x_1 \neq x_2\}$. Then the previous proposition yields the existence of infinite indiscernible sequences.

Proof. (of Proposition 2.6) Consider the following set of L-formulas

$$\begin{aligned}\Gamma(x_1, x_2, \ldots) = \{&\varphi(x_{i_1}, \ldots, x_{i_n}) \leftrightarrow \varphi(x_{j_1}, \ldots, x_{j_n}) : \\ &i_1 < \ldots < i_n,\ j_1 < \ldots < j_n \in \omega, \varphi \in L\} \\ &\cup \bigcup_n \Sigma_n(x_1, \ldots, x_n).\end{aligned}$$

By Fact 2.5, it is enough to prove that every finite subset of Γ is consistent. Let Γ' be a finite subset of Γ. By choosing n large enough, we can assume the only variables appearing in Γ' are $x_1, \ldots, x_n$. Let $\varphi_1, \ldots, \varphi_r$ be the L-formulas appearing in Γ'. For $i = 1, \ldots, r$, let

$$\begin{aligned}X_{i,1} &= \{(j_1, \ldots, j_n) : \quad j_1 < \ldots < j_n \in \omega, \quad \bar{M} \models \varphi_i(a_{j_1}, \ldots, a_{j_n})\}, \\ X_{i,2} &= \{(j_1, \ldots, j_n) : \quad j_1 < \ldots < j_n \in \omega, \quad \bar{M} \models \neg\varphi_i(a_{j_1}, \ldots, a_{j_n})\}.\end{aligned}$$

By Ramsey's theorem, there exists an infinite $Y \subseteq \mathbb{N}$ such that $\forall i = 1, \ldots, r$, $Y^{[n_i]}$ is either contained in $X_{i,1}$ or in $X_{i,2}$. Write $Y = \{k_1 < k_2 < \ldots\}$. Interpret each x_i as a_{k_i} to satisfy Γ'. □

Note. Proposition 2.6 also works when we consider tuples instead of elements, and when we consider formulas over some small fixed $A \subset \bar{M}$.

Another to obtain indiscernibles is the use of special types, called coheirs. We will define them now. First recall that if A is any set of parameters, then $S_{\bar{x}}(A)$ denotes the class of complete types over A.

Definition 2.7. Let $M \prec N \prec \bar{M}$ be models, and $p(\bar{x}) \in S_{\bar{x}}(N)$. We say p is finitely satisfiable in M, or $p(\bar{x})$ is a *coheir* of $p \restriction M \in S_{\bar{x}}(M)$, if every $\varphi(\bar{x}) \in p(\bar{x})$ is satisfied by some $\bar{a} \in M$, or equivalently, every finite subset of $p(\bar{x})$ is realized by some $\bar{a} \in M$.

Note. Any $\varphi(\bar{x}) \in p$ is realized in N as $\varphi(\bar{x}) \in L_N$ and $N \prec \bar{M}$ and $\varphi(\bar{x})$ is realized in $\bar{M}$, but there is no reason to expect φ to be realized in M.

Remark. $p(\bar{x}) \in S_{\bar{x}}(N)$ is finitely satisfiable (f.s.) in M if and only if $p(\bar{x})$ is in the topological closure of $\{\mathrm{tp}(\bar{a}/N) : \bar{a} \in M\}) \subseteq S_{\bar{x}}(N)$.

Proof. Suppose first that $p(\bar{x})$ is finitely satisfiable. Let $\varphi(\bar{x}) \in p(\bar{x})$, and let $[\varphi(\bar{x})]$ be the corresponding open set. Then φ is realized in M, say by $\bar{a} \in M$. Hence $\mathrm{tp}(\bar{a}/N) \in [\varphi(\bar{x})]$.

For the other direction, assume $p(\bar{x})$ is in the closure of $\{\mathrm{tp}(\bar{a}/N) : \bar{a} \in M\}) \subseteq S_{\bar{x}}(N)$. Let $\varphi(\bar{x}) \in p(\bar{x})$, then by assumption, there is $\bar{a} \in M$ such that $\mathrm{tp}(\bar{a}/N) \in \varphi(\bar{x})$, so we obtain $M \models \varphi(\bar{a})$.

□

Lemma 2.8. *Suppose $p(\bar{x}) \in S_{\bar{x}}(M)$ and $M \prec N$. Then there is $p'(\bar{x}) \in S_{\bar{x}}(N)$ such that $p \subseteq p'$ and p' is finitely satisfiable in M.*

Proof. Consider $\Gamma(\bar{x}) = p(\bar{x}) \cup \{\neg\varphi(\bar{x}) : \varphi(\bar{x}) \in L_N$ and not realized in $M\}$. It is enough to show that $\Gamma(\bar{x})$ is consistent because any extension of such Γ to a complete type $p'(\bar{x}) \in S_{\bar{x}}(N)$ will work. Let Γ' be a finite subset of Γ, written in the form $\Gamma' = \{\Psi(\bar{x}), \neg\varphi_1(\bar{x}), \dots, \neg\varphi_r(\bar{x})\} \in p$. Then any solution $\bar{a}$ of Ψ in M satisfies Γ', as the $\varphi_1, \cdots, \varphi_r$ are not realized in M. □

Remark. The following is a proof of the above fact using analysis. Let i_M denote the map from $M^{\bar{x}}$ to $S_{\bar{x}}(M)$ such that $m \mapsto \mathrm{tp}(m/M)$. We define $i_N : M^{\bar{x}} \to S_{\bar{x}}(N)$ similarly. Let r denote the restriction map from $S_{\bar{x}}(N) \to S_{\bar{x}}(M)$. Note that $r \circ i_N = i_M$, and the set of types in $S_{\bar{x}}(N)$ that are finitely satisfiable in M is exactly the closure of $i_N(M^{\bar{x}})$ in $S_{\bar{x}}(N)$. Hence its image under the restriction map is closed. However the image must contain $i_M(M^{\bar{x}})$, which is dense in $S_{\bar{x}}(M)$. Therefore it must be onto, which proves the desired result.

Note. Lemma 2.8 also works when we consider formulas over some $A \subset M$, i.e. if $A \subset M$ and $p(\bar{x}) \in S_{\bar{x}}(M)$ is finitely satisfiable in A, then for any model $N \succ M$, there is $p'(\bar{x}) \in S_{\bar{x}}(N)$ such that $p \subseteq p'$ and p' is finitely satisfiable in A.

Proposition 2.9. *Let $p(\bar{x}) \in S_{\bar{x}}(M)$, $N \succ M$ be $|M|^+$-saturated, and $p'(\bar{x}) \in S_{\bar{x}}(N)$ a coheir of p. Let $\bar{a}_1, \bar{a}_2, \ldots \in N$ be defined as follows:*

$\bar{a}_1$ *realizes* $p(\bar{x})$,
$\bar{a}_2$ *realizes* $p'(\bar{x}) \restriction (M, \bar{a}_1)$,
$\bar{a}_3$ *realizes* $p'(\bar{x}) \restriction (M, \bar{a}_1, \bar{a}_2)$,
$\ldots$.

Then $(\bar{a}_i : i \in \omega)$ is indiscernible over M.

Proof. We prove by induction on k that for any $n \leq k$ and $i_1 < \ldots < i_n \leq k$ and $j_1 < \ldots < j_n \leq k$, we have

$$\mathrm{tp}_M(\bar{a}_{i_1}, \ldots, \bar{a}_{i_n}/M) = \mathrm{tp}_M(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}/M).$$

Assume this is true for k, and consider $k+1$. Let $i_i < \ldots < i_n \leq k$, $j_1 < \ldots < j_n \leq k$. We need to show that

$$\mathrm{tp}_M(\bar{a}_{i_1}, \ldots, \bar{a}_{i_n}, \bar{a}_{k+1}/M) = \mathrm{tp}_M(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}, \bar{a}_{k+1}/M).$$

Consider a formula $\varphi(\bar{x}_1, \ldots, \bar{x}_n, \bar{x}_{n+1}) \in L_M$. Assume by contradiction that

$$M \models \varphi(\bar{a}_{i_1}, \ldots, \bar{a}_{i_n}, \bar{a}_{k+1}) \ \wedge\ \neg\varphi(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}, \bar{a}_{k+1}).$$

But $\mathrm{tp}(\bar{a}_{k+1}/M, \bar{a}_1, \ldots, \bar{a}_k)$ is finitely satisfiable in M, so there is $\bar{a}' \in M$ such that

$$M \models \varphi(\bar{a}_{i_1}, \ldots, \bar{a}_{i_n}, \bar{a}') \ \wedge\ \neg\varphi(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}, \bar{a}'),$$

which contradicts the induction hypothesis. □

Finally, the Erdös-Rado theorem allows us to obtain indiscernible. We will not expand on this, but let us state the following :

Proposition 2.10. *Given T, there exists $\lambda \sim \beth_{(2^{|T|})^+}$ (assumed $< \bar{\kappa}$) such that if $(\bar{a}_i : i < \lambda)$ is a sequence of finite tuples in $\bar{M}$, then there exists an indiscernible sequence $(\bar{b}_i : i < \omega)$ in $\bar{M}$ (tuples of the same length as the $\bar{a}_i$) such that for all $n < \omega$, there are $\alpha_0 < \ldots < \alpha_n < \lambda$ satisfying*

$$\mathrm{tp}_{\bar{M}}(\bar{b}_0, \ldots, \bar{b}_n) = \mathrm{tp}_{\bar{M}}(\bar{a}_{\alpha_0}, \ldots, \bar{a}_{\alpha_n}).$$

2.2 Definability and Generalizations

Let $A \subseteq \bar{M}$. Recall an A-definable set in $\bar{M}$ is a subset $X \subseteq \bar{M}^n$ defined by an L_A-formula $\varphi(x_1, \ldots, x_n)$. We will call $\varphi(\bar{x}), \Psi(\bar{x}) \in L_{\bar{M}}$ equivalent if $\bar{M} \models (\forall \bar{x})\,(\varphi(\bar{x}) \leftrightarrow \Psi(\bar{x}))$. That is, the formulas Φ, Ψ define the same definable set.

Lemma 2.11. *Let $X \subseteq \bar{M}^n$ be definable (with parameters in $\bar{M}$). Then X is definable over A (i.e. X is definable by an L_A-formula) iff X is $\mathrm{Aut}(\bar{M}/A)$-invariant.*

Proof. Let us start with the forward implication. Given $\sigma \in \mathrm{Aut}(\bar{M}/A)$ and $\bar{a} \in A$,

$$\begin{aligned}
\sigma(X) &= \{\sigma(\bar{x}) : \varphi(\bar{x}, \bar{a})\} \\
&= \{\bar{x}' : \varphi(\bar{x}', \sigma(\bar{a}))\} \quad \text{(since } \varphi(\bar{x}, \bar{a}) \leftrightarrow \varphi(\sigma(\bar{x}), \sigma(\bar{a})) \text{ for automorphisms)} \\
&= \{\bar{x}' : \varphi(\bar{x}', \bar{a})\} \qquad \text{(since } \sigma(\bar{a}) = \bar{a}) \\
&= X.
\end{aligned}$$

To prove the converse, we need to use strong $\bar{\kappa}$-homogeneity of $\bar{M}$. Suppose X is defined by $\varphi(\bar{x}, \bar{b})$, where $\bar{b} \in M$ are parameters.

Claim. *Given $\bar{b}' \in \bar{M}$, if $\mathrm{tp}(\bar{b}'/A) = \mathrm{tp}(\bar{b}/A)$, then $\varphi(\bar{x}, \bar{b}')$ is equivalent to $\varphi(\bar{x}, \bar{b})$.*

Proof. Follows from strong $\bar{\kappa}$-homogeneity of $\bar{M}$, which yields some $\sigma \in \mathrm{Aut}(\bar{M}/A)$ such that $\sigma(\bar{b}) = \bar{b}'$.

$\square$

Let $p(\bar{y}) = \mathrm{tp}_{\bar{M}}(\bar{b}/A)$. The claim yields $p(\bar{y}) \models (\forall \bar{x})\ (\varphi(\bar{x}, \bar{y}) \leftrightarrow \varphi(\bar{x}, \bar{b}))$. Then compactness yields some $\psi(\bar{y}) \in p(\bar{y})$, $\psi \in L_A$ such that

$$\psi(\bar{y}) \models (\forall \bar{x})\ (\varphi(\bar{x}, \bar{y}) \leftrightarrow \varphi(\bar{x}, \bar{b})).$$

So X can be defined by the L_A-formula

$$\theta(\bar{x}) = (\exists \bar{y})\ (\varphi(\bar{x}, \bar{y}) \wedge \psi(\bar{y}))\ \in L_A.$$

$\square$

The equivalence of definablity and invariance under automorphisms can be generalized, as shown in the following definition and lemma.

Definition 2.12. $X \subseteq \bar{M}^n$ is definable almost over A if there is an A-definable equivalent relation E on $\bar{M}^n$ with finitely many classes and X is a union of some E-classes.

Lemma 2.13. *Let $X \subseteq \bar{M}$ be definable, and $A \subset \bar{M}$ be small. The following are equivalent :*

i X is definable almost over A.

ii $|\{\sigma(X) : \sigma \in \mathrm{Aut}(\bar{M}/A)\}| < \omega$.

iii $|\{\sigma(X) : \sigma \in \mathrm{Aut}(\bar{M}/A)\}| < \bar{\kappa}$.

Proof. ii $\Rightarrow$ iii is immediate.

i $\Rightarrow$ ii also follows from definition: Let $\varphi(\bar{x}_1, \bar{x}_2) \in L_A$ be the A-definable equivalence relation E, and let $\bar{b}_1, \ldots, \bar{b}_n \in \bar{M}$ be representatives in each equivalence class so that each class can be written as $[\bar{b}_i] = \{\bar{x} : \varphi(\bar{x}, \bar{b}_i)\}$. Given $\sigma \in \mathrm{Aut}(\bar{M}/A)$, since $\varphi(\bar{x}_1, \bar{x}_2) \leftrightarrow \varphi(\sigma(\bar{x}_1), \sigma(\bar{x}_2))$, the image of each $[\bar{b}_i]$ under σ will be

$$\sigma([\bar{b}_i]) = \{\sigma(\bar{x}) : \varphi(\bar{x}, \bar{b}_i)\} = \{\bar{x}' : \varphi(\bar{x}', \sigma(\bar{b}_i))\} = \{\bar{x} : \varphi(\bar{x}, \bar{b}_{j_i})\} = [\bar{b}_{j_i}],$$

for some $j_i \leq n$. Now X is a disjoint union of some $[\bar{b}_i]$'s, so $\sigma(X)$ is a disjoint union of some $[\bar{b}_j]$'s. Since there are only finitely many equivalence classes, there can only be finitely many possibilities for disjoint unions of these classes, which completes the proof.

ii $\Rightarrow$ i: Let $\varphi(\bar{x}, \bar{b}) \in L_A$ define X and let $p(y) = \mathrm{tp}(\bar{b}/A)$. From the above argument, given $\sigma \in \mathrm{Aut}(\bar{M}/A)$, we have $\sigma(X) = \{\bar{x}' : \varphi(\bar{x}', \sigma(\bar{b}))\}$. Then from assumption, there must be distinct $\bar{b} = \bar{b}_1, \ldots, \bar{b}_n \in \bar{M}$ so that

$$\{\sigma(X) : \sigma \in \mathrm{Aut}(\bar{M}/A)\} \; = \; \{\{\bar{x} : \varphi(\bar{x}, \bar{b}_i)\} : \; i \leq n\}.$$

Now if $\bar{b}' \in \bar{M}$ has the same type as $\bar{b}$, then strong $\bar{\kappa}$-homogeneity of $\bar{M}$ yields some $\sigma \in \mathrm{Aut}(\bar{M}/A)$ such that $\sigma(\bar{b}) = \bar{b}'$. Then the above argument again shows that $\varphi(\bar{x}, \bar{b}')$ defines $\sigma(X)$ for some $\sigma \in \mathrm{Aut}(\bar{M}/A)$. Thus

$$\{\bar{x} : \varphi(\bar{x}, \bar{b}')\} = \{\bar{x} : \varphi(\bar{x}, \bar{b}_i)\}$$

for some $i \leq n$. Therefore

$$p(\bar{y}) \models \bigvee_{i \leq n} (\forall \bar{x}) \; [\varphi(\bar{x}, \bar{y}) \leftrightarrow \varphi(\bar{x}, \bar{b}_i)].$$

Then compactness yields some $\psi(\bar{y}) \in p(\bar{y})$ such that

$$\psi(\bar{y}) \models \bigvee_{i \leq n} (\forall \bar{x}) \; [\varphi(\bar{x}, \bar{y}) \leftrightarrow \varphi(\bar{x}, \bar{b}_i)]. \tag{$*$}$$

Define the equivalence relation $E(\bar{x}_1, \bar{x}_2)$ by the formula

$$\theta(\bar{x}_1, \bar{x}_2) = (\forall \bar{y})[\psi(\bar{y}) \to (\varphi(\bar{x}_1, \bar{y}) \leftrightarrow \varphi(\bar{x}_2, \bar{y}))] \; \in L_A.$$

Now E has only finitely many equivalence classes from $(*)$, and X is the equivalence class represented by $\bar{b}$.

iii $\Rightarrow$ ii: Assume for contradiction that

$$|\{\sigma(X) : \sigma \in \mathrm{Aut}(\bar{M}/A)\}| = \lambda \geq \omega.$$

Using the same notation as in the proof of ii $\Rightarrow$ i, we can find λ-many elements $(\bar{b}_i : i < \lambda) \subset \bar{M}$ to represent the distinct images under automorphisms $\{\bar{x} : \varphi(\bar{x}, \bar{b}_i)\}$. Then the set of formulas

$$q(\bar{y}) = p(\bar{y}) \cup \{\neg(\forall \bar{x}) \; [\varphi(\bar{x}, \bar{y}) \leftrightarrow \varphi(\bar{x}, \bar{b}_i)] : \; i < \lambda\}$$

will be finitely satisfiable because $\lambda \geq \omega$. By $\bar{\kappa}$-saturation, $q(\bar{y})$ must be realized in $\bar{M}$ by some $\bar{b}' \in \bar{M}$. But such $\bar{b}'$ has the same type as $\bar{b}$ over A, and so strong $\bar{\kappa}$-homogeneity yields some $\sigma \in \mathrm{Aut}(\bar{M}/A)$ such that $\sigma(\bar{b}) = \bar{b}'$. Yet applying such σ on X gives the image

$$\sigma(X) = \{\sigma(\bar{x}) : \varphi(\bar{x}, \bar{b})\} = \{\bar{x}' : \varphi(\bar{x}', \sigma(\bar{b}))\} = \{\bar{x} : \varphi(\bar{x}, \bar{b}')\} = \{\bar{x} : \varphi(\bar{x}, \bar{b}_i)\}$$

for some $i < \lambda$, which contradicts $\bar{b}'$ satisfying the formula in variable $\bar{y}$

$$\neg(\forall \bar{x}) \; [\varphi(\bar{x}, \bar{y}) \leftrightarrow \varphi(\bar{x}, \bar{b}_i)].$$

$\square$

Recall that $S_{\overline{x}}(A)$ is the set of complete types in free variables $\overline{x}$ with parameters from A. (Equivalently, the complete types in the language L_A in free variables $\overline{x}$ extending the L_A-theory $Th(\overline{M})$.) Since $\overline{M}$ is saturated and A is small, this is also the set of all $\mathrm{tp}_{\overline{M}}(\overline{a}/A)$ for some $\overline{a} \in \overline{M}$.

The topology on $S_{\overline{x}}(A)$ is generated by the basic open sets

$$[\varphi(\overline{x})] = \{p(\overline{x}) \in S_{\overline{x}}(A) : \varphi(\overline{x}) \in p(x)\}$$

Since types are complete, these basic open sets are also closed because we have $S_{\overline{x}}(A) \setminus [\varphi(\overline{x})] = [\neg\varphi(\overline{x})]$. With this topology, $S_{\overline{x}}(A)$ is the Stone space of the Boolean algebra of L_A-formulas $\varphi(\overline{x})$, up to equivalence in $\overline{M}$. It is a compact, Hausdorff, totally disconnected space, e.g. a profinite space. (Not to be confused with a pseudofinite space.)

Proposition 2.14. *We can identify definable sets with continuous functions in a certain setting:*

1. *Formulas $\varphi(\overline{x})$, $\psi(\overline{x}) \in L_A$ are equivalent if and only if $[\varphi(\overline{x})] = [\psi(\overline{x})]$ in $S_{\overline{x}}(A)$.*

2. *The clopen subsets of $S_{\overline{X}}(A)$ are precisely the basic clopen sets.*

3. *Clopen subsets X of $S_{\overline{x}}(A)$ correspond exactly to continuous functions $f : S_{\overline{x}}(A) \to 2$, where $f(p(\overline{x})) = 1$ if $p(\overline{x}) \in X$ and 0 otherwise.*

4. *The definable subsets of $\overline{M}^n$ are in one-to-one correspondence with continuous functions from $S_{\overline{x}}(A)$ to 2.*

Proof. 1. Suppose $\varphi(\overline{x})$ and $\psi(\overline{x})$ are equivalent. Then $\overline{M} \models \forall\overline{x}\ \varphi(\overline{x}) \leftrightarrow \psi(\overline{x})$. In particular, $[\psi(\overline{x})] \subseteq [\varphi(\overline{x})]$ and $\varphi(\overline{x}) \subseteq [\psi(\overline{x})]$. Indeed, Suppose there was a type in one but not the other. Then by consistency, it is realized in the saturated model $\overline{M}$. That would imply that there is an $\overline{a} \in \overline{M}$ such that, for example, $\overline{M} \models \varphi(\overline{a}) \wedge \neg\psi(\overline{a})$, a contradiction.

   Suppose $[\varphi(\overline{x})] = [\psi(\overline{x})]$, but they are not equivalent. Then $\overline{M} \models \exists\overline{x}\ \varphi(\overline{x}) \wedge \neg\psi(\overline{x})$. Let $\overline{a} \in \overline{M}$ be a witness. Then $\mathrm{tp}_{\overline{M}}(\overline{a}/A) \in [\varphi(\overline{x})]$ but not is not in $[\psi(\overline{x})]$, a contradiction. Thus they must be equivalent in $\overline{M}$.

2. We only need to show that any clopen set is a basic clopen set. Let $X \subseteq S_{\overline{x}}(A)$ be a clopen set. Hence X and its complement are closed, so each is the intersection of some collection of basic clopen sets represented by sets P_0 and P_1 of L_A-formulas. Then $P_0 \cup P_1$ is inconsistent because

$$\bigcap_{\varphi(\overline{x}) \in P_0 \cup P_1} [\varphi(\overline{x})] = X \cap X^c = \emptyset$$

   Therefore, by compactness there are finite subsets $\{\varphi_0(\overline{x}), \ldots, \varphi_k(\overline{x})\} \subseteq P_0$ and $\{\psi_0(\overline{x}), \ldots, \psi_j(\overline{x})\} \subseteq P_1$ whose union is inconsistent. Therefore the L_A-formula $\tau(\overline{x})$ defined via

$$\bigwedge_{i \in k+1} \varphi_i(\overline{x}) \wedge \bigwedge_{i \in j+1} \neg\psi_i(\overline{x})$$

has $[\tau(\overline{x})] = X$, so X is a basic clopen set as desired.

3. Suppose $f : S_{\overline{x}}(A) \to 2$ is as stated in the proposition. Then $f^{-1}(\{1\})$ is open because 2 has the discrete topology and f is continuous. Similarly, $f^{-1}(\{0\}) = \overline{f^{-1}(\{1\})}$ is open, so $f^{-1}(\{1\})$ is clopen. Note that $f^{-1}(\{1\})$ trivially gives rise to f as described in the proposition.

   Now suppose X is a clopen set and define $f_X : S_{\overline{x}}(A) \to 2$ as described. Clearly f_X is continuous, as $f^{-1}(2) = S_{\overline{x}}(A)$, $f^{-1}(\emptyset) = \emptyset$, $f^{-1}(\{0\}) = X^c$, and $f^{-1}(\{1\}) = X$ are all open because X is clopen.

4. By the first part of the proposition, definable sets are in one-to-one correspondence with basic clopen subsets of $S_{\overline{x}}(A)$. By the second part, basic clopen sets are exactly all of the clopen subsets, so definable sets are in one-to-one correspondence with clopen sets. By the third part, clopen sets are in one-to-one correspondence with continuous functions $f : S_{\overline{x}}(A) \to 2$, so definable sets are in one-to-one correspondence with these continuous functions.

□

We can extend this idea to the general setting of continuous logic, which we shall introduce below. The above proposition shows that traditional logic is a special case of the following:

Definition 2.15.

1. Let C be a topological space. A C-valued formula in $\overline{x}$ over A is a continuous function $\varphi : S_{\overline{x}}(A) \to C$.

2. By a CL-formula (continuous logic) formula over A, we mean an $\mathbb{R}$-valued formula over A. Alternatively, the map can be into $[0, 1]$ or $\mathbb{R}_{\geq 0}$.

Remark 2.16.

- Suppose $\varphi : S_{\overline{x}}(A) \to C$ is a C-valued formula in $\overline{x}$ over A. Then φ gives rise to a map from $\overline{M}^n$ to C given by $f_\varphi(\overline{b}) = \varphi(\mathrm{tp}_{\overline{M}}(\overline{b}/A))$. We call f_φ an A-definable function from $\overline{M}^n$ to C.

- A definable function from $\overline{M}^n$ to C is an A-definable function from $\overline{M}^n$ to C for some small $A \subseteq \overline{M}$.

- For M an arbitrary model of T, a definable function from M^n to C is a function $f : M^n \to C$ which lifts to a continuous function from $S_{\overline{x}}(M)$ to C.

Recall that Lemma 2.11 showed that being a definable set and being invariant under automorphism are the same thing. This result adapts to continuous logic too.

Lemma 2.17. *Let f be a definable function from $\overline{M}^n$ to C. Let $A \subseteq \overline{M}$ be small. Then f is A-definable if and only if f is $Aut(\overline{M}/A)$-invariant, where the action of $Aut(\overline{M}/A)$ on f is $\sigma(f)(\bar{b}) = f(\sigma^{-1}(\bar{b}))$.*

Proof. Notice one direction is trivial. If f is A-definable, then f is $Aut(\overline{M}/A)$-invariant. We prove the other direction. Assume that f is $Aut(\overline{M}/A)$-invariant and definable. Then, it is B-definable for some small $B \subset \overline{M}$. Without loss of generality, we can assume that $A \subseteq B$. Let $f_B : S_{\overline{x}}(B) \to C$ be the map lifting f to $S_{\bar{x}}(B)$, that is, for all $p \in S_{\bar{x}}(B)$, we have $f_B(p) = f(a)$ for some (any) $a \models p$. Since f is B-definable, the map f_B is continuous.

Similarly, since f is A-invariant, the map $f_A : S_{\overline{x}}(A) \to C$ defined via $f_A(q) = f(c)$ where $c \models q$ is well-defined. Let $\pi_{B,A} : S_{\overline{x}}(B) \to S_{\overline{x}}(A)$ be the natural restriction map. Then, the following diagram commutes:

$$\begin{array}{ccc} S_{\overline{x}}(B) & & \\ \downarrow & \searrow & \\ S_{\overline{x}}(A) & \to & C \end{array}$$

The map $\pi_{B,A}$ is a topological quotient map, meaning that $\Omega \subset S_{\bar{x}}(B)$ is open if and only if $\pi_{B,A}^{-1}(\Omega)$ is open. Indeed, it is continuous, surjective and closed, as types spaces are compact Hausdorff.

By the universal property of quotient maps and continuity of f_b, we then obtain continuity of f_A. $\square$

Example 2.18. Ordinary first order formulas and definable sets are examples, with $C = \{0, 1\}$

Example 2.19. Consider the identity from $S_{\overline{x}}(A)$ to itself. Then this induces the tautological map that takes a tuple to its type, e.g. $f(\bar{b}) = \mathrm{tp}_{\overline{M}}(\bar{b}/A)$.

Example 2.20. Let $T = RCOF = Th(\mathbb{R}, +, x, 0, 1, -, <)$. This theory admits quantifier elimination, that is all formulas are equivalent to some quantifier free formula defining a semi-algebraic set. Recall that a semi-algebraic set is defined by a finite disjunction of formulas of the form $f(\overline{x}) = 0 \wedge g(\overline{x}) > 0$.

A real closed field is a model of T, equivalently any ordered field with the intermediate value property. In any such model M, the absolute value function on M is definable over the empty set in the usual first order sense of definability.

In the context of real closed field, we can also produce CL-formulas.

Example 2.21. Let $C = [-1, 1] \subseteq \mathbb{R}$ with the subspace topology, and let M be a model of T. Define the function $f : M \to C$ as follows:

If $|x|_M > 1$, then $f(x) = 0$. Otherwise, $f(x)$ is the unique element r of C such that for all $n \in \mathbb{N} \setminus \{0\}$ we have $|x - r|_M < \frac{1}{n}$.

We first claim that such an r exists and is unique. Assuming two such reals exist for a given x, equality follows immediately from the triangle inequality and the least upper bound property. Therefore we only need to show existence.

Suppose $|x|_M \leq 1$. Then since $M \models T$, $\mathbb{R}$ can be embedded in $\overline{M}$. Therefore, consider $\{r \in \mathbb{R} : r < x\}$. But $\mathbb{R}$ is complete as a metric space and $|x|_M| \leq 1$, so this set has a supremum r in $\mathbb{R}$. Notice that r has the required property: if not, there is some $n \in \mathbb{N} \setminus \{0\}$ with $|x - r|_M \geq \frac{1}{n}$, and there are two cases:

- $r < x$: Then $|x - r|_M = x - r \geq \frac{1}{n}$, so $r + \frac{1}{2n} < x$, contradicting the fact that r is the supremum of the aforementioned set.
- $x < r$: Then $|x - r|_M = r - x \geq \frac{1}{n}$, so $r - \frac{1}{2n} > x$, contradicting the fact that r is the supremum.

We still need to show that this map $f(x)$, which we shall now denote as $st(x)$ for the "standard part of x," is definable in any model $\mathcal{M}$ of T.

Note that any rational number q is definable over the empty set. Hence, for any type $p(x) \in S_1(M)$ with $|x| \leq 1$, we can consider the set $L_p = \{q \in \mathbb{Q}, q < x \in p\}$. It is a downward closed set, and it has a least upper bound, as it is bounded from above by 1.

This allows us to define the map :

$$
\begin{aligned}
f^* : & S_1(M) \to [0,1] \\
& p(x) \to \sup L_p \text{ if } |x| \leq 1 \in p(x) \\
& p(x) \to 0 \text{ else}
\end{aligned}
$$

Then this map clearly induces $st(x)$ by the above argument since the rationals are dense in the reals. Furthermore, it is continuous: without loss of generality assume $0 < a < b < 1 \in C$ (the other cases are similar.) Then

$$
(f^*)^{-1}((a,b)) = \bigcup_{q<r\in\mathbb{Q}\cap(a,b)} [q < x \leq r]
$$

Therefore the inverse image of basic open sets is open, so the function is continuous. Therefore $st(x)$ is definable as desired.

The key thing to note with this example is that the standard part map gives us a way to recover the usual topology on $\mathbb{R}$ from the Stone topology.

2.3 Imaginaries and T^{eq}

Recall that $\mathrm{acl}_M(A)$ (the algebraic closure of A in M) is the set of $\overline{b} \in M$ such that there exist an L_A-formula $\varphi(\overline{x})$ with $M \models \varphi(\overline{b})$ and $M \models \exists^{\leq k}\overline{x}\ \varphi(\overline{x})$. Similarly, the definable closure of A in M, denoted $\mathrm{dcl}_M(A)$, is the set of $\overline{b} \in M$ such that there exist an L_A-formula $\varphi(\overline{x})$ with $M \models \varphi(\overline{b})$ and $M \models \exists^{\leq 1}\overline{x}\ \varphi(\overline{x})$.

As an example, let the structure K be an algebraically closed field of characteristic 0 in the language of rings. Let $A \subseteq K$, and let k be the subfield of

K generated by A. Then $\bar{b} \in \mathrm{acl}_K(A)$ if and only if if it is an element of the algebraic closure of k in the algebraic sense. Similarly, $\bar{b} \in \mathrm{dcl}_K(A)$ if and only if $\bar{b} \in k$. This result is not totally trivial: for the full argument, see [Pil98].

Lemma 2.22. *Assume $M = \overline{M}$, $A \subseteq \overline{M}$ is small, and $\bar{b} \in M$.*

1. *$\bar{b} \in \mathrm{acl}_M(A)$ if and only if $\{f(\bar{b}) : f \in Aut(\overline{M}/A)\}$ is finite.*
2. *$\bar{b} \in \mathrm{dcl}_M(A)$ if and only id $f(\bar{b}) = \bar{b}$ for all $f \in Aut(\overline{M}/A)$.*

Proof.

1. Suppose $\bar{b} \in \mathrm{acl}_M(A)$ with witness $\exists^{\leq k}\varphi(\overline{x})$. Then the set defined by $\varphi(\overline{x})$ has at most k elements, e.g. is finite. Hence by Lemma 2.11, the set $\{\overline{m} : M \models \varphi(\overline{m})\}$ is $\mathrm{Aut}(\overline{M}/A)$-invariant, so $\{f(\bar{b}) : f \in \mathrm{Aut}(\overline{M}/A)\} \subseteq \{\overline{m} : M \models \varphi(\overline{m})\}$, therefore it is finite as desired.

   Suppose $\{f(\bar{b}) : f \in \mathrm{Aut}(\overline{M}/A)\}$ is finite. Since the composition of automorphisms is an automorphism, this set is $\mathrm{Aut}(\overline{M}/A)$-invariant. Therefore, by Lemma 1.11, it is definable by some formula $\varphi(\overline{x})$. Furthermore, $M \models \exists^{\leq k}\overline{x}\ \varphi(\overline{x})$ for some k since the set defined by $\varphi(\overline{x})$ is finite. Hence $\bar{b} \in \mathrm{acl}_M(A)$.

2. Suppose $\bar{b} \in \mathrm{dcl}_M(A)$ with witness $\exists^{\leq 1}\varphi(\overline{x})$. Then $\{\bar{b}\} = \{\overline{m} \in \overline{M} : \varphi(\overline{m})\}$ is a definable set, so by Lemma 2.11 it is $\mathrm{Aut}(\overline{M}/A)$-invariant. But any singleton which is invariant must be fixed pointwise, so $f(\bar{b}) = \bar{b}$ for all $f \in \mathrm{Aut}(\overline{M}/A)$, as desired.

   Suppose $f(\bar{b}) = \bar{b}$ for all $f \in \mathrm{Aut}(\overline{M}/A)$. Then $\{\bar{b}\}$ is $\mathrm{Aut}(\overline{M}/A)$-invariant, so by Lemma 1.11 it is definable over A via some formula $\varphi(\overline{x})$. Moreover, $M \models \exists^{\leq 1}\overline{x}\ \varphi(\overline{x})$ since $\varphi(\overline{x})$ defines a singleton, so $\bar{b} \in \mathrm{dcl}_{\overline{M}}(A)$.

□

Note the similarity of this lemma to lemmas 2.11 and 2.13. In fact, the the machinery of T^{eq} could be used to obtain these lemmas as a direct consequence of Lemma 2.11.

The first motivation to develop T^{eq} is dealing with quotient objects, without leaving the context of first order logic. That is, if E is some definable equivalence relation on some definable set X, we want to view X/E as a definable set.

We work in the setting of multi-sorted languages. Let L be a 1-sorted language and let T be a (complete) L-theory. We shall build a many-sorted language L^{eq} and an L^{eq}-theory T^{eq}. We will ensure that in some natural sense, L^{eq} contains L and T^{eq} contains T.

First we define L^{eq}. Consider the set L-formula $\varphi(\bar{x}, \bar{y})$, up to equivalence, such that T models that φ is an equivalence relation (e.g. reflexivity, transitivity, and symmetry). For each such φ, define s_φ to be a new sort in L^{eq}. Of particular importance is $s_=$, the sort given by the formula "$x = y$." This sort $s_=$ will yield, in each model of T^{eq}, a model of T.

Also define define f_φ to be a function symbol with domain sort $s_=^n$ (where φ has n free variables) and domain sort s_φ.

For each m-place relation symbol $R \in L$, make R^{eq} an m-place relation symbol in L^{eq} on $s_=^m$. Likewise for all constant and function symbols in L. Finally, for the sake of formality, we put a unique equality symbol $=_\varphi$ on each sort.

Remark 2.23. Let N be an L^{eq} structure. Then N has interpretations $s_\varphi(N)$ of each sort s_φ and $f_\varphi(N) : s_=(N)^{n_{f_\varphi}} \to s_\varphi(N)$ of each function symbol f_φ. Additionally, N will contain an L-structure consisting of $s_=$ and interpretations of the symbols of L inside of $s_=$.

Definition 2.24. T^{eq} is the L^{eq} theory which is axiomatized by the following:

1. T, where the quantifiers in the formulas of T now range over the sort $s_=$
2. For each suitable L-formula $\varphi(\bar{x}, \bar{y})$, the axiom $(\forall_{s_=} \bar{x})\ (\forall_{s_=} \bar{y})\ (\varphi(\bar{x}, \bar{y}) \leftrightarrow f_\varphi(\bar{x}) = f_\varphi(\bar{y}))$.
3. For each L-formula φ, the axiom $(\forall_{s_\varphi} y)\ (\exists_{s_=} \bar{x})\ f_\varphi(\bar{x}) = y$.

Note that axioms 2 and 3 simply state that f_φ is the quotient function for the equivalence relation given by φ.

Definition 2.25. Let $M \models T$. Then M^{eq} is the L^{eq} structure such that $s_=(M^{eq}) = M$, and for each suitable L-formula $\varphi(\bar{x}, \bar{y})$ of n variables, the sort $s_\varphi(M^{eq})$ is equal to $M^{n_{f_\varphi}}/E$, where E is the equivalence relation defined by $\varphi(\bar{x}, \bar{y})$, and $f_\varphi(M^{eq})(\bar{b}) = \bar{b}/E$.

One can now easily verify that $M^{eq} \models T^{eq}$, as expected. Moreover, passing from T to T^{eq} is a canonical operation, in the following sense :

Lemma 2.26. *1. For any $N \models T^{eq}$, there is an $M \models T$ such that $N \cong M^{eq}$.*

2. Suppose $M, N \models T$ are isomorphic, and let $h : M \xrightarrow{\sim} N$. Then h extends uniquely to $h^{eq} : M^{eq} \xrightarrow{\sim} N^{eq}$.

3. T^{eq} is a complete L^{eq}-theory.

4. Suppose $M, N \models T$ and let $\bar{a} \in M, \bar{b} \in N$ with $\mathrm{tp}_M(\bar{a}) = \mathrm{tp}_N(\bar{b})$. Then $\mathrm{tp}_{M^{eq}}(\bar{a}) = \mathrm{tp}_{N^{eq}}(\bar{b})$.

Proof. 1. Let $N \models T^{eq}$, we can take $M = s_=(N)$.

2. Suppose $M, N \models T$ are isomorphic, and let $h : M \xrightarrow{\sim} N$. Let $h^{eq} : M^{eq} \to N^{eq}$ be defined as $h^{eq}(f_\varphi(M^{eq})(\bar{b})) = f_\varphi(N^{eq})(h(\bar{b}))$ for each L-formula φ. This defines a function on M^{eq}, because $f_\varphi(M^{eq})$ is surjective by the T^{eq} axioms. Moreover h^{eq} is well-defined, bijective by the construction of M^{eq} and N^{eq}, and an isomorphism by definition.

3. Let $M, N \models T^{eq}$, we want to show that they are elementary equivalent. Assume the generalized continuum hypothesis (GCH). By GCH, there are $M', N' \models T^{eq}$ which are λ saturated of size λ, fro some large λ, with $M \preceq M'$ and $N \preceq N'$. Since we want to show elementary equivalence, we can replace M, N with M' and N'. By 1, we have $M = M_0^{eq}, N = N_0^{eq}$ for some $M_0, N_0 \models T$. Furthermore, M_0, N_0 are λ-saturated of size λ. By assumption, T is complete, so $M_0 \equiv N_0$, and therefore $M_0 \cong N_0$. By 2, $M \cong N$, and therefore $M \equiv N$.

   Remark that although convenient, the generalized continuum hypothesis is not necessary to prove this result. One could simply prove that there is a back and forth system between M and N, using such a system between $M \supset M_0 \models T$ and $N \supset N_0 \models T$

4. Let $M, N \models T^{eq}$, they are elementary small submodels of $\bar{M}$, with $\bar{M}$. Since $\mathrm{tp}_M(\bar{a}) = \mathrm{tp}_N(\bar{b})$, there exists an $h : M \xrightarrow{\sim} N$ automorphism of $\bar{M}$ such that $h(\bar{a}) = \bar{b}$. By 2, this automorphism h extends to $h^{eq} : M^{eq} \to N^{eq}$, and then $\mathrm{tp}_{M^{eq}}(\bar{a}) = \mathrm{tp}_{M^{eq}}(\bar{b})$.

□

Corollary 2.27. *Consider the Stone space $S_{(s_=)^n}(T^{eq})$. The forgetful map $\pi : S_{(s_=)^n}(T^{eq}) \to S_n(T)$ is a homeomorphism.*

Proof. Observe that it is continuous and surjective. By part 4 of the previous lemma, it is injective. Any continuous bijection from a compact space to a Hausdorff space is a homeomorphism. □

Proposition 2.28. *Let $\varphi(x_1, \dots, x_k)$ be an L^{eq} formula, where x_i is of sort S_{E_i}. There is an L-formula $\psi(\bar{y_1}, \dots, \bar{y_k})$ such that :*

$$T^{eq} \models (\forall \bar{y_1}) \ \dots \ (\forall \bar{y_k}) \ (\psi(\bar{y_1}, \dots, \bar{y_k}) \leftrightarrow \varphi(f_{E_1}(\bar{y_1}), \dots, f_{E_k}(\bar{y_k}))).$$

Proof. Let n be the length of $\bar{y_1} \dots \bar{y_k}$. Consider the set $\pi(\varphi(f_{E_1}(\bar{y_1}), \dots, f_{E_k}(\bar{y_k})))$, it is a clopen subset of $S_n(T)$ by the previous corollary, hence equal to $\psi(\bar{y_1}, \dots, \bar{y_k})$ for some formula ψ. One easily checks this formula is the one we're looking for.

□

Corollary 2.29. *1. Let $M, N \models T$, and let $h : M \to N$ be an elementary embedding. Then $h^{eq} : M^{eq} \to N^{eq}$ (defined as was done earlier) is also an elementary embedding.*

2. $\bar{M}^{eq}$ is also κ-saturated.

Remark 2.30. For $M \models T$, a definable set $X \subseteq M^n$ can be viewed as an element of M^{eq}. Suppose X is defined in M by $\varphi(\bar{x}, \bar{a})$ where $\bar{a} \in M$. Consider the equivalence relation E_ψ defined by $\psi(\bar{y_1}, \bar{y_2}) = (\forall \bar{x}) \ (\varphi(\bar{x}, \bar{y_1}) \leftrightarrow \varphi(\bar{x}, \bar{y_2}))$, and consider $c = \bar{a}/E_\psi = f_\psi(\bar{a}) \in M^{eq}$. Then X is defined in M^{eq} by $\chi(\bar{x}, c) = (\exists \bar{y}) \ (\varphi(\bar{x}, \bar{y}) \wedge f_\psi(\bar{y}) = c)$. Moreover, if $c' \in S_\psi(M^{eq})$, and $(\forall \bar{x}) \ (\chi(\bar{x}, c) \leftrightarrow \chi(\bar{x}, c')$, then $c = c'$. To see this, let $c' = f_\psi(\bar{a}')$, and let X' be defined in M by $\varphi(\bar{x}, \bar{a}')$. Then X' is defined in M^{eq} by $\chi(\bar{x}, c')$, so we have that $X = X'$ (in M^{eq}). And then $X = X'$ (in M), so $c = f_\psi(\bar{a}) = f_\psi(\bar{a}') = c'$.

Definition 2.31. With the above considerations in mind, given $M \models T$ and a definable set $X \subseteq M^n$, we call such a $c \in M^{eq}$ a code for X.

Remark 2.32. Any automorphism of $\bar{M}^{eq}$ fixes a definable set X set-wise if and only if it fixes a code for X. However, the choice of a code for X will depend on the formula φ used to define it, and so codes are not necessarily unique.

Definition 2.33. Let $A \subseteq M \models T$. Then $\operatorname{acl}^{eq}(A) = \{c \in M^{eq} : c \in \operatorname{acl}_{M^{eq}}(A)\}$, and $\operatorname{dcl}^{eq}(A)$ is defined similarly.

Remark 2.34. Suppose $A \subseteq M \prec N$. Then $\operatorname{acl}_{N^{eq}}(A), \operatorname{dcl}_{N^{eq}}(A) \subseteq M^{eq}$, so this notation is unambiguous.

Lemma 2.35. *Let $M \models T$, a definable subset X of M^n, and $A \subseteq M$. Then X is almost over A if and only if X is definable in M^{eq} by a formula with parameters in* $\operatorname{acl}^{eq}(A)$.

Proof. We can work in $\bar{M}$, since $M \prec \bar{M}$. Let c be a code for X. From Lemma 2.13, X is almost over A if and only if $|\{\sigma(X) : \sigma \in \operatorname{Aut}(\bar{M}/A)\}| < \omega$, if and only if $|\{\sigma(c) : \sigma \in \operatorname{Aut}(\bar{M}/A)\}| < \omega$, that is $c \in \operatorname{acl}(A)$. $\square$

Definition 2.36. Let $\bar{a}, \bar{b} \in \bar{M}$ have length n. Let $A \subseteq \bar{M}$. Then $\bar{a}, \bar{b}$ have the same strong type over A (written as $\operatorname{stp}_{\bar{M}}(\bar{a}/A) = \operatorname{stp}_{\bar{M}}(\bar{b}/A)$) if $E(\bar{a}, \bar{b})$ for any finite equivalence relation defined over A.

Remark 2.37. If $\varphi(\bar{x})$ is a formula over A, then it defines an equivalence relation with two classes by $E(\bar{x}_1, \bar{x}_2)$ if and only if $(\varphi(\bar{x}_1) \wedge \varphi(\bar{x}_2)) \vee (\neg\varphi(\bar{x}_1) \wedge \neg\varphi(\bar{x}_2))$. Hence, strong types are a refinement of types.

Lemma 2.38. *If $A = M \prec \bar{M}$, then $\operatorname{tp}_{\bar{M}}(a/M) \vdash \operatorname{stp}_{\bar{M}}(a/M)$.*

Proof. Let E be an equivalence relation with finitely many classes, defined over M, and $\bar{b}$ another realization of $\operatorname{tp}_{\bar{M}}(\bar{a}/M)$, we want to show $E(a, b)$. Since E has only finitely many classes, and M is a model, there are representants $e_1, \cdots, e_n$ of each E-class in M. Hence we must have $E(a, e_i)$ for some i, and therefore also $E(b, e_i)$, which yields $E(a, b)$, what we wanted.

$\square$

Lemma 2.39. *Let $A \subseteq M \models T$, and let $\bar{a}, \bar{b} \in M$. Then the following are equivalent:*

1. $\operatorname{stp}(\bar{a}/A) = \operatorname{stp}(\bar{b}/A)$
2. *$\bar{a}, \bar{b}$ satisfy the same formulas almost over A*
3. $\operatorname{tp}_{\bar{M}}(\bar{a}/\operatorname{acl}^{eq}(A)) = \operatorname{tp}_{\bar{M}}(\bar{b}/\operatorname{acl}^{eq}(A))$.

Proof. $2 \Leftrightarrow 3$ is a direct consequence of lemma 2.35, so we just need to prove $1 \Leftrightarrow 2$.

First, let's do the left to right implication. Assume $\bar{a}, \bar{b} \in M$ and $\operatorname{stp}(\bar{a}/A) = \operatorname{stp}(\bar{b}/A)$. Let X be definable almost over A, we want to show that $\bar{a} \in X$ if and only if $\bar{b} \in X$. By symmetry, it is enough to do one direction. So assume $\bar{a} \in X$.

Since X is almost over A, there is an A-definable equivalence relation E with finitely many classes, and $\bar{c}_1, \cdots, \bar{c}_n$ such that for all $\bar{x} \in M$, we have $\bar{x} \in X$ if and only if $M \models E(\bar{x}, \bar{c}_1) \vee \cdots \vee E(\bar{x}, \bar{c}_n)$. Hence $E(\bar{a}, \bar{c}_i)$ for some i, so by assumption $E(\bar{b}, \bar{c}_i)$. Therefore $b \in X$.

Now for the right to left direction, suppose $\bar{a}, \bar{b}$ satisfy the same almost over A formulas. Let E be an A-definable equivalence relation with finitely many classes, we want to show $E(\bar{a}, \bar{b})$. The set $X = \{\bar{x} \in M, E(\bar{x}, \bar{a})\}$ is definable almost over A. But $\bar{a} \in X$, so $\bar{b} \in X$, hence $E(\bar{a}, \bar{b})$.

□

Definition 2.40.
1. T has elimination of imaginaries (EI) if, for any model $M \models T$ and $e \in M^{eq}$, there is a $\bar{c} \in M$ such that $e \in \mathrm{dcl}_{M^{eq}}(\bar{c})$ and $\bar{c} \in \mathrm{dcl}_{M^{eq}}(e)$.
2. T has weak elimination of imaginaries if, as above, except $\bar{c} \in \mathrm{acl}_{M^{eq}}(e)$.
3. T has geometric elimination of imaginaries if, as above, except $e \in \mathrm{acl}_{M^{eq}}(\bar{c})$ and $\bar{c} \in \mathrm{acl}_{M^{eq}}(e)$.

Note that in particular, elimination of imaginaries imply the existence of codes for definable sets, which is a very useful property. Moreover, we have the following characterization :

Proposition 2.41. *The following are equivalent:*

1. *T has EI.*
2. *For some model $M \models T$, we have that for any $\varnothing$-definable equivalence relation E, there is a partition of M^n into $\varnothing$-definable sets $Y_1, \ldots Y_r$, and for each $i = 1, \ldots, r$ a $\varnothing$-definable $f_i : Y_i \to M^{k_i}$ where $k_i \geq 1$ such that for each $i = 1, \ldots, r$, for all $\bar{b_1}, \bar{b_2} \in Y_i$, we have $E(\bar{b_1}, \bar{b_2})$ iff $f_i(\bar{b_1}) = f_i(\bar{b_2})$.*
3. *For any model $M \models T$, we have that for any $\varnothing$-definable equivalence relation E, there is a partition of M^n into $\varnothing$-definable sets $Y_1, \ldots Y_r$, and for each $i = 1, \ldots, r$ a $\varnothing$-definable $f_i : Y_i \to M^{k_i}$ where $k_i \geq 1$ such that for each $i = 1, \ldots, r$, for all $\bar{b_1}, \bar{b_2} \in Y_i$, we have $E(\bar{b_1}, \bar{b_2})$ iff $f_i(\bar{b_1}) = f_i(\bar{b_2})$.*
4. *For any model $M \models T$, and any definable $X \subseteq M^n$ there is an L-formula $\varphi(\bar{x}, \bar{y})$ and $\bar{b} \in M$ such that X is defined by $\varphi(x, \bar{b})$ and for all $\bar{b}' \in M$ if X is defined by $\varphi(\bar{x}, \bar{b}')$ then $\bar{b} = \bar{b}'$. We call such a $\bar{b}$ a code for X.*

Proof. Note that properties in 2 and 3 concern only $\emptyset$-definable relations and functions. Hence, if it is true in some model, it is true in any model, so 2 and 3 are equivalent. As a consequence, we can and will work in $\overline{M}$ for the rest of the proof.

We start with $1 \Rightarrow 2$, and therefore consider some $\emptyset$-definable equivalence relation E. Let $\pi_E : S^n_= \to S_E$ the canonical definable quotient map. Let

$e \in S_E$. By assumption, there is $k \in \mathbf{N}$ and $\bar{c} \in \overline{M}^k$ such that e and $\bar{c}$ are interdefinable. In other words, there is a formula $\varphi_e(x, \bar{y})$ over $\emptyset$ such that $\varphi_e(e, \bar{c})$. Moreover, $\bar{c}$ is the unique tuple such that $\models \varphi_e(e, \bar{y})$, and e is the unique element such that $\models \varphi_e(x, \bar{c})$.

Let $X_e = \{x \in \overline{M}, \models (\exists! \bar{y} \varphi_e(\pi_e(x), \bar{y})) \wedge ((\forall z (E(x,z)) \leftrightarrow (\forall y (\varphi_e(\pi_E(x), \bar{y})) \leftrightarrow (\varphi_e(\pi_E(z), \bar{y})))))$. This means that φ_e defines a function on X_e, and that this function separates E-classes.

Then $\pi^{-1}(\{a\}) \subset X_e$. Indeed, let $\bar{a} \in \pi^{-1}(\{a\})$, then $\bar{c}$ is the unique realization of $\varphi_e(\pi_E(\bar{a}), \bar{y})$, hence the first half of the conjunction is true.

We now prove the equivalence of the right half of the conjunction. Suppose first $E(\bar{a}, \bar{b})$, then $\pi_E(\bar{a}) = e = \pi_E(\bar{b})$, and it is again the unique realization of $\varphi_e(\pi_E(\bar{a}), \bar{y})$, hence we get the left to right implication.

Conversely, suppose that $\models \forall \bar{y}(\varphi_e(\pi_E(\bar{a}), \bar{y}) \leftrightarrow \varphi_e(\pi_E(\bar{b}), \bar{y}))$, for some $\bar{b}$. By assumption, we have $\varphi_e(\pi_E(\bar{a}), \bar{c})$, hence $\varphi_e(\pi_e(\bar{b}), \bar{c})$. But by definition of φ_e, this implies that $e = \pi_E(\bar{b})$, and by definition of π_E, this yields $E(\bar{a}, \bar{b})$.

Since each X_e contains $\pi^{-1}(\{a\})$, we get $\overline{M}^n = \bigcup_{e \in \pi_E(\overline{M}^n)} X_e$, and by compactness, there are $e_1, \cdots, \cdots e_l$ such that $\overline{M}^n = \bigcup_{i=1}^{l} X_{e_i}$. We almost have what we wanted, by considering $f_i = \varphi_{e_i} \circ \pi_E$ on each X_{e_i}. However, the X_{e_i} are not disjoint.

But we can consider $Y_1, \cdots, Y_r$ to be the atoms of the boolean algebra generated by the X_i. These are disjoint, and we can pick, for each Y_j, an appropriate f_i, to get the result.

We now prove $3 \Rightarrow 4$. Let X be defined by $\psi(x, \bar{a})$. Consider the $\emptyset$-definable equivalence relation $E(\bar{y}, \bar{z}) \Leftrightarrow \forall x(\varphi(x, \bar{y}) \leftrightarrow \varphi(x, \bar{z}))$. Let Y_i and f_i be as in 2, and say $\bar{a} \in Y_1$, and let $\bar{b} = f_1(\bar{a})$. Then $\exists \bar{y} f_1(\bar{y}) = \bar{b} \wedge \varphi(x, \bar{y})$ defines X, call this formula $\psi(x, \bar{b})$.

We have to show that $\bar{b}$ is unique such. Let $\bar{b}'$ be such that $\exists \bar{y} f_1(\bar{y}) = \bar{b}' \wedge \varphi(x, \bar{y})$ also defines X, and let $\bar{a}_0$ be as the $\bar{y}$ in the formula. Then $\varphi(x, \bar{a}_0)$ defines X, hence $\bar{a}_0 E \bar{a}$, which implies $\bar{b}' = f_1(\bar{a}_0) = f_1(\bar{a}) = \bar{b}$.

Finally, we need to prove $4 \Rightarrow 1$. Let $e \in \mathcal{M}^{eq}$, then $e = \pi_E(\bar{a})$, for some $\bar{a} \in \overline{M}^n$ and some $\emptyset$-definable equivalence relation E.

The set $X = \{\bar{x} \in \overline{M}^n, \models E(\bar{x}, \bar{a})$ has a code $\bar{b} \in \overline{M}^k$, so that $X = \psi(\overline{M}^n, \bar{b})$. We are going to prove interdefinability of e and $\bar{b}$ using automorphisms of $\overline{M}$.

First suppose that $\sigma \in \mathrm{Aut}(\overline{M})$, and fixes e. We have $\overline{M}^{eq} \models \forall \bar{x}(\pi_E(\bar{x}) = \pi_E(\bar{a})) \leftrightarrow \psi(x, \bar{b})$. Applying σ, we get $\overline{M}^{eq} \models \forall \bar{x}(\pi_E(\bar{x}) = \pi_E(\bar{a})) \leftrightarrow \psi(x, \sigma(\bar{b}))$, and therefore $\overline{M}^{eq} \models \forall \bar{x}(\psi(x, \bar{b}) \leftrightarrow \psi(x, \sigma(\bar{b})))$. But $\bar{b}$ is a code for X, hence $\bar{b} = \sigma(\bar{b})$. This implies $\bar{b} \in \mathrm{dcl}(e)$.

Now suppose $\sigma \in \mathrm{Aut}(\overline{M})$, and fixes $\bar{b}$. Again $\overline{M}^{eq} \models \forall \bar{x}(\pi_E(\bar{x}) = \pi_E(\bar{a})) \leftrightarrow \psi(x, \bar{b})$, so applying σ we get $\overline{M}^{eq} \models \forall \bar{x}(\pi_E(\bar{x}) = \pi_E(\sigma(\bar{a}))) \leftrightarrow \psi(x, \bar{b})$. But $\varphi(\bar{a}, \bar{b})$, so $e = \pi_E(\bar{a}) = \pi_E(\sigma(\bar{a})) = \sigma(e)$. Hence $e \in \mathrm{dcl}(b)$.

$\square$

Note that condition 2 is somewhat unsatisfying, as we would like to have a quotient function for E, that is, have $r = 1$. Here is a condition for this to be true :

Proposition 2.42. *Suppose T eliminates imaginaries. We get $r = 1$ in condition 2 if and only if* $\operatorname{dcl}(\emptyset)$ *has at least two elements.*

Proof. First, suppose that $r = 1$. Consider the equivalence relation on $\overline{M}^2$ given by $E((x,y),(x',y'))$ if and only if $x = y \wedge x' = y'$ or $x \neq y \wedge x' \neq y'$. In other words, the E classes are the diagonal and its complement. Then $\pi_E(\overline{M}^2)$ has two elements, and they belong to $\operatorname{dcl}^{eq}(\emptyset)$. But because T eliminates imaginaries, this implies that there is also two elements in $\operatorname{dcl}(\emptyset)$.

Second, suppose that $\operatorname{dcl}(\emptyset)$ contains two constants a and b. Let Y_i, f_i be as in condition 2. Using a and b, we can find some number k and functions $g_i : \overline{M}^{k_i} \to \overline{M}^k$ such that the $g_i(\overline{M}^{k_i})$ are pairwise disjoint. We can check that the $\emptyset$-definable function $f : \overline{M}^n \to \overline{M}^k$ sending $y \in Y_i$ to $g_i(f_i(y))$ has all the required properties. □

Remark 2.43. Elimination of imaginaries also makes sense for many sorted theories. Barring this in mind, we will now give a lemma.

Lemma 2.44 (Assume T 1-sorted)**.** *T^{eq} has elimination of imaginaries.*

Proof. We prove a strong version of (2) in proposition 2.41. Let E' be a $\emptyset$-definable equivalence relation on a sort s_E in some model M^{eq} of T^{eq}. By proposition 2.28, there is an L-formula $\psi(\overline{y}_1, \overline{y}_2)$ ($\overline{y}_i$ the appropriate length) such that for all $\overline{a}_1, \overline{a}_2 \in M$, $M \models \psi(\overline{a}_1, \overline{a}_2)$ if and only if $M^{eq} \models E'(f_E(\overline{a}_1), f_E(\overline{a}_2))$. So $\psi(\overline{y}_1, \overline{y}_2)$ is an L-formula defining an equivalence relation on M^k for the suitable length k. Consider the map h, taking $e \in S_E$ to $f_\psi(\overline{a})$[1] for any $\overline{a} \in M^k$ such that $f_E(\overline{a}) = e$. Suppose $f_E(\overline{a}) = e = f_E(\overline{a'})$, we easily see that $f_\psi(\overline{a}) = f_\psi(\overline{a'})$, hence the map h is well defined, and satisfied (2) of 2.41. □

In applied model theory, a substantial amount of work has been done on proving some relative elimination of imaginaries. That is, given a theory T, is there some natural imaginary sorts one can add to T so that the resulting theory has elimination of imaginaries ? The interested reader can look at what has been achieved in the case of algebraically closed valued fields.

2.4 Examples and counterexamples

Even in the simplest case, there is some non-trivial result :

Example 2.45. The theory of an infinite set has weak elimination of imaginaries but not full elimination of imaginaries.

[1] $f_\psi(\overline{a}) \in S_\psi$.

Proof. First, we show that T has weak elimination of imaginaries. Let M be an infinite set and let $e \in M^{eq}$ be an imaginary element. Suppose that $e = \ulcorner X \urcorner$ for some definable set X. Let $A \subset M$ be a finite set over which X is definable. Consider the set

$$\hat{A} := \bigcap_{\substack{\sigma \in \mathrm{Aut}(M) \\ \sigma(X)=X}} \sigma(A).$$

Since A is finite, there are $\sigma_1, \ldots, \sigma_n$ such that $\hat{A} = \cap_i \sigma_i(A)$. Observe that X is definable over $\hat{A}$; since automorphisms of M are just permutations, for any finite sets B_1 and B_2, $\mathrm{Aut}(M/B_1 \cap B_2)$ is generated by $\mathrm{Aut}(M/B_1)$ and $\mathrm{Aut}(M/B_2)$. Therefore if X is definable over B_1 and B_2, then X is definable over $B_1 \cap B_2$. Now, it suffices to show that $\hat{A} \subseteq \mathrm{acl}^{eq}(e)$. Let $a \in \hat{A}$. For any $\sigma \in \mathrm{Aut}(M^{eq})$ fixing e, we have that $\sigma(X) = X$ and so $\sigma(\hat{A}) = \hat{A}$, i.e. $\sigma(a) \in \hat{A}$. Thus we have the orbit $\mathrm{Aut}(M^{eq}/e) \cdot a$ is contained in $\hat{A}$ (which is finite) and so, by definition, $a \in \mathrm{acl}^{eq}(e)$.

To see that T does not have full elimination of imaginaries, observe that there is never a code for any finite set. Indeed, if M is an infinite set, $X \subset M$ a finite subset, and $\bar{a}$ an arbitrary tuple from M, we can find a permutation of M which fixes X as a set but does not fix $\bar{a}$, meaning $\bar{a}$ could not be a code for X. $\square$

The previous example shows that non-complex theories do not necessarily have elimination of imaginaries. And the other way around, some very complex theories may have elimination of imaginaries :

Example 2.46. Let $T = Th(M, <, ...)$ where $<$ is a total well-ordering (e.g. True arithmetic). Then, T has elimination of imaginaries.

Proof. Every definable set has a first element since $<$ is a total well-ordering. We verify (2) in 2.41. Let E be an $\emptyset$-definable equivalence relation on M^n. Let $f : M^n \to M^n$ such that for any $\overline{a}$, $f(\overline{a})$ is the least element of the E-class of $\overline{a}$, with respect to the lexicographic ordering. Notice that f is $\emptyset$-definable, and for all $\bar{a}, \bar{b}$, we have $f(\bar{a}) = f(\bar{b})$ if and only if $E(\bar{a}, \bar{b})$. $\square$

Therefore, one should not think of elimination of imaginaries as measuring the complexity of a theory.

Definition 2.47. A theory T is strongly minimal if for any model M of T and any definable set $X \subset M$, either X is finite or $M \setminus X$ is finite.

Example 2.48. The following theories are strongly minimal :

- the theory of equality
- the theory of an algebraically closed fields of any characteristic
- the theory $\mathrm{Th}(\mathbb{Q}, +)$

From the point of view of model theory, strongly minimal theories are very well understood, and in particular, we have :

Lemma 2.49. *Let T be strongly minimal and* $\mathrm{acl}(\emptyset)$ *be infinite (in some, any model). Then T has weak elimination of imaginaries.*

Proof. Fix a model M. Let $e \in M^{eq}$, so $e = \overline{a}/E$ for some $\overline{a} = (a_1, ..., a_n)$ and E some some $\emptyset$-definable equivalence relation. Let $A = \mathrm{acl}_{M^{eq}}(e) \cap M$. A is infinite as it contains $\mathrm{acl}(\emptyset)$.

We first prove that there exists some $\overline{b} \subset A$ such that $E(\overline{a}, \overline{b})$. Let $X_1 = \{y_1 \in M : M \models \exists y_2, ..., y_n(\overline{y} E \overline{a})\}$, it is definable over e. If X_1 is finite, any b_1 in X_1 then belongs to A. Otherwise, X_1 is cofinite, hence meets the infinite set A. Either way, $X_1 \cap A \neq \emptyset$ and we have $b_1 \in X_1 \cap A$.

Now let $X_2 = \{y_2 \in M : M \models \exists y_3, ..., y_n(b_1 \overline{y} E \overline{a})\}$. We remark that $X_2 \neq \emptyset$ since $b_1 \in X_1$. Now, X_2 is either finite or cofinite since T is strongly minimal. By the same argument above, we may find $b_2 \in X_2 \cap A$. Then, repeating this process, we may find $\overline{b} \subset A$. Therefore, $\overline{b} \in \mathrm{acl}_{M^{eq}}(e)$.

Finally, notice that that $e \in \mathrm{dcl}_{M^{eq}}(\overline{b})$ since $\overline{b}/E = \overline{a}/E = e$. ☐

Example 2.50. The theory ACF_p has elimination of imaginaries, for any p.

Proof. By Lemma 2.49, ACF_p has weak elimination of imaginaries. Therefore, it suffices to show that that every finite set can be coded. Let K be an algebraically closed field and let $X = \{c_1, \ldots, c_n\} \subseteq K$. Consider the polynomial

$$\begin{aligned} P(x) &= \prod_{i=1}^{n}(x - c_i) \\ &= x^n + e_{n-1}x^{n-1} + \ldots + e_1 x + e_0. \end{aligned}$$

Then we may take the tuple $\bar{e} = (e_{n+1}, \ldots, e_1, e_0)$ to be our code for X. Indeed, if an automorphism of K permutes the elements of X, then certainly the polynomial $P(x)$ is fixed and so $\bar{e}$ is fixed. On the other hand, if some automorphism fixes $\bar{e}$, then $P(x)$ is fixed and so at most the automorphism permutes the roots of $P(x)$.

☐

Example 2.51. The theory of real closed fields has elimination of imaginaries. To see this, we prove that $T = RCF$ has definable choice.

Definition 2.52. A theory T has definable choice if for any formula $\varphi(\bar{x}, \bar{y})$ and any model M of T, there is a definable function $f_\varphi : M^{\bar{\bar{y}}} \to M^{\bar{\bar{y}}}$ satisfying :

- $M \models \forall \bar{y} \varphi(f_\varphi(\bar{y}), \bar{y})$
- if $\varphi(\mathcal{R}, \bar{a}) = \varphi(\mathcal{R}, \bar{b})$, then $f_\varphi(\bar{a}) = f_\varphi(\bar{b})$

We can then use the general property :

Proposition 2.53. *If T has definable choice, then it eliminates imaginaries.*

Proof. Let E be an equivalence relation, definable over the empty set. Consider the definable function f_E. Then for any $\bar{a}, \bar{b}$ we check that $E(\bar{a}, \bar{b})$ if and only if $f_E(\bar{a}) = f_E(\bar{b})$, which yields elimination of imaginaries by Lemma 2.41.

First suppose $E(\bar{a}, \bar{b})$, then since $\bar{M} \models \forall \bar{y} E(f_E(\bar{y}), \bar{y})$, we have $E(f_E(\bar{a}), f_E(\bar{b}))$, hence $E(\bar{M}, f_E(\bar{a})) = E(\bar{M}, f_E(\bar{b}))$, so $f_E(\bar{a}) = f_E(\bar{b})$.

Now suppose $f_E(\bar{a}) = f_E(\bar{b})$. We have $E(f_E(\bar{a}), \bar{a})$ and $E(f_E(\bar{b}), \bar{b})$, but also $E(f_E(\bar{a}), f_E(\bar{b}))$. Hence $E(\bar{a}, \bar{b})$.

□

Proposition 2.54. *The theory of real closed fields has definable choice.*

Proof. We proceed by induction on $|\bar{x}|$.

$|\bar{x}| = 1$: By quantifier elimination, for any $\bar{b}$, the definable set $\varphi(\mathcal{R}, \bar{b})$ is a finite union of intervals (we will assume points are intervals as well). Let I be the left-most interval of $\varphi(\mathcal{R}, \bar{b})$. We define $f(\bar{b})$ in cases depending on the shape of I:

- if $I = \{a\}$ is a singleton, set $f(\bar{b}) = a$,
- if $I = \mathcal{R}$, set $f(\bar{b}) = 0$,
- if the interior of I is of the form $(c, +\infty)$, set $f(\bar{b}) = c + 1$,
- if the interior of I is of the form $(-\infty, c)$, set $f(\bar{b}) = c - 1$,
- if the interior of I is of the form (c, d), then set $f(\bar{b}) = \frac{c+d}{2}$.

Note that f is definable, since for any interval I, the supremum and infimum of I are definable (provided they exist).

$|\bar{x}| = n+1$: Write $\bar{x} = (x_0, \bar{x}_1)$, where $|\bar{x}_1| = n$. By the induction hypothesis, there is a definable function $f(\bar{y})$ such that

- $\mathcal{R} \models \forall \bar{y} \exists x_0 \varphi(x_0, f(\bar{y}), \bar{y})$, and
- $\exists x_0 \varphi(x_0, \mathcal{R}, \bar{a}) = \exists x_0 \varphi(x_0, \mathcal{R}, \bar{b})$, then $f(\bar{a}) = f(\bar{b})$.

Now, by the case of $|\bar{x}| = 1$, we have that there is a definable function $g(\bar{x}_1, \bar{y}) : \mathcal{R}^{|y|+1} \to \mathcal{R}$ such that the map $\bar{y} \mapsto (g(f(\bar{y}), \bar{y}), f(\bar{y}))$ satisfies the requirements of the proposition. □

There is more to the interaction between definable choice and elimination of imaginaries :

Definition 2.55. T has skolem functions if for each formula $\varphi(\bar{x}, \bar{y})$ there is some definable (over the empty set) function $f_\varphi(\bar{y})$ such that for any model M of T, we have $M \models \forall y((\exists \bar{x} \varphi(\bar{x}, \bar{y})) \to (\varphi(f_\varphi(\bar{y}), \bar{y})))$.

Note that definable choice implies definable skolem functions, but the converse is false. Indeed, the theory $\mathrm{Th}(\mathbb{Q}_p, +, \times)$ has skolem functions, but does not have elimination of imaginaries, and hence does not have definable choice.

However we have :

Fact 2.56. *A theory T has definable choice if and only if T^{eq} has definable skolem functions.*

We conclude this section by mentioning the following questions, which, as far as we know, are completely unexplored :

Question 1. Consider the theories ZF and ZFC for sets. Which completion of these theories have skolem function, elimination of imaginaries, or definable choice ?

3 Stability

Throughout this chapter we will fix a complete theory T in some language L. Moreover, we will have no problem in working in T^{eq} (that is to say, to assume $T = T^{\text{eq}}$), at least for the general theory. In the context of specific examples, however, we will only deal with 1-sorted theories.

Before we begin, a little bit of history. The origin of Stability Theory can be traced back to Morley's work in the sixties regarding uncountable categoricity. He proved the following theorem:

Theorem. *Suppose T is a countable theory. Then T is κ-categorical for some $\kappa > \aleph_0$ if and only if T is κ-categorical for all $\kappa > \aleph_0$.*

A key step in the proof of this theorem is to show that, if T is κ-categorical for some $\kappa > \aleph_0$, then T is ω-stable (a property called "totally transcendental" by Morley). This property, ω-stability, is a strong form of stability, defined as follows: for all $n < \omega$ and all countable models $M \models T$, the cardinality of $S_n(M)$ is at most countable (which in turn implies $|S_n(M)| \leq |M|$ for *all* $M \models T$).

Assuming ω-stability, additional machinery was developed (namely, Morley rank) to, together with indiscernibles, deduce the theorem. Subsequently, it was seen that, in fact, a theory T is κ-categorical for some/any $\kappa > \aleph_0$ if and only if T is ω-stable and unidimensional (i.e. any two types are nonorthogonal), a characterisation of $\aleph_1$-categorical theories without any mention of an uncountable cardinal.

Later on, Shelah, who was working in Classification theory[2], considered the "test question" of what forms could the spectrum function take. For a given theory T, the spectrum function is given as

$$I(T, -) : \text{cardinals} \longrightarrow \text{cardinals}$$
$$I(T, \lambda) = \#\text{ of models of } T \text{ of cardinality } \lambda \text{ (up to isomorphism)}$$

In solving this question, he invented Stability theory. For instance, if T is unstable, then $I(T, \lambda) = 2^\lambda$ for all $\lambda > \aleph_0$, i.e. T has the greatest possible

[2]Shelah was attempting to classify theories, that is to say, to find meaningful division lines between them (for example, in former times, theories were divided between decidable and undecidable). In this endeavor, he was trying to formulate "test questions" that could help classify first order theories.

number of models in each uncountable cardinality. Hence, if we want to study theories through their spectrum, we can restrict our attention to stable theories.

Before we start with local stability, a quick additional historical remark. The notion of stability, as many others in mathematics, appeared more or less simultaneously in two different contexts. In this case, we have Grothendieck's thesis in which, in the context of Functional Analysis, the notion under consideration makes an appearance, as does a proof which we will give shortly (cf. Proposition 3.7).

3.1 Local Stability

We begin by recalling that, for an L-formula φ, to write φ as $\varphi(x_1, \dots, x_n)$ means that the free variables of φ are among $x_1, \dots, x_n$ (which we assume are distinct). Similarly, writing φ as $\varphi(\overline{x}, \overline{y})$ means $\overline{x}, \overline{y}$ form a partition of $x_1, \dots, x_n$. (Note that $\varphi(\overline{x}, \overline{y})$ thus given defines a bipartite graph.)

Definition 3.1.

(i) Let $M \models T$. We say $\varphi(\overline{x}, \overline{y})$ is *stable in M* if it is *not* the case that there are $\overline{a}_i, \overline{b}_i$ in M, for $i < \omega$, such that *either*, for all $i, j < \omega$, $M \models \varphi(\overline{a}_i, \overline{b}_j)$ if and only if $i \leq j$, *or*, for all $i, j < \omega$, $M \models \neg\varphi(\overline{a}_i, \overline{b}_j)$ if and only if $i \leq j$.

(ii) We say $\varphi(\overline{x}, \overline{y})$ is *stable* (for T) if it is stable in M for all $M \models T$.

(iii) Finally, we say T is *stable* if every L-formula $\varphi(\overline{x}, \overline{y})$ is stable (for T).

Remark 3.2. A formula $\varphi(\overline{x}, \overline{y})$ is stable for T if and only if it is *not* the case that there are $\overline{a}_i, \overline{b}_i$ in the monster model $\overline{M}$, $i < \omega$, such that $\overline{M} \models \varphi(\overline{a}_i, \overline{b}_i)$ if and only if $i \leq j$ for all $i, j < \omega$.

Before we go on, a remark on notation. For simplicity, from now onwards we will write tuples simply as x and a, instead of the more cumbersome $\overline{x}$ and $\overline{a}$. Thus, in general, when say we write $\varphi(x, y)$, we understand x and y not as single variables but as tuples (possibly of length 1) of variables.

We start by proving a few elementary facts about stable formulas :

Lemma 3.3.

(i) Suppose $\varphi(x, y)$ and $\psi(x, z)$ are stable for T. Then $\neg\varphi(x, y)$, $(\varphi \vee \psi)(x, yz)$ and $(\varphi \wedge \psi)(x, yz)$ are also stable.

(ii) Given $\varphi(x, y)$, let $\varphi^*(y, x)$ be $\varphi(x, y)$. Then, $\varphi(x, y)$ stable for T implies $\varphi^*(y, x)$ is stable for T as well.

(iii) The formula $\varphi(x, y)$ is stable for T if and only if there is $n < \omega$ such that $\varphi(x, y)$ is n-stable: it is not *the case that there are a_i, b_i (in $\overline{M}$ or in some/any $M \models T$), $i \leq n$, such that $\models \varphi(a_i, b_i)$ if and only if $i \leq j$ for all $i, j \leq n$.*

(iv) There are T, $M \models T$ and $\varphi(x, y)$ such that $\varphi(x, y)$ is stable in M but it is not stable for T.

In particular, as a consequence of the lemma, notice that (iii) implies that stability is expressed by a sentence of the theory.

Proof. (i) We will only prove it for the negation, the other cases being similar. Enrich the language L with tuples of constant symbols $\overline{c}_i, \overline{d}_i$, $i < \omega$, and consider the set $\Sigma = T \cup \{\neg\varphi(\overline{c}_i, \overline{d}_j) : i \leq j\} \cup \{\varphi(\overline{c}_i, \overline{d}_j) : i > j\}$ of formulas in the new language. For a finite subset $\Sigma' \subseteq \Sigma$, let n be the largest index occurring as a subscript of a new constant tuple in some sentence of Σ'. Then, interpreting $\overline{c}_i$ as $\overline{a}_{n+1-i}$ and $\overline{d}_i$ as $\overline{b}_{n-i}$ in $\overline{M}$ for all $i \leq n$, and the rest of new constants arbitrarily, yields a model of Σ'. In fact, this is true because $i \leq j$ implies $n+1-i > n-j$, and $i > j$ implies $n+1-i \leq n-j$. By compactness we can thus obtain a model of Σ, which will yield the desired tuples.

(ii) Suppose $\varphi^*(y, x)$ is not stable, so by (i) $\neg\varphi^*(y, x)$ is also unstable. Let a_i, b_i be witnesses in $\overline{M}$ of the latter. Then, $a'_i = b_i$ and $b'_i = a_{i+1}$, $i < \omega$, witness the instability of $\varphi(x, y)$, as $j + 1 > i$ holds if and only if $i \leq j$. It follows that $\varphi^*(y, x)$ must be stable.

(iii) Any witnesses of the failure of stability for $\varphi(x, y)$ yield witnesses of the failure of n-stability for every $n < \omega$. Thus, n-stable for some n implies stable.

On the other hand, if $\varphi(x, y)$ is n-unstable for all n, then by an straightforward compactness argument, we can prove that φ is unstable.

(iv) Consider the graph G, disjoint union of all finite graphs. Then the edge relation E is stable in G. Indeed, if it wasn't, we would in particular have a vertex x_0 and infinitely many vertices $\{y_i, i \in \mathbb{N}\}$ such that $E(x_0, y_i)$ for all i. But this would imply, if G_0 is the graph containing x_0, that $y_i \in G_0$ for all i, which is impossible since G_0 is finite.

But by (iii), the edge relation is not stable in $\mathrm{Th}(G)$. $\square$

Definition 3.4. Fix $\varphi(x, y)$ in L. By a *complete φ-type over M*, $M \models T$, we mean a maximal consistent set of instances of φ and $\neg\varphi$ over M, namely L_M-formulas of the form $\varphi(x, b)$, $\neg\varphi(x, b)$ for $b \in M$. We write $S_\varphi(M)$ for the set of such complete φ-types over M.

Remark 3.5.

(i) By a *φ-formula over M* we mean a Boolean combination of instances (over M) of φ and $\neg\varphi$. For example, $(\varphi(x, c) \wedge \varphi(x, b)) \vee \neg\varphi(x, d)$ is a φ-formula.

(ii) Any type $p(x) \in S_\varphi(M)$ decides any φ-formula $\psi(x)$ over M, that is to say $p(x) \models \psi(x)$ or $p(x) \models \neg\psi(x)$, so in fact $p(x)$ extends to a unique maximal consistent set of φ-formulas over M

(iii) By defining the basic open sets of $S_\varphi(M)$ to be $\{p(x) \in S_\varphi(M) : \psi(x) \in p\}$ for ψ a φ-formula, $S_\varphi(M)$ becomes a compact totally disconnected space, where in addition the clopen sets are precisely given by φ-formulas, i.e. they are the basic clopen sets.

(iv) Any $p(x) \in S_\varphi(M)$ extends to some $q(x) \in S_x(M)$ such that $p = q \restriction \varphi$, where $q \restriction \varphi$ is the set of φ-formulas in $q(x)$ (or instances of $\varphi, \neg\varphi$ in $q(x)$).

Definition 3.6.

(i) Let $p(x) \in S_x(M)$ be a complete type over M. We say that $p(x)$ is *definable* if, for each $\varphi(x, y)$ in L, there is an L_M-formula $\psi(y)$ such that for all $m \in M$, we have $M \models \psi(m)$ if and only if $\varphi(x, b) \in p$ (note that such $\psi(y)$ is unique up to equivalence). We say the type $p(x)$ is *definable over* $A \subseteq M$ if each such $\psi(y)$ is over A.[3]

(ii) Likewise, we speak of the φ-type $p(x) \in S_\varphi(M)$ being *definable* when $\{b \in M : \varphi(x, b) \in p(x)\}$ is defined by a formula $\psi(y)$ of L_M. (Note that in this case $\psi(y)$ determines $p(x)$.)

As we will see later, a theory T is stable if and only if all types over *all* models of T are definable. Note that there are unstable theories for which all the types over certain models are definable. For instance, in the case of dense linear orders, all types over $\mathbb{R}$ (the standard model) are definable.

Indeed, by quantifier elimination, any non-realized 1-type over any model of DLO corresponds to a cut in its order. But in the case of $\mathbb{R}$, the order is complete, so for any cut, there will in fact exist a real number r such that the cut is of the form $(\{l \in \mathbb{R}, l < r\}, \{d \in \mathbb{R}, d > r\})$. Using this real number r, one can easily show definability of 1-types over $\mathbb{R}$. A similar proof would work for higher arities.

Another example of such a model, in the the theory of p-adically closed fields, is $\mathbb{Q}_p$.

Now we come to one of the fundamental results of the subject, linking definability of types and stability. This is the most general version of this result, only assuming stability of a formula in a given model. Note that Grothendieck gave, in his thesis, a result which can be interpreted as a functional analysis version of this (and indeed, our proposition is a consequence of it).

Proposition 3.7. *Fix a model $M \models T$ and an L-formula $\varphi(x, y)$. Then the following are equivalent:*

(i) $\varphi(x, y)$ is stable in M.

(ii) Whenever $M^* \succ M$ is $|M|^+$-saturated and $\mathrm{tp}(a^*/M^*)$ is finitely satisfiable in M, then $\mathrm{tp}_\varphi(a^*/M^*)$ is definable over M and, moreover, it is defined by some φ^*-formula, i.e. a Boolean combination of $\varphi(a, y)$'s, $a \in M$.

Proof. We start by proving (i) $\Rightarrow$ (ii). Fix some $p^*(x) = \mathrm{tp}(a^*/M^*)$ finitely satisfiable in M. We want to prove $\mathrm{tp}_\varphi(a^*/M^*)$ is definable over M by a φ^*-formula. Note first that, as p^* is finitely satisfiable in M, whether or not some $\varphi(x, b)$, $b \in M^*$, is in in p^* depends only on $\mathrm{tp}(b/M)$; in fact, even only on $\mathrm{tp}_{\varphi^*}(b/M) = q(y) \in S_{\varphi^*}(M)$. To see this, suppose, by way of contradiction, that we had $b' \in M^*$ such that $\mathrm{tp}_{\varphi^*}(b'/M) = \mathrm{tp}_{\varphi^*}(b/M)$, but $\varphi(x, b) \in p^*$ while $\neg\varphi(x, b') \in p^*$. Then we would have $\models \varphi(a^*, b) \wedge \neg\varphi(a^*, b')$, so by finite satisfiability there would be $a \in M$ such that $\models \varphi(a, b) \wedge \neg\varphi(a, b')$. Yet this

[3]This definition also occurred in the work of Gaifman regarding models of arithmetic.

contradicts the fact that b and b' have the same φ^*-type over M. Hence, we may write $\varphi(x, q) \in p^*$ whenever $\varphi(x, b) \in p^*$ and $q = \mathrm{tp}_{\varphi^*}(b/M)$.

Claim. *In order for $p^* \restriction \varphi$ to be definable by a φ^*-formula over M, it suffices to prove (*) : for any type $q(y) \in S_{\varphi^*}(M)$, if $\varphi(x, q) \in p^*$, then there is some φ^*-formula $\psi(y) \in q(y)$ such that, for all $b \in M$ satisfying $\psi(y)$, we have $\varphi(x, b) \in p^*$. Similarly if $\neg\varphi(x, q) \in p^*$.*

Proof of claim. Suppose for all q we find $\psi_q(y)$ as above. By compactness, finitely many $\psi_q(y)$ cover $S_{\varphi^*}(M)$. Thus, $M \models (\forall y)(\psi_{q_1}(y) \vee \cdots \vee \psi_{q_n}(y))$ for some $q_1, \ldots, q_n \in S_{\varphi^*}(M)$, and we may assume the ψ_{q_i} are pairwise contradictory. Let $\psi(y)$ be the disjunction of those $\psi_{q_i}(y)$ for which $\varphi(x, q_i) \in p^*$. Then $\psi(y)$ is a φ^*-formula which defines $p^* \restriction \varphi$.

To prove this, we need to show that for any q, we have $\varphi(x, q) \in p^*$ if and only if $\psi(y) \in q$. The right to left direction is by construction of the ψ_{q_i}.

For the other direction, suppose $b^* \in M^*$, $q = \mathrm{tp}_{\varphi^*}(b^*/M)$ and $\varphi(x, b^*) \in p^*$. We want to show $\models \psi(b^*)$; by contradiction, assume instead we have $\models \neg\psi(b^*)$. Then $\psi_q(y) \wedge \neg\psi(y) \in q$ and we may take some $b \in M$ such that $\models \psi_q(b) \wedge \neg\psi(b)$. Now, by assumption, since $\models \psi_q(b)$, we must have $\varphi(x, b) \in p^*$. Also $\neg\varphi(x, b) \in p^*$, because $\models \neg\psi(b)$ implies $\models \psi_{q_i}(b)$ for some q_i such that $\neg\varphi(x, q_i) \in p^*$ by choice of ψ. Of course, this is impossible, so $\models \psi(b^*)$. On the other hand, if $\models \psi(b^*)$, then again the choice of ψ and the assumption about the ψ_q's tell us that $\varphi(x, b^*) \in p^*$. The claim is proved.

$\square$

Now we prove that (*) holds. Suppose it fails for some $q \in S_{\varphi^*}(M)$ and let $b^* \in M^*$ realize q. Without loss of generality, we may assume $\varphi(x, q) \in p^*$. Then, for every φ^*-formula $\psi(y) \in q$ there is $b \in M$ such that $\models \psi(b)$ and $\neg\varphi(x, b) \in p^*$. With this we will inductively construct $a_i, b_i \in M$, $i < \omega$, with the following properties:

1) $\models \varphi(a_i, b_j)$ if and only if $i \leq j$,

2) $\models \neg\varphi(a^*, b_i)$ for all i, and

3) $\models \varphi(a_i, b^*)$ for all i.

Suppose we have found a_i, b_i, $i < n$, with these properties. As $\mathrm{tp}(a^*/M^*)$ is finitely satisfiable in M, using 2) we can find $a_n \in M$ such that $\models \varphi(a_n, b^*)$ and $\models \neg\varphi(a_n, b_i)$ for all $i < n$. Now, note that we have $\models \bigwedge_{i \leq n} \varphi(a_i, b^*)$, so the φ^*-formula $\bigwedge_{i \leq n} \varphi(a_i, y)$ is in q. Thus, since we assumed (*) fails for q, there exists some $b_n \in M$ such that $\models \bigwedge_{i \leq n} \varphi(a_i, b_n)$, but $\neg\varphi(x, b_n) \in p^*$, i.e. $\models \neg\varphi(a^*, b_n)$. Hence, we get properties 1), 2) and 3) for $i < n + 1$ and the induction carries on. The resultant sequences witness the fact that φ is not stable in M, a contradiction. Thus, we have shown that (i) implies (ii).

We now will prove (ii) implies (i), so assume (ii), and by way of contradiction, assume that φ is not stable in M. Then there are sequence $a_i, b_j \in M$, with $i, j < \omega$, such that $M \models \varphi(a_i, b_j)$ if and only if $i \leq j$.

Consider the set of formulas :

$$\{x \neq a_i, i \in \mathbb{N}\} \cup \{\theta(x) \in L_{M^*}, \text{ for all } i, M \models \theta(a_i)\}$$

it is finitely satisfiable (just pick an appropriate a_i), hence satisfiable by compactness. Let a^* be a realization, and $p = \mathrm{tp}(a^*/M)$. By construction, this type p is finitely satisfiable in $\{a_i, i \in \mathbb{N}\}$, so in particular in M. By (ii), it ought to be definable over M, so there is an M-formula $\psi(y)$ such that for all $b \in M^*$, we have $\varphi(a^*, b)$ if and only if $M^* \models \varphi(b)$.

Suppose that $\psi(b_j)$ for some j, then if $j' > j$, we have to have $\psi(j')$ as well. Else, we would have $\varphi(a, b_{j'})$, and since p is finitely satisfiable in $\{a_i, i \in \mathbb{N}\}$, some $i \in \mathbb{N}$ such that $\neg\varphi(a_i, b_j) \wedge \varphi(a_i, b_{j'})$, which is impossible because $j < j'$.

So there are two possibilities : either $\neg\psi(b_j)$ for all j, or there is j such that $\varphi(b_{j'})$ for all $j' > j$. In the second case, by cutting the sequence at j, we can assume $\psi(b_j)$ for all j.

We will only treat the possibility $\psi(b_j)$ for all j, the other being similar. Consider the set of formula :

$$\psi(x) \cup \{\neg\varphi(a_i, x), \text{ for all } i \in \mathbb{N}\}$$

it is finitely satisfiable by what we just proved, hence satisfiable, let $b \in M^*$ be a realization (which exists by $|M|^+$ saturation of M^*). Then $M^* \models \neg\varphi(a, b)$. Indeed, suppose on the contrary that $M^* \models \varphi(a, b)$. Again, since p is finitely satisfiable in $\{a_i, i \ni \mathbb{N}\}$, there has to be some i such that $M \models \varphi(a_i, b)$, which contradicts the choice of b.

So we have both $\psi(b)$ and $\neg\psi(a, b)$, which is a contradiction, as ψ is supposed to be the φ-definition of p.

$\square$

Remark 3.8. In the last step of the proof of the first implication, given $q \in S_{\varphi^*}(M)$ such that $\varphi(x, q) \in p^*$, we actually proved a stronger statement. We proved that, if for all *positive* formulas $\psi(y) \in q$ we have some $b \in M$ such that $\models \psi(b)$ and $\neg\varphi(x, b) \in p^*$, then $\varphi(x, y)$ is unstable. Thus, for such q we may find positive $\psi_q(y)$ as in (*). Since for proving the claim made in the proof, we took as a formula defining $p^* \restriction \varphi$ a disjunction of such $\psi_q(y)$, we may assert that indeed the defining formula of $p^* \restriction \varphi$ in (ii) can be taken to be positive.

We now consider the consequences of 2.7 when $\varphi(x; y)$ is stable for T.

Proposition 3.9. *Let $\varphi(x; y)$ be an L-formula. Then the following are equivalent:*

1. *$\varphi(x; y)$ is stable (for T).*
2. *For all $M \models T$ and $p(x) \in S_\varphi(M)$, $p(x)$ is definable, i.e. there is some L_M-formula $\psi(y)$ such that for all $b \in M$, $M \models \psi(b)$ if and only if $\varphi(x, b) \in p(x)$.*

3. *for all cardinal $\lambda \geq |T|$ and $M \models T$ of cardinality λ, we have $|S_\varphi(M)| \leq \lambda$.*

4. *There is $\lambda \geq |T|$ such that for all $M \models T$ of cardinality λ, we have $|S_\varphi(M)| \leq \lambda$.*

Proof. 1 $\Rightarrow$ 2: This is a consequence of Proposition 3.7. Indeed, fix $p(x) \in S_\varphi(M)$. Let $p'(x) \in S_x(M)$ extend $p(x)$. Let $M \prec M^*$ be $|M|^+$-saturated and $p^* \in S_x(M)$ a coheir of p' (that is, extending p' and finitely satisfiable in M). By Proposition 3.7, we obtain that $p^*|_\varphi$ is definable over M, so $p(x) = p^*|_\varphi|_M$ is definable.

2 $\Rightarrow$ 3: By assumption, each $p(x) \in S_\varphi(M)$ is determined by its definition $\psi_p(y) \in L_M$. But there are at most λ choices for such a definition, hence the result.

3 $\Rightarrow$ 4: Immediate.

4 $\Rightarrow$ 1: We prove the contrapositive, so assume $\varphi(x; y)$ is unstable. We denote $\lambda = |T|$. Let $\mu \leq \lambda$ be the least cardinal such that $2^\mu > \lambda$ (so in particular $\mu \leq \lambda$).

Recall that ${}^\mu 2$ denotes the set of functions $\mu \to 2 = \{0, 1\}$. For $f, g \in^\mu 2$, we define $f < g$ if and only if there exists an ordinal $\alpha < \mu$ such that $f|_\alpha = g|_\alpha$ but $f(\alpha) = 0$, $g(\alpha) = 1$, it is the lexicographical orderings on ${}^\mu 2$.

Let $X \subseteq {}^\mu 2$ be the set of eventually constant functions in ${}^\mu 2$. By choice of μ we see that $|X| \leq \lambda \cdot \mu = \lambda$. Moreover one easily checks that X is dense in ${}^2\mu$.

Recall that we assumed that φ is unstable, hence we get sequences $(a_i)_{i \in \mathbb{N}}$ and $(b_j)_{j \in \mathbb{N}}$ in $\bar{M}$ such that $\varphi(a_i, b_j)$ if and only if $i \leq j$. Hence by compactness, we can find $\{a_f, b_f, f \in X\}$ in M such that $\varphi(a_f, a_g)$ if and only if $f \leq g$. By downward Löwenheim-Skolem, we can find such a sequence in $M \prec \overline{M}$ of cardinality λ. For each f, g in ${}^\mu 2$ such that $f \neq g$, we have that $\mathrm{tp}(a_f/M) \neq \mathrm{tp}_\varphi(a_g/M)$. Indeed, if for example $f < g$, we can find some f_1 so that $f < f_1 < g$ and therefore $\models \varphi(a_f, b_{f_1}) \wedge \neg\varphi(a_g, b_{f_1})$. Hence $|S_\varphi(M)| \geq 2^\mu > \lambda$. $\square$

For the next results, we will need a new tool : Cantor-Bendixson rank, of CB rank. Let us define it :

Definition 3.10 (Cantor-Bendixson Rank)**.** Let X be a topological space. The Cantor Bendixson rank is a function $CB_X : X \to \mathrm{On} \cup \{\infty\}$ (where On is the class of ordinals). Let $p \in X$, then :

(i) $CB_X(p) \geq 0$,

(ii) $CB_X(p) = \alpha$ if $CB_X(p) \geq \alpha$ and p is isolated in the (closed) subspace $\{q \in X : CB_X(q) \geq \alpha\}$.

(iii) $CB_X(p) = \infty$ if $CB_X(p) > \alpha$ for every ordinal α.

For example, $CB_X(p) = 0$ if p is isolated, equivalently if $\{p\}$ is open. $CB_X(p) \geq 1$ otherwise.

Note that *(ii)* claims that the subspace $\{q \in X : CB_X(q) \geq \alpha\}$ is closed for all α. This is a consequence of the fact that the set of isolated points of any topological space form an open set, as a union of open sets.

Fact 3.11. *Suppose X is compact and $CB_X(p) < \infty$ for every p in X. Then there exists a maximal element α of $\{CB_X(p) : p \in X\}$ and $\{p \in X : CB_X(p) = \alpha\}$ is finite and non empty.*

Proof. Assume there is no maximal element. Then, for each ordinal α, there exists some p_α in X such that $CB_X(p_\alpha) > \alpha$. The set $\{p_\alpha, \alpha \in \mathrm{On}\}$ must have an accumulation point p in the compact set X, which cannot be isolated in any of the $\{q \in X : CB_X(q) \geq \alpha\}$. Hence $CB_X(p) = \infty$, a contradiction.

Let $\alpha = \sup\{CB_X(p) : p \in X\}$. We want to show that $X_\alpha = \{p \in X : CB_X(p) = \alpha\}$ is non-empty. By way of contradiction, assume it is empty. If $\alpha = 0$, this is an obvious contradiction. Else, there are two possibilities.

If α is a successor ordinal, then $\sup\{\beta < \alpha\} < \alpha$, which, if X_α is empty, contradicts that $\alpha = \sup\{CB_X(p) : p \in X\}$.

Else, the ordinal α is limit. For each β less than α, we consider $X_{<\beta} = \{p \in X : CB_X(p) < \beta\}$. Now, the collection $\mathcal{C} = \{X_\beta : \beta < \alpha\}$ is an open cover of X which clearly has no finite subcover, as α is a limit ordinal. This contradicts the assumption that X is compact and therefore X_α is non-empty.

The subset $\{p \in X : CB_X(p) \geq \alpha\}$ is closed, so compact. Since α is maximal, all points in $\{p \in X : CB_X(p) \geq \alpha\}$ are isolated. Therefore, $\{p \in X : CB_X(p) \geq \alpha\}$ is finite. □

Lemma 3.12. *Suppose $\varphi(x, y)$ is stable for T. Let $M \models T$. Let $X = S_\varphi(M)$. Then, $CB_X(p) < \infty$ for each p in X.*

Proof. Define $X_\alpha = \{p \in X : CB_X(p) \geq \alpha\}$. Assume that there exists some q such that $CB_X(q) = \infty$, then for some α, $X_\alpha \neq \emptyset$ and has no isolated points. Indeed, if not, then each X_α has at least one isolation point, and we could conclude that $CB_X(p) \leq |X|$ for any $p \in X$.

We now fix such an α. Since there are no isolated points in X_α, we can find $p_0, p_1 \in X_\alpha$ where $p_0 \neq p_1$. Since $S_\varphi(M)$ is Hausdorff, we can find $\psi_0(x)$ such that $\psi_0(x) \in p_0$ and $\psi_1(x) = \neg\psi_0(x) \in p_1$. Notice that $\{p : p \in X_\alpha\} \cap \{p \in S_\varphi(M) : \psi_0(x) \in p\}$ and $\{p : p \in X_\alpha\} \cap \{p \in S_\varphi(M) : \psi_1(x) \in p\}$ have no isolated points. We will apply this construction recursively to construct a tree of formulas.

Let $\psi_\eta(x)$, for some $\mu \in 2^{\leq\omega}$ finite sequence, be a formula such that $A_\eta = \{p : p \in X_\alpha\} \cap \{p \in S_\varphi(M) : \psi_\eta(x) \in p\}$ is non-empty and contains no isolated points. Then, just as we did in the previous paragraph, we can find $p_{\eta 0}, p_{\eta 1}$ in A_η and a formula $\psi_{\eta 0}(x)$ such that $\psi_{\eta 0}(x) \in p_{\eta 0}$ and $\psi_{\eta 1}(x) = \neg\psi_{\eta 0}(x) \in p_{\eta 1}$.

Hence, we can inductively build a tree $\{\psi_\mu, \mu \in 2^{\leq\omega}\}$ of formula. By construction, each path in T is consistent.

We will now show that we can find $2^{\aleph_0}$ many types over a countable set. If γ is a path in $2^{\leq\omega}$, we let $A_\gamma = \{\psi_\eta(x) : \eta \in \gamma\}$. Let γ, ρ be two distinct paths in T. Let n be the smallest natural number such that $\eta 0 \in \gamma$, $\eta 1 \in \rho$, and length of η is n. $\psi_0(x) \in A_\gamma$ and $\psi_1(x) = \neg\psi_0(x) \in A_\rho$. Therefore, $A_\gamma \cup A_\rho$ is inconsistent. Since each A_γ may be extended to a complete φ-type, no two paths may extend to the same φ-types, and since there are $2^{\aleph_0}$ many paths in $2^{\leq\omega}$, we have that that $|S_\varphi(M)| \geq 2^{\aleph_0}$.

We now have to do this construction over a countable set of parameters. Let C be the collection of parameters contained in each ψ_η for $\eta \in 2^{\leq \omega}$, this is a countable set. Consider L_φ, which is the language L restricted to φ, it is a countable language. By the downward Löwenheim-Skolem theorem applied to L_φ and the L_φ structure M, we note that we can find some countable $N \prec M$ containing C. By the argument above, we have that $|S_\varphi(N)| \geq 2^{\aleph_0}$, but $|N| = \aleph_0$. This contradicts the assumption that φ is stable. □

Remark 3.13. A similar proof shows that if $\varphi(x; y)$ is stable in M, then every p in $S_\varphi(M)$ has CB-rank, using 3.7.

We have seen that if $\varphi(x; y)$ is stable, $p(x) \in S_\varphi(M)$, and $M \models T$, then $p(x)$ is definable. We will now work over sets instead of models, which will matters more complex.

Definition 3.14. If $A \subset M$ we say that $p(x) \in S_\varphi(M)$ is definable over A if there is $\psi(y)$, a definition for $p(x)$ (i.e. for every b we have that $\models \psi(b)$ if and only if $\varphi(x, b) \in p(x)$), which is equivalent to a L_A-formula (not necessarily a φ^*-formula of A).

Lemma 3.15. *Assume $\varphi(x, y)$ is stable. Let $M \models T$, $A \subseteq M$ and let $p(x) \in S_x(A)$. Then there is $q(x) \in S_\varphi(M)$ such that $p(x) \cup q(x)$ is consistent and $q(x)$ is definable 'almost over A', i.e. over $\mathrm{acl}^{eq}(A)$.*

Proof. We may assume that M is $||T| + |A||^+$-saturated and strongly homogeneous, since every model of T is contained in such a model.

Consider the set $X = \{q(x) \in S_\varphi(M) : p(x) \cup q(x)$ is consistent$\}$. Let $f_1 : S_x(A) \to S_\varphi(A)$ and $f_2 : S_x(M) \to S_\varphi(M)$ be the obvious restriction maps, then $X = f_2^{-1}(f_2(p))$. Since these maps are continuous, X is closed and therefore compact. Moreover, all p in $S_\varphi(M)$ have CB-rank, so each element q in X has CB_X rank.

We can therefore apply Fact 3.11 to X. Let $X_0 \subseteq X$ be the finite set of elements of max CB_X-rank α. Let $q(x) \in X_0$ and let $\psi(y)$ be its definition (i.e. $M \models \psi(b)$ if and only if $\varphi(x, b) \in q$), which is a priori over M. Note that $\mathrm{Aut}(M/A)$ acts continuously on X, and X_0 is invariant under this action. Hence $q(x)$ has only finitely many images under $\mathrm{Aut}(M/A)$. So $\psi(y)$ has finitely many images under $\mathrm{Aut}(M/A)$, and Lemmas 2.13 and 2.35 yield that $\psi(y)$ is almost over A. □

Remark 3.16. When $A = M_0 \prec M$, the previous lemma is immediate from 3.7.

Corollary 3.17. *If $A = \mathrm{acl}^{eq}(A)$ and $A \subset M$, every φ-type over A has an A-definable extension to M.*

We will now prove various properties of the formulas defining types, which will translate later into properties of forking, the fundamental tool in studying stable theories.

Theorem 3.18 (Symmetry)**.** *Let $\varphi(x,y)$ be stable, $p(x) \in S_\varphi(M)$ and $q(y) \in S_{\varphi^*}(M)$. Let $\psi(y)$ be the φ^*-formula defining $p(x)$ and $\chi(x)$ the φ-formula defining $q(y)$. Then $\psi(y) \in q(y) \Leftrightarrow \chi(x) \in p(x)$.*

Proof. Note first that p, q are finitely satisfiable in M, so ψ, χ exist and are φ^*, φ formulas, respectively. It is enough to show that $\psi \in q \to \chi \in p$, the other implication being prove exactly the same way.

Let $M \prec M^*$ where M^* is $|M|^+$-saturated. Extend p to $\bar{p} \in S_x(M)$ and then to $p^* \in S_x(M^*)$, a finitely satisfiable type in M, which we can do by Lemma 2.8. The φ-type $p^* \upharpoonright \varphi$ is then definable over M by some $\psi'(y)$. This formula also defines $p = (p^* \upharpoonright \varphi) \upharpoonright M$, so ψ and ψ' are logically equivalent. Therefore ψ defines $p^* \upharpoonright \varphi$ as well.

Let a^* realize p^*. Since $p = \mathrm{tp}_\varphi(a^*/M)$, it is enough to show that $M^* \models \chi(a^*)$. For the sake of a contradiction, suppose not, so $M^* \models \neg\chi(a^*)$.

Let $b \in M^*$ realize q. Since $\psi \in q$, we know that $M^* \models \psi(b)$, and because ψ defines p^*, we have $M^* \models \varphi(a^*, b)$. Then $M^* \models \varphi(a^*, b) \wedge \neg\chi(a^*)$, so by finite satisfiablility of $p^* = \mathrm{tp}_\varphi(a^*/M^*)$, there is an $a \in M$ such that $M^* \models \varphi(a,b) \wedge \chi(a)$. Therefore $\varphi(a, y) \in q(y)$ and also $M^* \models \neg\chi(a)$, which contradicts the fact that χ defined $q \in S_{\varphi^*}(M)$. $\square$

Remark 3.19. The above theorem can be also be proved just by using the stability of $\varphi(x,y)$ in M.

Definition 3.20.

1. Let $A \subseteq \bar{M}$ (or $\bar{M}^{eq}$), and let $\varphi(x,y)$ be a formula. By a φ-formula over A, we mean a φ-formula which is equivalent to a φ-formula with parameters in A.
2. Any φ-type p over M restricts to a $\varphi-$type over A, denoted $p \upharpoonright A$, consisting just of the formulas in p with parameters from A. If $p = \mathrm{tp}_\varphi(a/M)$, we write $\mathrm{tp}_\varphi(a/A)$ for $p \upharpoonright A$, and let $S_\varphi(A)$ be the set of all such types.
3. A φ-type p over M is said to definable over A if it is definable by a φ^*-formula over A.

Lemma 3.21 (Uniqueness)**.** *Let $\varphi(x,y)$ be stable in M and let $A \subseteq M$. If p_1, p_2 are φ-types, definable over A and $p_1 \upharpoonright \mathrm{acl}^{eq}(A) = p_2 \upharpoonright \mathrm{acl}^{eq}(A)$, then $p_1 = p_2$.*

Proof. It is enough to show that for all $b \in M$, $\varphi(x,b) \in p_1 \Leftrightarrow \varphi(x,b) \in p_2$. Fix an arbitrary $b \in M$, let $q(y) = \mathrm{tp}_{\varphi^*}(b/A)$ and it extend to $q' \in S_x(A)$. By Lemma 3.15, there is some $q'' \in S_{\varphi^*}(M)$ consistent with q and almost definable over A. By maximality of types, $q \subseteq q''$ and $q = q'' \upharpoonright \varphi \upharpoonright A$.

Let $\chi(x)$ be the defining formula for q'', which equivalent to some χ' definable over $\mathrm{acl}^{eq}(A)$. Let ψ_i be the φ^*-formula defining $p_i(x)$ over A. Then

$$\begin{aligned}\varphi(x,b) \in p_1 \Leftrightarrow & M \models \psi_1(b) \\ \Leftrightarrow & \psi_i(y) \in q \text{ (by definition of q)} \\ \Leftrightarrow & \psi_1(y) \in q''\end{aligned}$$

$$\begin{aligned}
&\Leftrightarrow \chi(x) \in p_1 \text{ (by 3.18)}\\
&\Leftrightarrow \chi(x) \in p_2 \text{ (since } p_1 \upharpoonright \operatorname{acl}^e q(A) = p_2 \upharpoonright \operatorname{acl}^e q(A))\\
&\Leftrightarrow \psi_2 \in q''\\
&\Leftrightarrow \psi_2 \in q\\
&\Leftrightarrow M \models \psi_2(b)\\
&\Leftrightarrow \varphi(x,b) \in p_2.
\end{aligned}$$

□

Remark. Together with existence, uniqueness shows that if A is an algebraically closed subset of M and $p \in S_\varphi(A)$ then p has a unique extension $p'(x) \in S_\varphi(M)$ which is definable over A.

Theorem 3.22 (Conjugacy and the Finite Equivalence Relation Theorem)**.** *Suppose that $A \subseteq M$, $p \in S_\varphi(A)$, and X is the set of all $\varphi-$types extending q which are definable almost over A. Then*

1. *X is finite*
2. $\operatorname{Aut}(M/A)$*, the group of automorphisms of M which fix A point-wise, acts transitively on X if M is sufficiently homogeneous and saturated.*
3. *There exists an A-definable equivalence relation E such that for any $q_1, q_2 \in X$, we have $q_1 = q_2$ if and only if $q_1(x) \cup q_2(x) \models E(x_1, x_2)$.*

Proof. Note first that by uniqueness, each $q \in S_\varphi(M)$ which is almost definable over A is the unique extension of $q \upharpoonright \operatorname{acl}^{eq}(A)$, so X is in bijective correspondence with $Y := \{q \upharpoonright \operatorname{acl}^{eq}(A) : q \in X\}$. Moreover, Y is still acted upon by $\operatorname{Aut}(M/A)$.

Proof of (ii). It suffices to show that $\operatorname{Aut}(M/A)$ acts transitively on Y. Recall that an action of a group G on a set Y is transitive if for all $y_1, y_2 \in Y$, some element of $\operatorname{Aut}(M/A)$ sends y_1 to y_2. So let $q_1, q_2 \in Y$.

Take $p' \in S_x(M)$ extending p. We claim that each $p' \cup q_i$ is consistent. Let this be granted for now. If so, then let $m_1 \in M$ realize $p' \cup q_1$ and $m_2 \in M$ realize $p' \cup q_2$, by saturation. In particular, this means that $q_1 = \operatorname{tp}_\varphi(m_1/\operatorname{acl}^{eq}(A))$ and $q_2 = \operatorname{tp}_\varphi(m_2/\operatorname{acl}^{eq}(A))$.

Since $\operatorname{tp}(m_1/A) = \operatorname{tp}(m_2/A)$, by homogeneity there is some automorphism $\sigma \in Aut(M/A)$ with $\sigma(m_1) = m_2$. But $\operatorname{acl}^{eq}(A)$ is invariant (as a set) under the action of σ since it fixes A pointwise. Therefore σ sends every statement true of m_1 with parameters from $\operatorname{acl}^{eq}(A)$ to a statement true of m_2 with (possibly different) parameters from $\operatorname{acl}^{eq}(A)$, so $\sigma(q_1) \subseteq q_2$ and by maximality, $\sigma(q_1) = q_2$. Therefore $\operatorname{Aut}(M/A)$ acts transitively on Y.

We will now prove the consistency claim.

Claim. *For each i, $p' \cup q_i$ is consistent.*

Proof of claim. Otherwise, by compactness, there is a formula $\sigma(x) \in q_i(x)$ such that $p' \cup \sigma$ is inconsistent. Hence each A-conjugate $\sigma'(x)$ of σ over A is inconsistent with $p'(x)$.

But σ is a formula over $\mathrm{acl}^{eq}(A)$ and as such, as only finitely many $\mathrm{Aut}(M/A)$ conjugates. Let $\psi(x)$ be the disjunction of these, it is inconsistent with p' as well.

But $\psi(M)$ is $\mathrm{Aut}(M/A)$ invariant by construction, hence ψ is over A. Moreover, it is implied by $\sigma(x)$, and thus it is in $p(x)$. But it is inconsistent with p', which extends p, a contradiction.

□

□

Proof of (i). It suffices to show that Y is finite. Let $q \in Y$, it is $\mathrm{acl}^{eq}(A)$-definable, and so has only finitely many images under $\mathrm{Aut}(M/A)$. Since the action of $\mathrm{Aut}(M/A)$ is transitive, Y is the orbit of q, which is finite. □

Proof of (iii). Each two types $q_i, q_j \in Y$ are separated by some formula $\theta_{i,j}(x)$, i.e. $\theta_{i,j}(x) \in q_i$ and $\neg\theta_{i,j}(x) \in q_j$. Consider the set of formulas $\{\theta_{i,j} | q_i, q_j \in Y, i \neq j\}$, it is finite. We can close this set under $\mathrm{Aut}(M/A)$, to obtain a new set of formula, which we denote Θ, finite as well.

Let $E(x_1, x_2)$ be the equivalence relation defined by $E(x_1, x_2)$ if and only if $\bigwedge_{\theta \in \Theta} \theta(x_1) \leftrightarrow \theta(x_2)$. It is easily checked that this equivalence relation has the required property.

□

We will now introduce notions of forking and dividing, which are essential to develop stability theory.

Definition 3.23. Let T be a complete theory, let $A \subset \bar{M}$, and let $\varphi(x, b)$ be a formula, where $b \in \bar{M}$.

(i) We say that $\varphi(x, b)$ divides over A if there is an A-indiscernible sequence $(b_i)_{i \in \mathbb{N}}$ such that $b_0 = b$ and the set $\{\varphi(x, b_i) : i \in \omega\}$ is inconsistent. By compactness and indiscernibility, we can see that it is the same as saying the above set is k-inconsistent for some k.

(ii) We say that $\varphi(x, b)$ forks over A if there are $\psi_1(x, b_1), ..., \psi_k(x, b_k)$ each dividing over A and $\models \varphi(x, b) \to \bigvee_i \psi_i(x, b_i)$.

Remark. If $b \in A$, then $\varphi(x, b)$ does not divide over A, as any A-indiscernible sequence containing b will just repeat b.

Remark 3.24. We can extend the above definition to partial types over B. Let $\pi(x)$ be a partial type over B, which we assume, without loos of generality, to be closed under conjunction. Then $\pi(x)$ divides/forks over A if there is $\varphi(x) \in \pi(x)$ that divides/forks over A.

For the benefit of the reader seeing these definitions for the first time, let us try to explain the intuition behind the notion of dividing.

Note that since $(b_i)_{i \in \mathbb{N}}$ is A-indiscernible, each b_i is the image, under an A-automorphism, of b. Hence, we can obtain "copies" $\varphi(\bar{M}, b_i)$ of the definable set $\varphi(\bar{M}, b)$. Dividing states that there is $k \in \mathbb{N}$ such that the intersection of any k different such copies is empty.

Now think about the type definable set given by $X = \{\exists y \varphi(x, y)\} \cup \mathrm{tp}(b/A)$. It contains all the copies of $\varphi(\bar{M}, b)$, which are k-inconsistent. What this means, geometrically, is that the "dimension" of $X(\bar{M})$ is strictly bigger than that of $\varphi(\bar{M}, b)$.

Example. Let k be a small subfield of $\mathbb{C}$. Let V be a variety, which we can identify it with a definable set. Then V does not divide over k if and only if V does not fork over k if and only if V is defined over $\mathrm{acl}(k)$.

Example. Let T be the theory of an equivalence relations with infinitely many infinite classes. Then for every b, he formula $E(x, b)$ divides over $\emptyset$. Moreover, in this theory, let $p \in S_E(M)$ and $A \subset M$. Then p is defined almost over A if and only if p does not divide over A if and only if p does not fork over A.

The following lemma will be useful in proofs :

Lemma 3.25. *Let $(b_i)_{i \in \mathbb{N}}$ be an A-indiscernible sequence. Then for all i, we have $\mathrm{tp}(b_i/\mathrm{acl}^{eq}(A)) = \mathrm{tp}(b_0/\mathrm{acl}^{eq}(A))$.*

Proof. Suppose not. Then there exists an A-definable finite equivalence relation $E(y, z)$ and some $i > 0$ such that $\neg E(b_0, b_i)$. By indiscernibility, this implies that $\neg E(b_j, b_k)$ for all $j, k \in \mathbb{N}$, which contradicts E being finite.

□

Corollary 3.26. *If $\varphi(x, b)$ is over $\mathrm{acl}^{eq}(A)$, then it does not divide over A.*

We will also need canonical extensions of definable types :

Lemma 3.27. *Let $p \in S_\varphi(M)$ be a definable type, and $M \prec M'$ an elementary extension of M. There is an unique extension $p' \in S_\varphi(M')$ with same definition as p.*

Proof. Given existence, unicity is immediate. Let $\psi(y)$ de the defining formula of p, and consider the family of formulas $p' = \{\varphi(x, b) | b \in M', M' \models \psi(b)\}$. We need to show this is maximal and consistent, which will yield the required extension of p.

Maximality is immediate, so let us prove consistency. Suppose that it isn't, then there are $\varphi(x, b_1), \cdots, \varphi(x, b_n)$ formulas in p', which are inconsistent. Hence, we have $M' \models \exists y_1, \cdots y_n((\bigwedge_{i=1}^{n} \psi(y_i)) \wedge (\neg(\exists x \bigwedge_{i=1}^{n} \varphi(x, y_i))))$, and since it is an elementary extension, the model M will satisfy this formula as well, which would imply inconsistency of p, a contradiction.

□

Lemma 3.28. *Let $\varphi(x, y)$ be stable, and $b \in \bar{M}$. The following are equivalent:*

(1) There is some M containing $A \cup \{b\}$ and $p(x) \in S_\varphi(M)$ such that $p(x)$ is defined almost over A and $\varphi(x, b) \in p$.

(2) $\varphi(x,b)$ does not divide over A.

Proof. (i) $\Rightarrow$ (ii) We may assume that M is sufficiently saturated($(|T|+|A|)^+$-saturated will suffice). Let $(b_i)_{i\in\mathbb{N}}$ be an A-indiscernible sequence in $\bar{M}$, with $b_0 = b$. We need to show that $\{\varphi(x,b_i) : i \in \omega\}$ is consistent.

Clearly, the above only depends on the infinite type $\mathrm{tp}((b_i)_{i\in\mathbb{N}}/A)$. By saturation, we may assume that the b_i's are in M. By assumption, there is $p(x) \in S_\varphi(M)$ such that $p(x)$ is defined almost over A and $\varphi(x,b) \in p$. Let $\psi(y)$ be the φ-definition of p, note that ψ is over $\mathrm{acl}^{\mathrm{eq}}(A)$. Furthermore, we know that $M \models \psi(b)$, which implies $M \models \psi(b_i)$ for all i by the Lemma 3.25. Hence $\varphi(x,b_i) \in p$ for all i, and in particular the $\varphi(x,b_i)$ form a consistent collection of formulas.

(ii)$\Rightarrow$(i) We will show the contrapositive. We may assume that A is algebraically closed. Hence we are given $\varphi(x,b)$ such that any $p(x) \in S_\varphi(M)$ definable over A contains $\neg\varphi(x,b)$. We want to show $\varphi(x,b)$ divides over A.

Let $q(y) \in S_y(M)$ be an extension of $\mathrm{tp}(b/A)$ with $q|\varphi^*$ definable over A (it exists by Corollary 3.17). Let M^* be a sufficiently saturated model extending M and let $q^*(y)$ be a coheir extension of q to M^*. By Lemma 3.7, we know that $q^*|\varphi^*$ is definable over M. Note further that it is actually definable over A since $q^*|\varphi^*|M = q|\varphi^*$ and $q|\varphi^*$ is definable over A. Let $\chi(x)$ be the φ^*-definition of q^* and q over A.

Consider a sequence $(b_i)_{i\in\mathbb{N}}$ such that $b_i \models q^*|M \cup \{b_0,...,b_{i-1}\}$. We see that $(b_i)_{i\in\mathbb{N}}$ is M-indiscernible, hence A-indiscernible. To complete the proof, it will be sufficient to show the following :

Claim. *The set of formulas $\{\varphi(x,b_i) : i \in \mathbb{N}\}$ is inconsistent.*

Proof of claim. If the above is consistent, let $a \models \varphi(x,b_i)$ for all i. Let $p = \mathrm{tp}_\varphi(a/A)$. Let $p^*(x) \in S_\varphi(M^*)$ be an extension of $p(x)$ that is definable over A. Let $B = \mathrm{acl}^{eq}(A \cup (b_i)_{i\in\mathbb{N}})$.

Since $p \cup \{\varphi(x,b_i) : i \in \mathbb{N}\}$ is consistent, there is an extension to $r(x) \in S_\varphi(B)$. By existence, we can further get $r \subset r^*(x) \in S_\varphi(M^*)$ that is definable over B. Let $\psi_1(y)$ be φ-definition of p^*, which can be assumed to be a φ^*-formula over A. And let $\psi_2(y)$ be the φ-definition of $r^*(X)$, which can be assumed to be over A.

Recall that $\chi(x)$, the definition of q and q^*, is a φ-formula over A. Note that both $p^*(x)$ and $r^*(x)$ contain $p(x) \in S_\varphi(A)$. Hence $\chi(x) \in p^*$ if and only if $\chi(x) \in r^*$. By symmetry it follows that $\psi_1(y) \in q^*|\varphi^*$ if and only if $\psi_2(y) \in q^*|\varphi^*$.

By assumption, and because p^* is definable over A, we have that $\neg\varphi(x,b) \in p^*$, hence $\models \neg\psi_1(b)$ since $\psi_1(y)$ defines p^*. Moreover $\psi_1(y)$ is a φ^*-formula over A, so $\neg\psi_1(y) \in q^*|\varphi^*$. Therefore $\neg\psi_2(y) \in q^*|\varphi^*$.

But ψ_2 defines r^* over B, say $\psi_2(y)$ is over $e \in B$, with $e \in \mathrm{acl}^{eq}(A, b_1, ..., b_k)$ for some $k \in \mathbb{N}$. Since $b_{k+1} \models q^*|M \cup \{b_0, \cdots, b_k\}$, we have $\models \neg\psi_2(b_{k+1})$. Since $\psi_2(y)$ defines r^*, we obtain $\neg\varphi(x, b_{k+1}) \in r^*$, a contradiction to the choice of r^*. $\square$

This finishes the proof, as $(b_i)_{i\in\mathbb{N}}$ is the required indiscernible sequence. $\square$

With the same method, we can conclude :

Lemma 3.29. *Assume $\varphi(x,y)$ is stable, fix A small and $p \in S_\varphi(A)$. Let $M \supset A$ and $b \in M$. The following are equivalent :*

1. *(i) $\varphi(x,b) \in p^*(x)$ for some $p^*(x) \in S_\varphi(M)$ extending p and definable almost over A.*
2. *(ii) $p(x) \cup \varphi(x,b)$ does not divide over A.*

Here, by non-dividing we mean that for every $b = b_0, b_1, ..., b_n, ...$ A-indiscernible sequence, we have some $a \in \bar{M}$ and $a \models p(x) \cup \{\varphi(x,b_i) : i \in \omega\}$.

We just proved a few interesting facts about dividing, but what about forking ? Recall that a formula $\varphi(x,\bar{c})$ is said to fork over A if there are $\varphi_1(x,\bar{c}_1), \cdots, \varphi_r(x,\bar{c}_r)$ such that $\models \forall x(\varphi(x,\bar{c}) \to \bigvee_{i=1}^{r} \varphi_i(x,\bar{c}_i))$, and each φ_i divides over A.

Similarly, if $\Phi(x,\bar{c})$ is a partial type, we say it forks over A if $\Phi(x,\bar{c})$ implies a finite disjunction of formulas, each dividing over A.

Remark that dividing trivially implies forking, but the converse is not true in general.

Example 3.30. Consider the unit circle $\mathbb{S}^1$, together with the ternary B relation of betweeness : we say that $B(x,y,z)$ is the counterclockwise arc from x to z goes through y (we leave the writing of a more rigorous formulation to the reader).

Then the formula $x = x$ does not divide over the empty set, as it is over the empty set. However, it does fork over the empty set. Indeed, it implies, for any $a, b \in \mathbb{S}^1$, the disjunction $B(b,x,a) \vee B(a,x,b)$, and each of these two formulas is easily verified to fork over the empty set.

However, the equivalence between forking and dividing will be true for stable formulas, using a local definition of forking. Let us start with :

Remark 3.31. Suppose the partial type $\Phi(x,\bar{c})$ does not fork over A. Let $\overline{M}$ be a saturated model containing $\bar{c}$. Then there is a complete type $p(x) \in S(\overline{M})$, containing $\Phi(x,\bar{c})$ and not dividing over A.

Proof. Consider the set of formulas $\Sigma = \Phi(x,\bar{c}) \cup \{\neg\varphi(x), \varphi \text{ over } M, \varphi \text{ divides over } A\}$. It is consistent. Indeed, if it wasn't, by compactness we would have $\Phi(x,\bar{c}) \models \varphi_1(x) \vee \cdots \vee \varphi_n(x)$, for some φ_i over $\overline{M}$, dividing over A.

Therefore, we can extend Σ to a complete type, which is the type we were looking for.

$\square$

Now define what we mean by local forking :

Definition 3.32. Fix a formula $\varphi(x, y)$, and let $\varphi(x, b)$ be an instance of it. We say $\varphi(x, b)$ φ-forks over A if it implies a finite disjuction of φ-formulas, each dividing over A.

We can now state and prove the announced result :

Theorem 3.33. *Suppose $\varphi(x, y)$ is stable. Then for any b, the formula $\varphi(x, b)$ divides over A if and only if it φ-forks over A.*

Proof. The left to right direction is immediate. For the other direction, suppose that $\varphi(x, b)$ does not divide over A. Let M^* be a $(|T| + |A|)^+$-saturated model containing $A \cup \{b\}$. By Lemma 3.28, there is $p(x) \in S_\varphi(M^*)$ definable almost over A, containing $\varphi(x, b)$. In particular it is invariant under $\mathrm{Aut}(M^*/\operatorname{acl}^{eq}(A))$, because any $\sigma \in \mathrm{Aut}(M^*/\operatorname{acl}^{eq}(A))$ will fix the definition of p.

By way of contradiction, suppose that $\varphi(x, b) \models \psi_1(x, b_1) \vee \ldots \vee \psi_r(x, b_r)$, with each ψ_i a φ-formula dividing over A. Note that if $\mathrm{tp}(b_i/Ab) = \mathrm{tp}(b'_i/Ab)$ for some b'_i, then $\psi(x, b'_i)$ divides over A as well. Hence, we can assume that $b_i \in M^*$ for all i, by saturation of M^*.

As $\varphi(x, b) \in p$, we have $\psi_i(x, b_i) \in p$ for some i. But again by Lemma 3.28, this implies that $\psi_i(x, b_i)$ does not divide over A, a contradiction.

□

We now return to ranks, specifically the Cantor-Bendixson rank. Recall that this rank was defined in 3.10, and some of its elementary properties were proven. For stable formulas, the Cantor-Bendixson rank is finite :

Lemma 3.34. *Suppose $\varphi(x, y)$ is stable, and let $X = S_\varphi(\overline{M})$. Then $\mathrm{CB}(X)$ is finite.*

Proof. Suppose for contradiction that there is $p \in S_\varphi(\overline{M})$ with $\mathrm{CB}(p) \geq \omega$. As in lemma 3.12, we can build, for any n, a tree of formulas $\{\varphi_\mu(x), \mu \in^{n\geq} 2\}$ such that $\varphi_\emptyset(x)$ is the true formula, for $0 < l(\mu) < n$, the formulas $\varphi_{\mu\wedge 0}$ and $\varphi_{\mu\wedge 1}$ are of the form $\varphi(x, b), \neg\varphi(x, b)$ or $\neg\varphi(x, b), \varphi(x, b)$ and moreover, for each $\mu \in^n 2$, the set of formulas $\{\varphi_{\mu|_i}(x), i \leq n\}$ is consistent.

By compactness and saturation, we find a tree $\{\varphi_\mu, \mu \in^{\omega\geq} 2\}$ with consistent branches. Since we only needed the formula φ, and a countable set of parameters, to define all the formulas in this tree, this contradicts stability of φ (see the proof of 3.12 for more details on this type of argument).

□

The Cantor-Bendixson rank can be extended to define a rank for partial types :

Definition 3.35. Fix $\varphi(x, y)$ a stable formula. Let Φ be a partial φ-type over some small set of parameters. We define $R_\varphi(\Phi) = \max\{\mathrm{CB}_X(p), p \in S_\varphi(\overline{M}), \Phi \subset p\}$, where $X = S_\varphi(\overline{M})$.

Note that there are only finitely many p realizing this maximum. This rank can be used to witness dividing, more precisely :

Proposition 3.36. *Let $\varphi(x,y)$ be stable, a type $p \in S_\varphi(A)$, with $A \subset M$, a $(|T|+|A|)^+$-saturated model. Let $p'(x)$ be an extension of $p(x)$ over M. Then $p'(x)$ is definable almost over A if and only if $R_\varphi(p) = R_\varphi(p')$.*

To prove this, we will need one more fact about local stability theory.

Lemma 3.37. *Let $\varphi(x,y)$ be stable, and $\varphi(x,b)$ be some instance of it. Then $\varphi(x,b)$ does not divide over A if and only if some finite positive boolean combination of $\mathrm{Aut}(\overline{M}/A)$ conjugates of $\varphi(x,b)$ is consistent and over A.*

Proof. Suppose $\varphi(x,b)$ does not divide over A. Let $M^* \supset A$ be sufficiently saturated, with $b \in M^*$. By lemma 3.28, there is $p^*(x) \in S_\varphi(M^*)$, definable almost over A and containing $\varphi(x,b)$.

Let $\psi(y)$ be the definition of p^*, it is a φ-formula almost over A. Let $q^* \in S_{\varphi^*}(M^*)$ be definable almost over A, consistent with $\mathrm{tp}(b/\operatorname{acl}^{eq}(A))$. Let $\tau(x)$ be the formula, over $\operatorname{acl}^{eq}(A)$, defining q^*.

By remark 3.8, the formula $\tau(x)$ is a positive boolean combination of $\operatorname{acl}^{eq}(A)$-conjugates of $\varphi(x,b)$. Notice that $\models \psi(b)$, so $\psi(y) \in q^*$, and by symmetry, we get $\tau(x) \in p^*$. So τ is consistent.

Let e, χ be such that $\tau(x) \leftrightarrow \chi(x,e)$, with $e \in \operatorname{acl}^{eq}(A)$. Let $e = e_1, \cdots e_n$ be the realizations of $\mathrm{tp}(e/A)$. Then the formula $\chi(x,e_1) \vee \cdots \vee \chi(x,e_n)$ is over A, consistent, and is equivalent to a positive boolean combination of $\varphi(x,b)$, as τ is.

Conversely, let $\tau(x)$ be some consistent, over A, finite positive boolean combination of $\mathrm{Aut}(\overline{M}/A)$ conjugates of $\varphi(x,b)$. Let $p(x) \in S_\varphi(A)$ containing $\tau(x)$. Let M^* be sufficiently saturated so that $\tau(x)$ is a positive boolean combination of $\mathrm{Aut}(M^*/A)$ conjugates of $\varphi(x,b)$. Let $p'(x) \in S_\varphi(M^*)$ extend $p(x)$, definable over A.

So in particular, the formula $\tau(x)$ belongs to $p'(x)$. Since p' is a complete type, this implies that for some b' with $\mathrm{tp}(b'/A) = \mathrm{tp}(b/A)$, the formula $\varphi(x,b')$ belong to p'. Hence $\varphi(x,b')$ does not divide over A, and neither does $\varphi(x,b)$. $\square$

We now are ready to prove proposition 3.36 :

Proof of proposition 3.36. First, suppose that p' is definable almost over A. We can pick $\psi(x) \in p'(x)$ such that $R_\varphi(\psi) = R_\varphi(p')$.

Indeed, if not, for all formulas $\psi(x)$ contained in $p'(x)$, we would have $R_\varphi(\psi) > R_\varphi(p') = \mathrm{CB}(p')$. But because M is $(|A|+|T|)^+$-saturated, we have $R_\varphi(\psi) = \max\{\mathrm{CB}(q), q \in S_\varphi(M), \psi \subset q\}$. Hence, for all formula $\psi(x)$ contained in $p'(x)$, there is a type $q \in S_\varphi(M)$, containing $\psi(x)$, with $\mathrm{CB}(q) > \mathrm{CB}(p)$. Now, if we consider the open sets $\{[\psi(x)], \psi(x) \in p'(x)\}$, this implies, because ordinals are well ordered, that $\bigcap[\psi(x)]$ contains a type $q \in S_\varphi(\overline{M})$, with $\mathrm{CB}(q) > \mathrm{CB}(p)$. But this intersection is precisely $\{p'\}$, a contradiction.

By 3.37 and 3.28, some positive boolean combination of A-conjugates of $\psi(x)$ is over A. Let $\psi_1(x) \vee \cdots \vee \psi_r(x)$ be this boolean combination. Each of the $\psi_i(x)$ is consistent with $p(x)$, as $\psi(x)$ is. Hence, their disjunction also is consistent with $p(x)$. But since it is over A, it has to be equivalent to a formula in $p(x)$, say $\tau(x)$. In particular :

$$\begin{aligned} R_\varphi(p) &\le R_\varphi(\tau) \\ &= R_\varphi(\psi_1(x) \vee \cdots \vee \psi_r(x)) \\ &= \max\{R_\varphi(\psi_i)\} \\ &= R_\varphi(\psi) \\ &= R_\varphi(p') \end{aligned}$$

so $R_\varphi(p) \le R_\varphi(p')$. But we always have, if p' is an extension of p, that $R_\varphi(p) \ge R_\varphi(p')$ (to see this, use the restriction map $\pi : S_\varphi(M) \to S_\varphi(A)$ and an induction on CB rank). Hence $R_\varphi(p) = r_\varphi(p')$.

Conversely, suppose that $p'(x)$ is not definable almost over A. Let $\psi(y)$ be the definition of $p'(x)$, which is therefore not over $\mathrm{acl}^{eq}(A)$. In particular, by saturation of M, it has infinitely many images under $\mathrm{Aut}(M/A)$, giving rise to infinitely many distinct conjugate $p'_i \in S_\varphi(M)$ of p', all containing p.

This implies $\alpha = R_\varphi(p') < R_\varphi(p)$. Indeed, consider the compact set $[p(x)] \subset S_\varphi(M)$. This yields an infinite number of p'_i, all of CB rank α and contained in $[p(x)]$ (their CB rank is equal to their φ rank by saturation of M). Hence, the set $\{p'_i, i \in I\}$ must have an accumulation point q in $[p(x)]$. And because they all have rank α, we get $CB(q) > \alpha$. But we also have $CB(q) < CB(p)$, as q extends p, hence $R_\varphi(p) = CB(p) > \alpha = R_\varphi(p')$.

□

This concludes our exploration of local stability. In the next section, we will apply the tools we developed to stable theories.

3.2 Stable Theories

A theory T is said to be stable if every formula $\varphi(x, y)$ is stable for T. Unless otherwise stated, in this section we will assume that the theory T is stable, and that $\overline{M} \models T$ is some large, sufficiently saturated and homogeneous model.

Let $\Delta = \{\varphi_1(x, y_1), \ldots, \varphi_n(x, y_n)\}$ be a finite set of stable formulas. By a Δ-formula, we mean a boolean combination of instances of $\varphi_i(x, y_i)$, $1 \le i \le n$. For some set A, a complete Δ-type over A, $p(x) \in S_\Delta(A)$ is a maximal, consistent set of Δ-formulas *over* A (i.e. Δ-formulas equivalent modulo T to a formula over A). As usual, if $A = M \models T$, then $p(x) \in S_\Delta(M)$ is determined by the instances of $\varphi_i(x, y_i)$ and $\neg\varphi_i(x, y_i)$ (with parameters in M) that appear in $p(x)$. Everything from Section 2.1 holds for Δ-formulas and Δ-types. In particular, we have

Lemma 3.38. *Let $\Delta(x) = \{\varphi_1(x, y_1), \ldots, \varphi_n(x, y_n)\}$ be a finite set of stable formulas.*

(i) Any $p(x) \in S_\Delta(M)$ is definable (i.e. every $\varphi_i(x, y_i)$ has a defining formula).

(ii) Suppose $q(x) \in S_x(A)$ is a type, $A \subseteq M \models T$. Then there is $p(x) \in S_\Delta(M)$ such that $p(x)$ is definable almost over A and $p(x) \cup q(x)$ is consistent.

From this, we can conclude the following:

Proposition 3.39. *Let T be a stable theory, and let $A = \operatorname{acl}^{eq}(A) \subseteq \overline{M} \models T$ and $p(x) \in S_x(A)$. Then, for any $M \supseteq A$ (in particular, for M sufficiently saturated), there is $p'(x) \in S_x(M)$ such that $p(x) \subseteq p'(x)$ and $p'(x)$ is definable over A. Moreover, $p'(x)$ is the unique such type.*

Proof. Fix $M \supseteq A$. By Lemma 3.15, for every formula $\varphi(x, y)$, there is a unique φ-type $p'_\varphi(x) \in S_\varphi(M)$ that is consistent with $p(x)$ and definable over A. Consider the set of formulas

$$p'(x) = \bigcup_{\varphi(x,y)} p'_\varphi(x).$$

We claim that $p'(x)$ is consistent. By compactness, it is enough to show that for any finite set of formulas $\Delta = \{\varphi_1(x, y_1), \ldots, \varphi_n(x, y_n)\}$, the set

$$p'_\Delta(x) = \bigcup_{i=1}^{n} p'_{\varphi_i}(x)$$

is consistent. By Lemma 3.38*(ii)* there is a Δ-type $q(x) \in S_\Delta(M)$ such that $p(x) \cup q(x)$ is consistent, and such that $q(x)$ is definable over A. By the uniqueness of p'_{φ_i}, we get $q \restriction_{\varphi_i} = p'_{\varphi_i}$ for $1 \leq i \leq n$, and hence $p'_\Delta(x) \subset q(x)$ and so $p'_\Delta(x)$ is consistent.

Is is maximal because each p'_φ is a complete φ-type, and uniqueness comes from the uniqueness of the p'_φ. □

Proposition 3.40. *Let $p(x) \in S_x(A)$, $q(x) \in S_x(B)$ and $p(x) \subseteq q(x)$. The following are equivalent:*

(i) there is $M \supseteq B$ and $q'(x) \in S_x(M)$ such that $q'(x) \supseteq q(x)$ and $q'(x)$ is definable almost over A,

(ii) $q(x)$ does not fork over A.

Proof. Assume *(i)*. As $q'(x)$ is definable almost over A, we can take an extension to some $(|T| + |B|)^+$-saturated model. Hence, we may assume that M is $(|T| + |B|)^+$-saturated.

Let $\varphi(x, b) \in q(x) \subseteq q'(x)$ and suppose, for a contradiction, that $\varphi(x, b)$ forks over A. Then there exists $\psi_1(x, b_1), \ldots, \psi_n(x, b_n)$ over $\overline{M}$, each dividing over A, with $\varphi(x, b) \vdash \bigvee_{i=1}^n \psi_i(x, b_i)$. As M is sufficiently saturated, we may assume that the b_i's are in M. Since $q'(x)$ is complete and $\varphi(x, b) \in q'(x)$, it follows that $\bigvee_{i=1}^n \psi_i(x, b_i) \in q'(x)$ and so there is i such that $\psi_i(x, b_i) \in q'(x)$. Without loss of generality we may assume $i = 1$. Let $b_1^1, b_1^2, b_1^3, \ldots \in M$ be indiscernible over A with $b_1^1 = b_1$. Since $q'(x)$ is definable over $\operatorname{acl}^{eq}(A)$ and since for all $k \geq 1$, $\operatorname{tp}(b_1^k / \operatorname{acl}^{eq}(A)) = \operatorname{tp}(b_1 / \operatorname{acl}^{eq}(A))$ by Lemma 3.25, we get that $\psi_1(x, b_1^k) \in q'(x)$ for all $k \geq 1$. Hence $\{\psi_1(x, b_1^k) : k \geq 1\}$ is consistent, contradicting the fact that $\psi_1(x, b_1)$ divides over A.

For the other direction, assume that $q(x)$ does not fork over A. Let $M \supseteq B$ be sufficiently saturated. Consider the set

$$\Sigma(x) = q(x) \cup \{\neg\psi(x) : \psi(x) \text{ is over } M \text{ and divides over } A\}.$$

We claim that $\Sigma(x)$ is consistent. If not, then by compactness there is a formula $\varphi(x) \in q(x)$ and formulas $\psi_1(x), \ldots, \psi_n(x)$ over M that divide over A such that $\varphi(x) \vdash \bigvee_{i=1}^n \psi_i(x)$, which contradicts the fact that $\varphi(x)$ does not fork over A. Let $q'(x) \in S_x(M)$ be a complete type extending $\Sigma(x)$. Then $q'(x)$ does not divide over A.

It remains to show that $q'(x)$ is almost definable over A. Fix a formula $\varphi(x, y)$. Suppose that $\varphi(x, b) \in q'(x)$ and let $p'(x) = q'(x) \restriction_{\mathrm{acl}^{eq}(A)}$. Since $q'(x)$ does not divide over $\mathrm{acl}^{eq}(A)$, we get that $p'(x) \cup \{\varphi(x, b)\}$ does not divide over $\mathrm{acl}^{eq}(A)$. By Lemma 3.29, $\varphi(x, b) \in p^*(x)$ for some $p^*(x) \in S_\varphi(M)$, definable over $\mathrm{acl}^{eq}(A)$ and extending $p' \restriction_\varphi$ (in fact, p^* is unique and depends only on $p' \restriction_\varphi$), and so $q' \restriction_\varphi = p^*(x)$. Since p^* is definable almost over A, it follows that $q'|\varphi$ is definable almost over A. Since we can do this for any formula φ, we obtain that q is definable almost over A. □

Remark 3.41. Suppose T is stable. Then $q(x) \in S(B)$ forks over A if and only if $q(x)$ divides over A.

Proof. By definition, if $q(x)$ divides over A, then it forks over A. Suppose that $q(x)$ does not divide over A. Then for every formula $\varphi(x, b) \in q(x)$, $\varphi(x, b)$ does not divide over A. Let M^* be a $(|T| + |A| + |B|)^+$-saturated model containing $A \cup B$. Suppose that there is $\varphi(x, b) \in q(x)$ that forks over A. We can mimic the proof of Theorem 3.33 to derive a contradiction. □

Here we introduce Makkai's anchor notation for independence. We write $\bar{a} \mathop{\smile\!\!\!\!\!\!\,|}_B C$ to mean that $\mathrm{tp}(\bar{a}/BC)$ does not fork over B. We extend this notation to sets by writing $A \mathop{\smile\!\!\!\!\!\!\,|}_B C$ to mean $\bar{a} \mathop{\smile\!\!\!\!\!\!\,|}_B C$ for every finite tuple $\bar{a} \in A$.

Proposition 3.42 (Properties of $\smile\!\!\!\!\!\!\,|$/forking)**.** *Let T be a stable theory.*

1. *(Existence) Let $p(x) \in S_x(A)$ and $A \subseteq B$. Then there exists $q(x) \in S_x(B)$, with $q(x) \supseteq p(x)$ and $q(x)$ does not fork over A.*

2. *(Transitivity) Let $A \subseteq B \subseteq C$ and $p(x) \in S_x(A)$, $q(x) \in S_x(B)$, $r(x) \in S_x(C)$, with $p(x) \subseteq q(x) \subseteq r(x)$. Then $r(x)$ does not fork over A iff $r(x)$ does not fork over B and $q(x)$ does not fork over A.*

3. *(Symmetry) Given $\bar{a}$, and $A \subseteq B$, $\mathrm{tp}(\bar{a}/B)$ does not fork over A iff $\mathrm{tp}(\bar{b}/A\bar{a})$ does not fork over A for all finite $\bar{b}$ from B.*

4. *(Local Character) For all $q(x) \in S_X(B)$, there is $A \subseteq B$, $|A| \leq |T|$ such that $q(x)$ does not fork over A.*

5. *If $q(x) \in S_x(B)$ is algebraic (i.e. has only finitely many realizations) and $q(x)$ does not fork over A, then $q \restriction_A$ is algebraic.*

6. *(Uniqueness) If $A = \mathrm{acl}^{eq}(A)$ and $p(x) \in S_x(A)$, then for all $B \supseteq A$ there is unique $q(x) \supseteq p(x)$ that does not fork over A.*

7. *(Conjugacy and the finite equivalence relation theorem) If $p(x) \in S_x(A)$ and $B \supseteq A$, then p has at most $2^{|T|}$ many non-forking extensions $q(x) \in S_x(B)$. Moreover, if $B = M \models T$ is $|T| + |A|$-saturated and strongly homogeneous, then $\mathrm{Aut}(M/A)$ acts transitively on the set of non-forking extensions $q(x) \in S_x(M)$, and if q_1 and q_2 are distinct such non-forking extensions of p, there is an A-definable finite equivalence relation $E(x_1, x_2)$ such that $q_1(x_1) \cup q_2(x_2) \vdash \neg E(x_1, x_2)$.*

We can restate these properties using the anchor notation as follows:

1. *(Existence) For all $A \subset B$ and tuple e, there a tuple e' satisfying both $\mathrm{tp}(e/A) = \mathrm{tp}(e'/A)$ and $e \mathop{\smile\!\!\!\!\!\!\mid}_A B$.*

2. *(Transitivity) For all A, B, C, D, we have $A \mathop{\smile\!\!\!\!\!\!\mid}_B CD$ if and only if $A \mathop{\smile\!\!\!\!\!\!\mid}_B C$ and $A \mathop{\smile\!\!\!\!\!\!\mid}_{BC} D$.*

3. *(Symmetry) $A \mathop{\smile\!\!\!\!\!\!\mid}_C B$ if and only if $B \mathop{\smile\!\!\!\!\!\!\mid}_C A$.*

4. *(Local Character) For all finite $\bar{a}$, if B is any set, there is $C \subseteq B$ with $|C| \leq |T|$ such that $\bar{a} \mathop{\smile\!\!\!\!\!\!\mid}_C B$.*

5. *(Uniqueness) If $A = \mathrm{acl}^{eq}(A)$, $\mathrm{tp}(a_1/A) = \mathrm{tp}(a_2/A)$, $a_1 \mathop{\smile\!\!\!\!\!\!\mid}_A B$, and $a_2 \mathop{\smile\!\!\!\!\!\!\mid}_A B$, then $\mathrm{tp}(a_1/B) = \mathrm{tp}(a_2/B)$.*

Proof. The proof follows from the local case. $\square$

Remark 3.43. 1. The properties of $\mathop{\smile\!\!\!\!\!\!\mid}$, other than uniqueness and conjugacy, characterize a broader class of theories called simple theories. In fact, the local character property can be taken as a definition of simple theories. Examples of simple theories include the theory of the random graph, the theory of pseudofinite fields, and the model companion to the theory of algebraically closed fields with an automorphism (ACFA).

2. Uniqueness of forking is characteristic of stable theories (it fails for simple theories). We may express uniqueness as follows: let $A \subseteq M$ and let $\mathcal{F}_A$ be the set of formulas $\varphi(x)$ over M that fork over A. Then $\mathcal{F}_A$ is a proper ideal in the boolean algebra of formulas over M (if $\varphi(x) \wedge \psi(x) \in \mathcal{F}_A$, then $\varphi(x) \in \mathcal{F}_A$ or $\psi(x) \in \mathcal{F}_A$). Furthermore, if $A = \mathrm{acl}^{eq}(A)$ and $p(x) \in S_x(A)$, then there is a unique $q(x) \in S_x(B)$ extending $p(x)$ and avoiding $\mathcal{F}_A$,

3. In many natural examples of stable theories, $\mathop{\smile\!\!\!\!\!\!\mid}$ has a natural interpretation. For example, in ACF, $\bar{a} \mathop{\smile\!\!\!\!\!\!\mid}_k \bar{b}$ iff $\mathrm{tr.deg.}(k(\bar{a})/k) = \mathrm{tr.deg.}(k(\bar{a}, \bar{b})/k(\bar{b}))$.

There is also a more geometrical interpretation: let k be any field (not necessarily algebraically closed) and let V be an irreducible k-variety. There is a complete ACF type $p_V(\bar{x})$, which says that $\bar{x} \in U$ for every Zariski open $U \subseteq V$ defined over k. Then for any field $F \supseteq k$, then $p_V(\bar{x})$ has a unique non-forking extension to F if and only if V remains irreducible as a variety over F. Note then that $p_V(\bar{x})$ has a unique non-forking extension to k^{alg} if and only if V is absolutely irreducible.

4. Let T be stable, $\overline{M} \models T$, and $A \subseteq \overline{M}$ a small set. We present some examples of forking calculus, which will highlight the usefulness of the anchor notation.

   (i) Suppose that $(a_\alpha : \alpha < \kappa)$ are tuples in $\overline{M}$ such that $a_\alpha \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A \{a_\beta : \beta < \alpha\}$ for all $\alpha < \kappa$. Then $a_\alpha \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A \{a_\beta : \alpha \neq \beta < \kappa\}$ and we say that $\{a_\alpha : \alpha < \kappa\}$ is an A-independent set.

   Proof. By local character, it is enough to prove, for any $\alpha < \kappa$, any n and $\beta_1, \cdots, \beta_n < \kappa$, that $a_\alpha \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A a_{\beta_1} \cdots a_{\beta_n}$. We will do so by induction on n. If $n = 1$, it is an immediate consequence of the assumption.
   Now assume $n > 1$. If $\alpha > \beta_n$ for all n, this is the assumption. Else, we can assume, by reordering the β_i, that if $i < j$ then $\beta_i < \beta_j$. By assumption, we get $a_{\beta_n} \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A a_\alpha a_{\beta_1}, \cdots a_{\beta_{n-1}}$. Applying transitivity, this yields $a_{\beta_n} \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{A a_{\beta_1} \cdots a_{\beta_{n-1}}} a_\alpha$, and by the induction hypothesis (and transitivity) we get $a_{\beta_n} \cdots a_{\beta_1} \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A a_\alpha$, what we wanted to prove. □

   (ii) (Weight) Suppose that $\{a_\alpha : \alpha < |T|^+\}$ is an A-independent set of finite tuples. Then for any b, there is α such that $b \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A a_\alpha$.

   Proof. For a contradiction, suppose that there is b such that $b \not\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A a_\alpha$ for all $\alpha < |T|^+$. We claim that for all $\alpha < |T|^+$, $b \not\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{A \cup \{a_\beta : \beta < \alpha\}} a_\alpha$. To see this, fix some $\alpha < |T|^+$. By assumption, $a_\alpha \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A \{a_\beta : \beta < \alpha\}$. As $a_\alpha \not\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_A b$, it follows from transitivity that $a_\alpha \not\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{A \cup \{a_\beta : \beta < \alpha\}} b$. By symmetry, we have that $b \not\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{A \cup \{a_\beta : \beta < \alpha\}} a_\alpha$, as required.
   Now, by local character, there is $B_0 \subseteq \{a_\alpha : \alpha < |T|^+\}$ such that $b \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{AB_0} \{a_\alpha : \alpha < |T|^+\}$ and $|B_0| \leq |T|$. As $|T|^+$ is regular, there is $\gamma < |T|^+$ such that $B_0 \subseteq \{a_\alpha : \alpha < \gamma\}$. But then $b \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{AB_0} a_\gamma$, contradicting the fact that for all $\alpha < |T|^+$, $b \not\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_{A \cup \{a_\beta : \beta < \alpha\}} a_\alpha$. □

   Denote $p = \mathrm{tp}(b/A)$. By what we just proved, the size of a cardinal α such that there is and A-independent set of size α, is bounded by $|T|^+$. Hence, we can take the supremum of such cardinals, it is called the preweight of the type p.
   From this, we define the weight of p as the supremum of the weights of non forking extensions of p, it is also a well-defined cardinal.

3.3 A Survey of Classification Theory

Classification theory, in the sense of Shelah, is concerned with finding meaningful dividing lines among complete first-order theories. In the past, decidability/undecidability was considered a fundamental dividing line, but it is no longer the case, as dividing lines of a different nature became prominent. However, decidability is still a meaningful proprerty, and can sometimes be deduced from structural properties such as quantifier elimination.

In modern model theory, there are many different crosscutting notions of complexity, some forming dividing lines.

One such measure of complexity is called the spectrum of a theory T, denoted by $I(\kappa, T)$. Assume that T is a complete, countable theory with no finite models. Then $I(\kappa, T)$ is defined to be the cardinality of the set of isomorphism classes of models of cardinality κ. In 1965, Morley proved his well-known categoricity theorem, that is, if $I(\kappa, T) = 1$ for some uncountable κ, then $I(\kappa, T) = 1$ for all uncountable κ.

In the case of $\kappa = \aleph_0$, much less is known. It is a famous conjecture by Vaught that either $I(\aleph_0, T) \leq \aleph_0$ or $I(\aleph_0, T) = 2^{\aleph_0}$ (without assuming the continuum hypothesis, of course).

In the early eighties, Shelah proved his "main gap" conjecture. To state it, we need to recall some set theoretic notation:

Definition 3.44. We define, for any ordinal α, the cardinal $\beth_\alpha$ by induction :

- $\beth_0 = \aleph_0$
- if $\alpha = \beta + 1$, then $\beth_\alpha = 2^{\beth_\beta}$
- if α is limit, then $\beth_\alpha = \sup\{\beth_\beta, \beta < \alpha\}$

The main gap theorem then states that for all ordinal $\alpha > 0$, either $I(\aleph_\alpha, T) = 2^\kappa$ for all $\alpha \geq 1$, or $I(\aleph_\alpha, T) < \beth_{\omega_1}(|\omega + \alpha|)$. The idea is that either a theory T has a maximal number of models, or there is some kind of classification of the models of T. One reason for studying stable theories is the following theorem due to Shelah:

Theorem 3.45. *For all unstable T and all uncountable κ, $I(\kappa, T) = 2^\kappa$.*

So, from the the point of view of $I(\kappa, T)$ as a measure of complexity, we may assume that T is stable (note that when $\kappa = \aleph_0$, there is no analogous result; the theory of dense linear orders is unstable, but is $\aleph_0$-categorical).

Definition 3.46. Let T be a single-sorted theory. We say T is strongly minimal if, for any formula $\varphi(x)$ in one variable and any model M of T, $\varphi(M)$ is either finite or cofinite.

Note that every strongly minimal theory is stable. Suppose T is strongly minimal, and that M is a countable model. Then there are only countably many 1-types: the algebraic types, i.e. those which contain (and are isolated by) a formula with only finitely many realizations, and the unique (unrealized) type $p(x)$ expressing that x is in every cofinite set.

Example 3.47. The following examples can be show to be strongly minimal via quantifier-elimination (note that quantifier-elimination does not imply strong minimality in general, it just makes it easier to study the formulas in one variable):

(i) the theory of an infinite set in the language of equality,

(ii) the theory of an infinite vector space over a countable field F in the language of modules (i.e. the language of groups together with a function symbol λ_r for scalar multiplication by r for all $r \in F$),

(iii) algebraically closed fields,

(iv) $\mathrm{Th}(\mathbb{Z}_{p^\infty})$, the theory of the Prüfer p-group.

Fact 3.48. *Suppose T is strongly minimal, $M \models T$, and $A \subseteq M$. Then the algebraic closure over A in M is a pregeometry. i.e. for $a, b \in M$ if $b \in \mathrm{acl}(A, a) \setminus \mathrm{acl}(A)$, then $a \in \mathrm{acl}(A, b)$.*

Definition 3.49. A tuple $b_1, \cdots, b_n$ is said to be A-algebraically independent if for any i, we have $a_i \in \mathrm{acl}(A \cup (\{a_1, \cdots a_n\} \setminus \{a_i\}))$ if and only if $a_i \in \mathrm{acl}(A)$.

Note that by properties of algebraic closure, if $b_i \in \mathrm{acl}(A, \overline{a})$ for all $i = 1, \ldots, n$ and $c \in \mathrm{acl}(b_1, \ldots, b_n, \overline{a}, A)$, then $c \in \mathrm{acl}(A, \overline{a})$. This, and the previous fact, allow us to define, for $\bar{b}$ a finite tuple from $\overline{M}$, the dimension of $\bar{b}$ over A, denoted $\dim(\bar{b}/A)$, as the maximal size of an A-algebraically independent subtuple of $\bar{b}$.

With a bit of forking calculus, the enthusiasmic reader can prove that if $A \subset B$, then $\bar{b} \mathop{\smile\!\!\!\!\!|}_A B \Leftrightarrow \dim(\bar{b}/B) = \dim(\bar{b}/A)$.

Fact 3.50. *Let T be strongly minimal. Then*

$$I(\kappa, T) = \begin{cases} 1, & \textit{if } \kappa > \omega, \\ 0 \textit{ or } \aleph_0, & \textit{if } \kappa = \aleph_0. \end{cases}$$

Proof. (For when $\kappa > \omega$) Let $M \models T$ and $I \subseteq M$ be a maximal algebraically independent set over $\emptyset$. Then $|I|$ is well-defined and $M = \mathrm{acl}(I)$. Moreover, the type $\mathrm{tp}(I/\emptyset)$ is determined by $|I|$. If $|M| = \kappa > \omega$, then $|I| = \kappa$. If $|M| = \aleph_0$ however, there is a bit more work to do, see [Mar06] for some details. $\square$

Definition 3.51. T is *ω-stable* if for some countable $M \models T$ and finite tuple of variables x, $|S_x(M)| = \aleph_0$.

Note that strongly-minimal implies ω-stable, which in turn implies stable.

Fact 3.52. *T is ω-stable if and only for any finite tuple of variable x and $X = S_x(M)$, every $p \in X$ has ordinal valued Cantor-Bendxson rank. In that case, the Cantor-Bendixson rank is better know as $\mathrm{RM}(p)$, the Morley rank.*

We can, as was done in Definition 3.35, define the Morley rank of a formula.

Note. If T strongly minimal then $RM(x = x) = 1$. There are only finitely many $p \in S_x(\overline{M})$ which containing $\Phi(x)$ with have maximal Cantor-Bendixson rank. This finite number is called the Morley degree $\mathrm{dM}(\Phi(x))$ of Φ.

A lot of natural and interesting theories are ω-stable, let us give a few examples.

Example 3.53. The theory DCF_0 of differentially closed fields of characteristic zero, is ω-stable (see 3.68 for a definition). Moreover, the formula $x = x$ has Morley rank ω.

There is a "structure theory" for models of T where T is ω-stable even though there may be 2^κ models of cardinality κ for all $\kappa > \omega$.

Example 3.54. The theory $T = \mathrm{Th}\left(\mathbb{Z}_{p^\infty}^{(\omega)}, +\right)$ is also ω-stable of Morley rank ω. Moreover, the models of T are precisely given by $\mathbb{Z}_{p^\omega}^{(\kappa)} \oplus \mathbb{Q}^{(\lambda)}$, where $\kappa > \omega$ and $\lambda > 0$.

A useful property of ω-stable theories is that there are prime models over all sets, i.e. for any $M \models T$ and $A \subseteq M$, the theory $\mathrm{Th}(M, a)_{a \in A}$ has a unique prime model.

The continuum hypothesis skeptical reader will be interested to learn that Vaught's Conjecture is known for ω-stable theories T.

Another class of stable theories, not as easily tamed as ω-stable ones, but no completely wild either, are given by superstable theories.

Definition 3.55. T is *superstable* if T is stable and for every $p(x) \in S_x(B)$ there is a finite set $A \subseteq B$ such that $p(x)$ does not fork over A.

Fact 3.56. *1. T is ω-stable if and only if*

- *T is superstable*
- *$S(T)$ $(= \bigcup_x S_x(T))$ is countable*
- *every $p(x) \in S(B)$ has finite multiplicity (i.e. has only finitely many non-forking extensions to a given $M \supseteq B$)*

2. If T is not superstable, then $I(\kappa, T) = 2^\kappa$ for all $\kappa > \omega$.

3. Vaught's Conjecture is still open for superstable theories.

Example 3.57. $\mathrm{Th}(\mathbb{Z}, +, 0)$ is superstable.

Exercise. If $M = (\mathbb{Z}, +, 0)$, prove that $|S_1(M)| = 2^{\aleph_0}$.

Solution. Since the language is countable, $S_1(M) \leq 2^{\aleph_0}$. For each $i \in \omega$, define

$$\begin{aligned} \varphi_i(x) &:= \text{"}x \text{ is a multiple of the } i^{\text{th}} \text{ prime"} \\ &= (\exists y)\ [\underbrace{y + \ldots + y}_{i^{\text{th}} \text{ prime times}} = x]. \end{aligned}$$

Then for each $\sigma \in 2^{<\omega}$, define $p_\sigma(x) := \{\varphi_i(x) : \sigma(i) = 1\} \cup \{\neg\varphi_i(x) : \sigma(i) = 0\}$. This definition can be extended in the natural way to define $p_f(x)$ for each $f \in 2^\omega$. Now each $p_\sigma(x)$ is satisfiable, thus by Compactness theorem, each $p_f(x)$ is also satisfiable. Therefore, by taking a completion of each $p_f(x)$, there are at least $2^{\aleph_0}$ many 1-types given by $\{p_f(x) : f \in 2^\omega\}$.

Even though this theory is not ω-stable, there is a description of the models. Let $\hat{\mathbb{Z}}$ be the profinite completion of $\mathbb{Z}$, i.e. $\hat{\mathbb{Z}} = \lim_{\overleftarrow{n}} \mathbb{Z}/n\mathbb{Z}$.
Then the $2^{\aleph_0}$-saturated models of T are precisely of the form

$$\hat{\mathbb{Z}} \oplus \mathbb{Q}^{[\delta]}, \quad \delta \geq 2^{\aleph_0}.$$

In fact the models of T are precisely of the form

$$G \oplus \mathbb{Q}^{[\delta]}, \quad \delta \geq 2^{\aleph_0}$$

for some $G < \hat{\mathbb{Z}}$.

Note. Let G be a saturated model of T. Then $G/G^0 = \hat{\mathbb{Z}}$, i.e. $G \Big/ \bigcap_n nG = \hat{\mathbb{Z}}$.

The last, and wildest, region in the land of stable theories, is composed of stable, not superstable theories.

Example 3.58. The following theories are stable but not superstable.

- $\mathrm{Th}(\mathbb{Z}, +)^{(\omega)}$
- $\mathrm{Th}(\mathcal{F}_p(t)^{\mathrm{Sep}})$
- $\mathrm{Th}(F_2, \cdot)$, where F_2 is the free group on 2 generators.

Remark 3.59. The two first statements are easy to prove, but the last one is not. It was proven by Sela, and opened a whole new area of exploration for applied model theory.

If fact, to prove the first two facts, one can use the following general statement: if T is a theory, G a group definable in a model M of T, and there is a sequence $G = G_0 \geq G_1 \geq G_2 \geq \cdots$ of definable subgroups where G_{i+1} has infinite index in G_i for all i, then T is not superstable.

Many theories that are interesting from a model theoretic perspective are actually unstable. For example, the theory of the reals $\mathrm{Th}(\mathbb{R}, +, \times)$ and the theory of the p-adics $\mathrm{Th}(\mathbb{Q}_{p,}, \times)$ are unstable.

These could be the subject of another course, but here, we will restrict our attention to stable theories. Therefore, let us fix (again) a stable theory T.

Definition 3.60. $p(x) \in S(A)$ is stationary if for all $B \subseteq A$ there is a unique nonforking extension $q(x) \in S(B)$ of p. For example, if $A = \mathrm{acl}^{\mathrm{eq}}(A)$, then any $p(x) \in S(A)$ is stationary.

Lemma 3.61. *Let $p(x) \in S(A)$ be stationary. Then there exists a unique smallest $A_0 \subseteq \mathrm{dcl}^{eq}(A)$ such that $p(x)$ does not fork over A_0, and $p \upharpoonright A_0$ is stationary. This set A_0 is denoted $\mathrm{Cb}(p)$, and called the canonical base of p.*

Proof. (Sketch) Let $M \supseteq A$, and let $q(x)$ be the nonforking extension of $p(x)$ to M. Then $q(x)$ is definable by stability. For each $\varphi(x, y) \in L$, let $\Psi_\varphi(y)$ over M be the φ-definition of p. Then $A_0 = \mathrm{dcl}\{\text{codes of } \Psi_\varphi(y) : \varphi \in L\}$. $\square$

3.4 Geometric Stability Theory

In geometric stability theory, we study the "complexity" of models in terms of dimension-like quantities that can be axiomatized in terms of combinatorial geometry. As an example, we explore questions such as: How does a family of subsets of $\mathbb{C} \times \mathbb{C}$

$$\{y = ax + b : a, b \in \mathbb{C}\} \subseteq \mathbb{C}^2$$

intersect with each other? Note that this notion of complexity has nothing to do with the complexity of computer science. When exploring the complexity of strongly minimal sets, it is convenient to discuss Zilber's conjecture, which we shall give below.

Let T be a stable theory, and X a definable set in a saturated model $\overline{M}$ of T. We call X strongly minimal (with respect to the ambient T or $\overline{M}$) if X is infinite and every definable (with parameters in $\bar{M}$) subset of X is finite or cofinite. Equivalently, $\mathrm{RM}(X) = \mathrm{dM}(X) = 1$.

Given any definable (without parameters) set X in $\overline{M}$, by X^{eq}, we mean the collection of all sorts S_E of T^{eq} (or $\overline{M}^{\mathrm{eq}}$) where E is a $\emptyset$-definable equivalence relation on $X \times \ldots \times X$ (n times for some n).

Zilber suggested that we could classify or describe strongly minimal sets, and that moreover, that any "rich enough" strongly minimal set arises from an algebraic object :

Conjecture. *(Zilber) Let X be a strongly minimal set (in $\overline{M} \models T$) definable without parameters. Then exactly one of the following holds:*

1. *The algebraic closure in X is "trivial" i.e. for $a_1, \ldots, a_n \in X$,*

$$\mathrm{acl}_{\overline{M}}(a_1, \ldots, a_n) \cap X = \bigcup_{i=1}^{n} \left(\mathrm{acl}(a_i) \cap X\right).$$

2. *X is biinterpretable with a strongly minimal group G, and the group G satisfies that any definable (with parameters) subset of $G \times \cdots \times G$ is a boolean combination of cosets of definable subgroups of $G \times \cdots \times G$*

3. *There is a strongly minimal field $(K, +, \times)$ definable in X^{eq}.*

These three alternatives measure the complexity of the strongly minimal set, from simplest to most complex. Unfortunatly, the Zilber conjecture was proven to be false. Nonetheless, it has been extremely influencial, and we will expose here some of the most important ideas it spawned.

First, let us define :

Definition 3.62. Let X be any definable set, it is said to be one based if for any tuple $\bar{a}$ from X, and any $B = \mathrm{acl}^{eq}(B)$, we have $\mathrm{Cb}(\mathrm{tp}(\bar{a}/B)) \in \mathrm{acl}(\bar{a})$.

Note the following proposition from general stability theory :

Proposition 3.63. *Let $p \in S(B)$ be a stationary type, and $(a_i)_{i\in\mathbb{N}}$ be a Morley sequence in p. Then $\mathrm{Cb}(p) \in \mathrm{acl}(\{a_i, i \in \mathbb{N}\})$.*

Hence, a set is one based if and only if only one realization of the type is needed to find the canonical base. From a geometric perspective, this is equivalent to the non-existence of "rich families of curves". For example, if one consider an algebraically closed field of characteristic zero, the formula $y = ax+b$ describes a rich family of curves as a and b vary, meaning that we need two points on the line $y = ax + b$ to find a and b back. So this structure is not one-based.

One can easily show that $\mathbb{Q}$ vector spaces are an example of a one-based structure.

This definition is linked with Zilber's conjecture by the following :

Theorem 3.64. *Let X be strongly minimal. Then X is one based if and only if case 1. or 2. hold.*

This is a non-trivial theorem, one of the main achievement of geometric stability theory.

As we mentioned, the trichotomy is false. In fact, it fails in a very strong sense :

Theorem 3.65 (Hrushovski)**.** *There is a strongly minimal set which is not one-based, but does not interpret any infinite group.*

Note that this indeed disproves the conjecture by Theorem 3.64, as such a set would have to satisfy alternative 3. of Zilber's trichotomy. But if it did, it would interpret an algebraically closed field, and this structure does not even interpret an infinite group !

To construct this strongly minimal set, a variant of Fraissé amalgamation was used, called Hrushovski construction. This proved to be a very useful tool to construct new structures. The curious reader is invited to consult [Wag10], for example, to learn more about these.

Even if false, the trichotomy still holds in a number of natural theories. By that we mean that if T is a theory, it satisfies Zilber's trichotomy if and only if strongly minimal sets definable in models of T satisfy it.

One of the most important example is DCF_0, the theory of differentially closed fields of characteristic zero, which we will now define.

Let $\mathcal{L}$ be the language of rings, with one additional function symbol ∂. A differential ring R is an $\mathcal{L}$-structure satisfying the ring axioms, and such that ∂ is a derivation, i.e. for all $a, b \in R$, we have $\partial(a + b) = \partial(a) + \partial(b)$ and $\partial(ab) = \partial(a)b + a\partial(b)$.

The theory of differential fields of characteristic zero has a model companion, which is the theory of differentially closed fields, denoted DCF_0. To describe its axioms, we will need to define differential polynomials.

Definition 3.66. Let (K, ∂_0) be a differential field. The ring $K\{x\}$ of differential polynomials over K is defined as the polynomial ring $K[\{\partial^i(x), i \in \mathbb{N}\}]$, with the differential ring structure given by $\partial = \partial_0$ on K, and $\partial(\partial^i(x)) = \partial^{i+1}(x)$.

For differential polynomials, we can define :

Definition 3.67. Let f be a differential polynomial. Then the order of f is the biggest i such that $\partial^i(x)$ appears in f.

We can now explicit the theory DCF_0 :

Definition 3.68. The theory DCF_0 is axiomatized by the following:

- axioms for ACF_0
- for any differential polynomials f, g, with the order of g strictly larger than the order of f, there is a such that $g(a) = 0$ and $f(a) \neq 0$

Equipped with these axioms, one can prove the following basic properties of DCF_0 :

- quantifier elimination
- elimination of imaginaries
- ω-stable of Morley rank ω

In particular, any definable set will be given as solution of differential equations and inequations.

Moreover, Zilber's trichotomy is true in DCF_0. This was first proven by Hrushovski and Zilber, using the abstract machinery of Zariski geometry, which we do not have time to describe here (see [Bou09] for an explanation). Later, Pillay and Ziegler, in [PZ03], obtained a second proof, as a corollary of proving a strong structural property of differentially closed fields, called the canonical base property.

Let us now fix a monster model $\mathbb{M}$ of DCF_0. Let $\mathcal{C} = \{x \in \mathbb{M}, \partial(x) = 0\}$ be the field of constants of $\mathbb{M}$. This is an algebraically closed field, and moreover, any definable subset (possibly with parameters) of $\mathcal{C}$ is already definable in $(\mathcal{C}, +, \times)$, with parameters from $\mathcal{C}$. We say that $\mathcal{C}$ is a purely stably embedded algebraically closed field.

Let us now give a more explicit version of the trichotomy in DCF_0. In case one, nothing more can be said about the strongly minimal set X. However, identifying exactly what differential equations give rises to trivial minimal sets is an ongoing project of the model theory community.

In the third case, the algebraically closed field definable in the set X can be proven to be definably isomorphic to $\mathcal{C}$.

The second case is the most complex. To expose it, we need to define some basic notions from algebraic geometry. For a more detailled introduction to this subject, we refer the reader to [SR94]. Let K be an algebraically closed field.

Definition 3.69. An affine algebraic variety is a subset of K^n, for some n, defined as the zero locus of some finite system of polynomial equations.

Definition 3.70. The projective n-space, denoted by $\mathbb{P}^n(K)$, is defined as $K^n \setminus \{0\}$, quotiented by the equivalence relation $\bar{a} \sim \bar{b}$ if and only if there is $\lambda \in K \setminus \{0\}$ such that $\bar{a} = \lambda \cdot \bar{b}$.

A projective algebraic variety is a subset of $\mathbb{P}^n(K)$ defined as the zero locus of a finite system of homogeneous polynomials. Note that this zero set is well defined because if P is an homogeneous polynomial, then there exist $k \in \mathbb{N}$ such that for all $\bar{a}, \lambda$, we have $P(\lambda \cdot \bar{a}) = \lambda^k \cdot P(\bar{a})$.

Morphisms of affine algebraic varieties are given by polynomial maps. Morphisms of projective algebraic varieties are also given by polynomial, under the condition that they are well defined on the projective space.

Definition 3.71. An abelian variety is an irreducible projective variety G equipped with a group operation $m : G \times G \to G$ which is a morphism of projective variety.

Definition 3.72. An abelian variety is simple if it has not proper nontrivial abelian subvariety.

It can be showed that all abelian variety are commutative groups, justifying the terminology. The following will be of importance :

Definition 3.73. Let A be an algebraic variety defined over K, and $k \subset K$ another algebraically closed field. We say that A descends to k is there is an abelian variety A_0, defined over k, isomorphic to A.

Let A be an abelian variety defined in $(\mathbb{M}, +, \times)$, its Morley rank can be computed as $\mathrm{RM}(A) = \omega \cdot \dim(A)$, where $\dim(A)$ is the dimension of A as a projective algebraic variety. Note that A can indeed be viewed as a definable set, by elimination of imaginaries.

Fact. *A has a unique smallest Zariski dense definable in $(\mathbb{M}, +, \times, \partial)$ subgroup $A^{\#}$. If A is simple, then $A^{\#}$ is strongly minimal.*

We can now states the second alternative of the trichotomy. If X interprets a group satisfying the properties in case 2, then there exist a simple abelian variety A, which does not descend to $\mathcal{C}$, and G can be taken equal to $A^{\#}$.

One of the most celebrated application of model theory was to use this to prove function field Mordell-Lang in characteristic zero :

Theorem 3.74. *Let $k \subset K$ be algebraically closed subfields of characteristic zero. Let A be an abelian variety defined over K, with k trace zero, i.e. no abelian subvariety of A descends to k.*

Let Γ be a finitely generated subgroup of A, and let X be an algebraic subvariety of A. Assume $X \cap \Gamma$ is Zariski dense in X. Then X is a coset of a subgroup $N \subset A$.

3.5 Keisler Measures and Combinatorics

We shall discuss the notion of Keisler measures on definable sets and their applications to combinatorics. References for this section include [HP11], [Pil18], and [Kei87]. We still use the conventions that T is a complete theory and $\overline{M}$ is κ-saturated, strongly κ-homogeneous for a sufficiently large κ, however we no longer assume T is stable. For convenience, we will occasionally identify a formula $\varphi(\overline{x})$ with the set in $\overline{M}$ it defines, i.e. $\varphi(\overline{M}) = \{\overline{m} \in \overline{M} : \varphi(\overline{m}) \text{ holds}\}$.

Definition 3.75. A Keisler measure in $\overline{x}$ (a collection of free variables/sorts) over a set of parameters A is a finitely additive probability measure $\mu(\overline{x})$ on the Boolean algebra of sets defined by A-formulas in variables $\overline{x}$, i.e. if $\varphi(\overline{x})$ defines X, then $0 \leq \mu(X) \leq 1$ is a real number. Furthermore, $\mu(\overline{M}^{|\overline{x}|}) = 1$, $\mu(\emptyset) = 0$, and if X, Y are definable and $X \cap Y = \emptyset$, then $\mu(X \cup Y) = \mu(X) + \mu(Y)$. We will feel free to abuse notation and write $\mu(\varphi(\overline{x}))$ to mean $\mu(X)$ where $\varphi(\overline{x})$ defines X in $\overline{M}$ (that is $X = \varphi(\overline{M})$.)

Remark 3.76.

1. A complete type $p(\overline{x}) \in S_{\overline{x}}(A)$ induces the Keisler measure μ_p, where $\mu(\varphi(\overline{x})) = 1$ if $\varphi(\overline{x}) \in p$ and 0 otherwise. This requires a short proof.

   Because types are finitely consistent, we have $x = x \in p$, hence $\mu_p(\overline{M}^{|\overline{x}|}) = \mu_p(x = x) = 1$. Similarly, $\mu_p(\emptyset) = \mu_p(\neg(x = x)) = 0$.

   Since p is a complete and finitely consistent, then if $\varphi(\overline{M}) \cap \psi(\overline{M}) = \emptyset$, at most one of $\psi(\overline{x})$, $\varphi(\overline{x})$ is in p, so $\mu_p(\varphi(\overline{x}) \vee \psi(\overline{x})) = \mu_p(\varphi(\overline{M}) \cup \psi(\overline{M})) = \mu_p(\varphi(\overline{x})) + \mu_p(\psi(\overline{x}))$.

2. Suppose $\Delta(\overline{x})$ is a finite collection $\{\varphi_1(\overline{x}, \overline{y_1}), \ldots, \varphi_r(\overline{x}, \overline{y_r})\}$ of L-formulas. Then we have the Keisler Δ-measures over A, i.e. where the relevant Boolean algebra is the set of formulas $\psi(\overline{x}) \in L_A$ which are equivalent to Boolean combinations of A-instances of $\varphi_i(\overline{x}, \overline{y_i})$ for $1 \leq i \leq r$.

3. We shall usually assume A is a model of T, i.e. $A = M$ is an elementary substructure of $\overline{M}$.

Example 3.77.

1. Complete types induce Keisler measures as in Remark 3.76.

2. Let $T = RCF$. Let $A = M = \mathbb{R}$ be the standard model of T, which is an elementary substructure of $\overline{M}$. Consider μ, the Lebesgue measure on $[0,1]_{\mathbb{R}}$, the standard copy of the unit interval. (Not to be confused with $[0,1]_{\overline{M}}$.)

   Given $\varphi(x)$, make $\varphi'(x)$ be $\varphi() \wedge (0 \leq x \leq 1)$. Then since we have quantifier elimination for real closed ordered fields, φ' is a finite Boolean combination of polynomial equalities and inequalities, so $\varphi'(\overline{M})$ is a finite collection of intervals and points lying in the unit interval.

Then define $\mu'(\varphi(x)) = \mu(\varphi'(\mathbb{R}))$. Since $\mu([0,1]) = 1$ and $\mu(A) \leq \mu(B)$ for $\mathcal{A} \subseteq B$, μ' is $[0,1]$-valued, $\mu'(x = x) = \mu((x = x)') = 1$. Since $\mu(\emptyset) = 0$, $\mu'(\emptyset) = 0$ as well. Finite additivity for μ' follows from σ-additivity of μ.

Remark that the Lebesgue measure can also be used to defined a measure on the sort $\bar{M}^n$, for any n.

3. Let $\{\alpha_i\}_{i\in\omega}$ be such that $\Sigma_{i=0}^{\infty} \alpha_i = 1$. Then for Keisler measures $\{\mu_i(\overline{x})\}_{i\in\omega}$ over A, $\mu(\overline{x}) = \Sigma_{i=0}^{\infty} \alpha_i \mu_i(\overline{x})$ is clearly also a Keisler measure over A. This is called a weighted average of the μ_i's. Of special note is when each of the μ_i's is given by a complete type, then this is called a weighted average of types.

Definition.

- For a topological space X, the Borel σ-algebra of X, denoted $B(X)$, (i.e. closed under countably infinite unions) is the collection of sets generated by the open sets under the actions of complementation and taking countable unions, or equivalently, the smallest σ-algebra containing all the open sets of X.
- A Borel probability measure is a σ-additive measure $\mu : B(X) \to [0,1]$.
- a Borel probability measure μ is regular if for any $B \in B(X)$, we have

$$\mu(B) = \inf\{\mu(U) : B \subseteq U \text{ open}\} = \sup\{\mu(C) : C \subseteq B \text{ closed}\}$$

Note. Note that if X is compact and totally disconnected, i.e. a Stone space, then regularity of a Borel probability measure μ implies that for all closed $B \subseteq X$, $\mu(B) = \inf\{\mu(C) : B \subseteq C \text{ clopen}\}$. That is, μ is determined uniquely by its value on the clopen sets.

Fact 3.78. *A Keisler measure $\mu(\overline{x})$ over M an elementary substructure of $\overline{M}$ can be identified with a regular Borel probability measure on $S_{\overline{x}}(M)$. Similarly, a Keisler Δ-measure $\mu(\overline{x})$ over M can be identified with a regular Borel probability measure on $S_\Delta(M)$.*

More specifically, given a Keisler measure $\mu(\overline{x})$ over m, define μ' to be a Borel probability measure μ' on $S_{\overline{x}}(M)$ via $\mu'([\varphi(\overline{x})]) = \mu(\varphi(\overline{x}))$ for all clopen sets $[\varphi(\overline{x})]$. By the above note, this is enough to uniquely determine the value of μ' on all Borel subsets of $S_{\overline{x}}(M)$. (Similarly for Keisler Δ-measures and $S_\Delta(M)$.)

Given a regular Borel probability measure μ' on $S_{\overline{x}}(M)$, define the Keisler measure μ via $\mu(\varphi(\overline{x})) = \mu'([\varphi(\overline{x}])$.

Stability, Cantor-Bendixson rank and Keisler measures come together in the following:

Lemma 3.79. *(Essentially Keisler) Suppose $\Delta(\overline{x}, \overline{z}, \overline{y_1}, \ldots, \overline{y_r})$ is a finite collection of stable formulas over A. Then any Keisler Δ-measure over M is a weighted average of complete Δ-types.*

Proof. By the above fact, we can identify μ with a regular Borel probability measure on $S_\Delta(M)$. We argue by induction on the maximum rank of elements in the domain of μ. That is, we will prove that for any Stone space of ordinal valued Cantor-Bendixson rank, any Borel probability measure is a weighted average of Dirac measures (i.e. given by types, in the case of a type space). And we will prove this by induction on the maximal Cantor-Bendixson rank of an element (a type, in the case of a type space) of our space.

In the base case, the maximum rank is 0, so all types are isolated, and thus there are only finitely many types, so $\mu(\overline{x})$ can be written as a weighted average of these types. We will now take care of the inductive step.

Note that by Lemma 3.34 generalized to finite sets of formulas Δ, the rank $\mathrm{CB}(S_\Delta(\overline{M}))$ is finite. For any partial Δ-type $\Phi(\overline{x})$ over A, recall that:

$$R_\Delta(\Phi(\overline{x})) := \max\{\mathrm{CB}_{S_\Delta(\overline{M})}(p) : \Phi(\overline{x}) \subseteq p(\overline{x})\}$$

And from Proposition 3.36, we have :

$$R_\Delta(\Phi(\overline{x})) = \max\{R_\Delta(p(\overline{x})) : p \in S_\Delta(M) \wedge \Phi(\overline{x}) \subseteq p(\overline{x})\}$$

and there are finitely many types realizing this maximum.

Let $\mu(\overline{x})$ be a Keisler Δ-measure over M. Let $p_1(\overline{x}), \ldots, p_k(\overline{x})$ be the finitely many complete Δ-types over M with $R_\Delta(p_i(\overline{x})) = \mathrm{CB}(S_\Delta(\overline{M}))$. Without loss of generality, assume $\mu(\{p_i(\overline{x})\}) > 0$ for $1 \leq i \leq r$ for some $0 \leq r \leq k$ and $\mu(\{p_i(\overline{x})\}) = 0$ for all $r < i \leq k$. Define $\alpha_i = \mu(\{p_i(\overline{x})\}) \in (0, 1]$ for $1 \leq i \leq r$. Let $U = S_\Delta(M) \setminus (\{p_1(\overline{x}), \ldots, p_r(\overline{x})\})$. Then by the additivity of μ, we get

$$\mu(U) = 1 - \sum_{i=1}^{r} \alpha_i.$$

Set $\mu(U) = \beta$. If $\beta = 0$, then notice that

$$\mu(\overline{x}) = \sum_{i=1}^{r} \alpha_i p_i(\overline{x})$$

by finite additivity, where here $p_i(\overline{x})$ stands for the Keisler measure induced by the type $p_i(\overline{x})$ as discussed above.

Therefore, we now assume that $\beta > 0$. By regularity of μ and the above note, for any open set V, we have $\mu(V) = \sup\{\mu(C) : C \subseteq V \text{ clopen}\}$. In particular, $\mu(U)$ is the supremum of a set of reals obtained as the measure of clopen subsets, so we can find a countable sequence of clopen subsets $\emptyset = U_0 \subseteq U_1 \subseteq U_2 \subseteq \cdots \subseteq U$ such that $\lim_{n\to\infty} \mu(U_n) = \beta$, because the reals have a countable dense subset. Set $\beta_i = \mu(U_i)$ for all $i \geq 1$.

Each of $U_{i+1} \setminus U_i$ for $i \geq 0$ is a clopen, and hence Δ-definable over M, set with positive μ-measure. Furthermore, since all types realizing the maximum rank are not in U and thus not in $U_{i+1} \setminus U_i$ for any $i \geq 0$, the rank of the maximum element of each is strictly smaller. Therefore we can apply our induction hypothesis on each to get that

$$\mu(\overline{x}) \restriction (U_{i+1} \setminus U_i) = \sum_{j=1}^{\infty} \beta_{i,j} q_{i,j}(\overline{x})$$

where

$$\sum_{j=1}^{\infty} \beta_{i,j} = \mu(U_{i+1} \setminus U_i).$$

Then by σ-additivity of μ as a regular Borel probability measure, we see that

$$\mu(\overline{x}) = \sum_{i=1}^{r} \alpha_i p_i(\overline{x}) + \sum_{i=0}^{\infty} \sum_{j=1}^{\infty} \beta_{i,j} q_{i,j}(\overline{x})$$

Noting that

$$\sum_{i=0}^{\infty} \sum_{j=1}^{\infty} \beta_{i,j} + \sum_{i=1}^{r} \alpha_i = 1$$

we conclude that μ is indeed a weighted average of types. □

Remark. In a stable theory, every Keisler measure is locally a weighted average of complete types. However, this is not true of all Keisler measures: In particular, the Keisler measure from Example 3.77 given from the standard Lebesgue measure has a unique extension, which cannot be true of a Keisler measure locally equal to a weighted average.

We now discussion an application of Keisler measures to combinatorics. We shall work in the setting of finite bipartite graphs (L, R, E) where L and R are sets of vertices and $E \subseteq L \times R$ is the edge relation.

Definition. A finite bipartite graph (L, R, E) is said to be ϵ-regular for $\epsilon > 0$ if for any $A \subseteq L$, $B \subseteq R$ and $|A| \geq \epsilon|L|$, $|B| \geq \epsilon|R|$, we have:

$$|\frac{|E \cap (A \times B)|}{|A \times B|} - \frac{|E \cap |L \times R|}{|L \times R}| < \epsilon$$

That is to say, the difference between the density of edges in $(A, B, E \cap (A \times B))$ and the density of edges in the original graph is less than ϵ.

Theorem 3.80. *(Szemerédi) For all $\epsilon > 0$, there exists $N_\epsilon \in \omega$ such that for any finite bipartite graph (L, R, E) there are partitions $L = L_1 \cup \cdots \cup L_n$, $R = R_1 \cup \cdots \cup R_m$ for $n, m < N_\epsilon$ and a set of exceptions $\Sigma \subseteq \{1, \ldots, n\} \times \{1, \ldots, m\}$ such that $|(\bigcup_{(i,j) \in \Sigma} L_i \times R_j)| \leq \epsilon|L \times R|$ and for every $(i, j) \notin \Sigma$, $(L_i, R_j, E \cap (L_i \times R_j))$ is ϵ-regular.*

Note. One can view this ideologically as having three components:

- The structural component, which is the partition of the graph into substructures of the original.

- The error component, which is the small set of exceptions.
- The pseudo-randomness component, which is the ϵ-regularity asserting that "most" of the graph is regular

In combinatorics, there is a strong tradition of restricting the objects of studies to those having or not having a specific property. Given a finite bipartite graph $G = (L_0, R_0, E_0)$, we shall consider finite bipartite graphs (L, R, E) which omit G, i.e. they do not contain a substructure isomorphic to G. In doing so, we can obtain an improved regularity theorem where the ϵ-regularity is replaced by ϵ-homogeneity.

Definition. A finite bipartite graph (L, R, E) is said to be ϵ-homogenous if either $|(L \times R) \setminus E| \leq \epsilon |L \times R|$ (the graph is "almost" the complete graph) or $|(L \times R) \cap E| \leq \epsilon |L \times R|$. (The graph is "almost" the empty graph.)

Definition. The k-half graph (L, R, E) is the graph where $L = R = \{1, \ldots, k\}$ and $E(i, j)$ if and only if $i \leq j$.

This graph should be familiar to the reader, as the one used to define stability of a formula. And in fact, notice that a graph (L, R, E) omits the k-half graph if and only if its edge relation is k-stable.

Theorem 3.81. *(Stable Regularity Theorem) For all $k \in \mathbb{N}$ and $\epsilon > 0$, there exists $N_{\epsilon,k} \in \mathbb{N}$ such that whenever (L, R, E) is a finite bipartite graph which omits the k-half graph (i.e. is k-stable), then there are partitions $L = L_1 \cup \cdots \cup L_n$, $R = R_1 \cup \cdots \cup R_m$, $n, m < N_{\epsilon,k}$ such that each of $(L_i, R_j, E \cap (L_i \times R_j))$ is ϵ-homogeneous.*

Note. In the ideological viewpoint presented above, this theorem is purely structural.

We shall prove the Stable Regularity Theorem using a statement about infinite graphs and some pseudofinite methods to apply it to the finite setting.

Theorem 3.82. *Let (L, R, E) be a $\emptyset$-definable bipartite graph in the structure M. Assume $E(\overline{x}, \overline{y})$ is stable. Identify E with the formula defining it. Let $\mu(\overline{x})$ be any Keisler measure on L over M. Then for any $1 > \epsilon > 0$ we can find definable partitions $L = L_0 \cup \ldots L_n$, $R = R_1 \cup \ldots R_m$ such that for each $(i, j) \subseteq \{0, \ldots, n\} \times \{1, \ldots, m\}$ either:*

$$\forall \overline{b} \in R_j \;\; \mu(L_i \setminus E(\overline{x}, \overline{b})) \leq \epsilon \mu(L_i)$$

or

$$\forall \overline{b} \in R_j \;\; \mu(L_i \cap E(\overline{x}, \overline{b})) \leq \epsilon \mu(L_i)$$

Moreover, each L_i is defined by an E-formula and each R_j is defined by an E^*-formula.

Proof. Let $\Delta = \{E(\overline{x}, \overline{y})\}$ and let μ_0 be the Keisler Δ-measure over M obtained by restricting μ to Δ-formulas. By Lemma 3.79, there is I an initial segment of ω (i.e. either finite or infinite) such that $\mu_0 = \Sigma_{i \in I} \alpha_i p_i$ where p_i is the measure obtained from a complete Δ-type p_i and $\alpha_i \in (0, 1]$ for all $i \in I$. As before, we identify μ_0 with a regular Borel probability measure on $S_\Delta(M)$.

Note that for all $i \in I$, $\mu(\{p_i\}) = \alpha_i$. Indeed, the measure μ_0 is regular, so $\mu_0(\{p_i\}) = \inf\{\mu_0(C) : p_i \in C \text{ clopen}\}$. But for any finite collection of the p_j's, $p_j \neq p_i$, there is a clopen set containing p_i not containing any of the p_j's, so the $\alpha_j p_j$ factors of the weighted average contribute 0. Therefore $\mu_0(\{p_i\}) \leq \alpha_i + \zeta$ for any $\zeta > 0$. Lastly, we have $\mu_0(\{p_i\}) \geq \alpha_i$ by regularity and the fact that μ_0 is a weighted average of a collection of types containing p_i.

By regularity, for each i there exist a clopen set L_i such that $\mu_0(L_i) < \mu_0(\{p_i\}) + \epsilon\mu(L_i)$ because we can pick L_i with $\mu_0(L_i) < \frac{\alpha_i}{1-\epsilon}$. Then by additivity of μ_0, we get $\mu_0(L_i \setminus \{p_i\}) < \epsilon\mu_0(L_i)$.

Let $B = S_\Delta(M) \setminus \{p_i : i \in I\}$. Then B is Borel as the complement of a countable union of closed sets. Furthermore, since $\mu_0 = \Sigma_{i \in I} \alpha_i p_i$, we have $\mu_0(B) = 0$. Now let $\delta = \frac{\alpha_0}{1-\epsilon} - \mu_0(L_0)$. By regularity of μ_0, there is an open set $B \subseteq U$ such that $\mu_0(U) < \delta$.

We have obtained $\{U\} \cup \{L_i : i \in I\}$, an open cover of $S_\Delta(M)$. By compactness there is a finite subcover, which we can assume to be $\{U, L_0, \dots, L_n\}$ for some n. Furthermore, we can assume these are pairwise disjoint by letting L_i be $L_i \setminus (\bigcup_{j<i} L_{i-1})$ for $1 \leq i \leq n$. Each L_i is clopen, so $L_1 \cup \dots \cup L_n$ is clopen, hence its complement is clopen. Let $L_0' = (L_1 \cup \dots \cup L_n)^c$. Notice that $L_0 \subseteq L_0' \subseteq U \cup L_0$ and $\mu_0(L_0') \leq \mu_0(U) + \mu_0(L_0) < \delta + \mu_0(L_0) = \frac{\alpha_0}{1-\epsilon}$, so $\mu_0(L_0') < \frac{\alpha_0}{1-\epsilon}$ as we had before.

Therefore $L_0' \cup L_1 \cup \dots \cup L_n$ is a definable partition of $L(M)$ with $\mu_0(L_i \setminus \{p_i\}) < \epsilon\mu_0(L_i)$ for all $0 \leq i \leq n$. For each $0 \leq i \leq n$, let $\psi_i(\overline{y})$ be the Δ^*-formula defining $p_i \in S_\Delta(M)$, which exists by Proposition 3.39 generalized to finite sets of formulas.

For each $J \subseteq \{0, \dots, n\}$, let $R_J(\overline{y})$ be the formula $\bigwedge_{i \in J} \psi_i(\overline{y}) \wedge \bigwedge_{i \notin J} \neg\psi_i(\overline{y})$. Then the nonempty R_j's are clearly a partition of $R(M)$ by definition. Fix $0 \leq i \leq n$ and $J \subseteq \{0, \dots, n\}$ with $R_J(M)$ nonempty. Then if $i \in J$, for any $\overline{b} \in R_J(M)$ we have $E(\overline{x}, \overline{b}) \in p_i$. Hence for any $\overline{b} \in R_j(M)$, the set $[E(\overline{x}, \overline{b})]$ is clopen, containing p_i, and $L_i \setminus [E(M, \overline{b})] \subseteq L_i \setminus \{p_i\}$, so we obtain $\mu_0(L_i \setminus [E(\overline{x}, \overline{b})] < \epsilon\mu_0(L_i)$. If $i \notin J$, then we have $E(\overline{x}, \overline{b}) \notin p_i$. Apply the same argument to get $\mu_0(L_i \cap [E(\overline{x}, \overline{b})]) < \epsilon\mu_0(L_i)$. $\square$

Equipped with this theorem, we shall now prove the Stable Regularity Theorem using some pseudofinite methods. To entertain the experienced reader, we will provide, instead of a standard ultrafilter argument, a proof based on non-standard models of set theory.

Proof of Stable Regularity Theorem. Assume not for sake of contradiction. Then for some $1 > \epsilon > 0$, no $N \in \omega$ witnesses the existence of partitions with the desired property. Then for each N, there is a counterexample (L_N, R_N, E_N).

Hence $|L_N|$ and $|R_N|$ must tend to infinity, as for N large enough the statement is trivial for small enough vertex sets.

Let $\mathcal{L} = \{\in, l, r, e\}$ be the language of set theory together with constant symbols l, r, and e. Let M_n be the $\mathcal{L}$-structure $(\mathbb{V}, L_N, R_N, E_N)$, where $\mathbb{V}$ is the set-theoretic universe. Let Σ be the incomplete $\mathcal{L}$-theory consisting of all sentences satisified by cofinitely many M_N's. Then Σ is finitely consistent (if we have some finite subcollection, each of them is satisfied by cofinitely many M_N's, and the intersection of finitely many cofinite sets is cofinite), so it is consistent by compactness.

Let $M^* \models \Sigma$ be $(2^{\aleph_0})^+$-saturated, denote $M^* = (\mathbb{V}^*, L^*, R^*, E^*)$. Then (L^*, R^*, E^*) is k-stable since each M_N is k-stable, and this is witnessed by a formula.

Note that each definable subset A of L^* or R^* is assigned a cardinality $|A|$ in M^*, which is a (possibly non-standard) natural number. Indeed, in a model of set theory, cardinality is a definable function, so it lifts to the model M^*.

Hence, for any definable set A in say L^*, we get a nonstandard rational number $\frac{|A|}{|L^*|}$. Now define $\mu(A) := \mathrm{st}(\frac{|A|}{|L^*|})$, where $\mathrm{st} : \mathbb{R}^* \to \mathbb{R}$ is the standard part map, sending each finite non standard real to the unique standard real infinitesimaly close to it. This can be seen as a non-standard counting measure, and is indeed a Keisler measure, because the counting measure is.

Therefore we can apply the above theorem to (L^*, R^*, E^*) with μ and $\delta = \frac{\epsilon}{2}$, a real number outside of M^*. (If we apply the theorem in M^*, then the partitions might be indexed by nonstandard naturals, which will not work. One can verify, though, that the proof did not require us to use anything about standard models, so we can apply it outside of the model just as well.) We obtain $L^* = L_1^* \cup \cdots \cup L_n^*$ and $R^* = R_1^* \cup \cdots \cup R_m^*$ for standard naturals n, m, all L_i and R_j definable in M^*, and each $(L_i^*, R_j^*, E^* \restriction (L_i^* \times R_j^*))$ is δ-homogeneous with respect to μ. Then by definition of μ, we have that either for all $\bar{b} \in R_j^*$ $|L_i^* \setminus E^*(\bar{x}, \bar{b})| \leq \delta |L_i^*| < \epsilon |L_i^*|$ or for all $\bar{b} \in R_j^*$ $|L_i^* \cap E^*(\bar{x}, \bar{b})| \leq \delta |L_i^*| < \epsilon |L_i^*|$.

Notice that there is some sentence σ in $\mathcal{L}$ that expresses this, i.e. uses l to interpret L, r to interpret R, and uses e to interpret E, the language of set theory allows us to express the existence of such a partition by a first order sentence. Thus $M^* \models \sigma$ for infinitely many M_N's, as otherwise $\neg\sigma \in \Sigma$ as it would be true in cofinitely many M_N's. In particular, let $N > n, m$ for which such a partition exists, it is a contradiction, as (M_N, R_N, E_N) was assumed to be a counterexample to the regularity theorem, i.e. assumed to not possess such a partition. $\square$

References

[Bou09] Elisabeth Bouscaren. *Model Theory and Algebraic Geometry: An introduction to E. Hrushovski's proof of the geometric Mordell-Lang conjecture.* Springer, 2009.

[HP11] Ehud Hrushovski and Anand Pillay. On nip and invariant measures, 2011.

[Kei87] H. Jerome Keisler. Measures and forking, 1987.

[Mar06] David Marker. *Model theory: an introduction*, volume 217. Springer Science & Business Media, 2006.

[Pil96] Anand Pillay. *Geometric Stability Theory*. Oxford logic guides. Clarendon Press, 1996.

[Pil98] Anand Pillay. *Model Theory of algebraically closed fields*. Springer, Berlin, Heidelberg, 1998.

[Pil02] Anand Pillay. Lecture notes - model theory, 2002. https://www3.nd.edu/~apillay/pdf/lecturenotes_modeltheory.pdf.

[Pil03a] Anand Pillay. Lecture notes - applied stability theory, 2003. https://www3.nd.edu/~apillay/pdf/lecturenotes.applied.pdf.

[Pil03b] Anand Pillay. Lecture notes - stability theory, 2003. https://www3.nd.edu/~apillay/pdf/lecturenotes.stability.pdf.

[Pil08] Anand Pillay. *An introduction to stability theory*. Dover Publications, 2008.

[Pil18] Anand Pillay. Domination and regularity, 2018. https://arxiv.org/abs/1806.08806.

[PZ03] Anand Pillay and Martin Ziegler. Jet spaces of varieties over differential and difference fields. *Selecta Mathematica*, 9(4):579–599, 2003.

[SR94] Igor Rostislavovich Shafarevich and Miles Reid. *Basic algebraic geometry*, volume 2. Springer, 1994.

[TZ12] Katrin Tent and Martin Ziegler. *A course in model theory*, volume 40. Cambridge University Press, 2012.

[Wag10] Frank Olaf Wagner. Fields and fusions: Hrushovski constructions and their definable groups. *arXiv preprint arXiv:1008.4217*, 2010.