Math 526 Lecture Notes: Stability and Categoricity

Alex Kruckman

Last updated: April 26, 2022

Contents

Introduction: Categoricity

Once we have understood the distinction between syntax and semantics in logic, i.e., between theories and their models, a natural question arises: How much can the axioms of a theory actually tell us about its models? In the extreme case, we might hope that a theory could tell us essentially everything about a particular structure of interest, in the sense that it pins it down uniquely (up to isomorphism).

For example, every complete ordered field is isomorphic to the field of real numbers $\mathbb{R}$. In a sense, this justifies the common axiomatic approach to real analysis: after we give some construction of $\mathbb{R}$, we can then forget that real numbers are Dedekind cuts or equivalence classes of Cauchy sequences or whatever, working instead from the axioms of complete ordered fields.

Definition 0.1. A theory T is **categorical**[1] if it has exactly one model up to isomorphism.

We should specify what we mean by "theory". This is a class on the model theory of first-order logic, so we will focus on first-order theories. While the theory of complete ordered fields alluded to above is classic example of a categorical theory, it is not first-order. The completeness axiom requires us to quantify over *subsets* of the field, so it is naturally second-order. More expressive logics are interesting (and parallels to some of the theory developed in these notes are possible, e.g. see the literature on abstract elementary classes (AECs)), but they are not within the scope of this course.

Unfortunately (or fortunately, depending on your perspective), the notion of categoricity is rather trivial for first-order theories. If a first-order L-theory T has an infinite model M, then, by the Löwenheim–Skolem theorem, T has models of every infinite cardinality $\kappa \geq |L|$. Models of different cardinalities cannot be isomorphic, so T is not categorical.

Exercise 1. Let M be a structure. Show that $\mathrm{Th}(M)$ is categorical if and only if M is finite.

Hint: If M is infinite, apply Löwenheim–Skolem as suggested above. If M is finite, show that every model of $\mathrm{Th}(M)$ is isomorphic to M. Think about the case of a finite language first, and then try to generalize to an infinite language.

A natural next question is this: Are there first-order theories for which the Löwenheim–Skolem is essentially the only obstruction to categoricity? That is, can a theory T pin down an infinite model M uniquely up to isomorphism once we specify the cardinality of M?

Definition 0.2. Let κ be an infinite cardinal. A theory T is **κ-categorical** if T has exactly one model of cardinality κ up to isomorphism. That is, if $M, N \models T$ with $|M| = |N| = \kappa$, then $M \cong N$.

[1] This terminology is due to Oswald Veblen in 1904, who borrowed the term "categorical" from philosophy. It has nothing to do with category theory, which would not be invented for another 40 years or so.

The answer is yes, κ-categorical theories exist. Here are some examples:

(1) The theory T_∞ of infinite sets in the empty language. A structure in the empty language is just a set, and two sets are isomorphic if and only if they have the same cardinality. So T_∞ is κ-categorical for every infinite cardinal κ.

(2) $\mathrm{VS}_{\mathbb{F}_q}$, the theory of infinite $\mathbb{F}_q$-vector spaces, where $\mathbb{F}_q$ is a finite field. An infinite $\mathbb{F}_q$-vector space has dimension equal to its cardinality, and two vector spaces are isomorphic if and only if they have the same dimension. So $\mathrm{VS}_{\mathbb{F}_q}$ is κ-categorical for every infinite cardinal κ.

(3) $\mathrm{VS}_{\mathbb{Q}}$, the theory of $\mathbb{Q}$-vector spaces. The $\mathbb{Q}$-vector spaces of dimensions 1, 2, 3, $\ldots$, $\aleph_0$ are all countably infinite and non-isomorphic. But an uncountable $\mathbb{Q}$-vector space has dimension equal to its cardinality. So $\mathrm{VS}_{\mathbb{F}_q}$ is κ-categorical for every uncountable cardinal κ.

(4) ACF_0, the theory of algebraically closed fields of characteristic 0. Every field of characteristic 0 has a transcendence degree over $\mathbb{Q}$ (the maximum cardinality of an algebraically independent subset), and two algebraically closed fields of characteristic 0 are isomorphic if and only if they have equal transcendence degrees. The fields of transcendence degree 0, 1, 2, 3, $\ldots$, $\aleph_0$ are all countably infinite and non-isomorphic. But an uncountable field has transcendence degree equal to its cardinality. So ACF_0 is κ-categorical for every uncountable cardinal κ.

(5) DLO, the theory of dense linear orders without endpoints. By a back-and-forth argument, DLO is $\aleph_0$-categorical. But it is not κ-categorical for any uncountable κ.

(6) T_{RG}, the complete theory of the random graph (the Fraïssé limit of the class of finite graphs). By a back-and-forth argument, T_{RG} is $\aleph_0$-categorical. But it is not κ-categorical for any uncountable κ.

(7) RCF, the theory of real closed fields, which is the complete theory of $\mathbb{R}$. This theory is not κ-categorical for any infinite κ.

(8) TA, true arithmetic, the complete theory of $(\mathbb{N}; 0, 1, +, \times)$. This theory is not κ-categorical for any infinite κ.

In 1954, Jerzy Łoś observed that among complete theories in a countable[2] language, he could only find four behaviors with respect to categoricity, and he conjectured that these were the only four.

- **Totally categorical**: κ-categorical for all infinite κ. Examples 1 and 2 above.

[2] In these notes, "countable" means "finite or countably infinite".

- **Uncountably categorical** (but not totally categorical): κ-categorical if and only if $\kappa > \aleph_0$. Examples 3 and 4 above.
- **Countably categorical** (but not totally categorical): κ-categorical if and only if $\kappa = \aleph_0$. Examples 5 and 6 above.
- **Non-categorical**: κ-categorical for no infinite κ. Examples 7 and 8 above.

Łoś's conjecture was proved by Michael Morley in his 1962 PhD thesis.

Theorem 0.3 (Morley's Categoricity Theorem)**.** *Let T be a complete theory in a countable language. If T is κ-categorical for some uncountable cardinal κ, then T is λ-categorical for every uncountable cardinal λ.*

More interesting than the statement of the theorem was the method of proof. It turns out that the hypothesis of κ-categoricity for some uncountable κ allows us to deduce a suprising amount of structural information about models of T. Morley's thesis contains a wealth of original ideas that would influence the development of modern model theory: the Morley rank, totally transcendental (ω-stable) theories, the importance of prime and saturated models, the distinction between indiscernible sequences and sets, etc.

One of the goals of this course is to prove Morley's theorem. I will not follow his original proof, but rather an improved presentation due to John Baldwin and Alistair Lachlan (from Baldwin's 1971 PhD thesis under Lachlan). We will also not take the most direct route to the proof, preferring to take our time situating the material, as much as time permits, in the context of more recent developments.

Morley's theorem is about uncountably categorical theories, e.g., Examples 1–4 above. In each of these examples, uncountable categoricity is explained by the existence of a notion of dimension, which classifies models up to isomorphism. For sets, the "dimension" is just cardinality. For vector spaces, it is ordinary linear dimension. And for algebraically closed fields, it is transcendence degree over the prime field. The Baldwin–Lachlan innovation is to find a way to assign an abstract dimension to any model of an uncountably categorical theory. They did this by finding something called a *strongly minimal set* in any such model, building on William Marsh's 1966 PhD thesis, which introduced strongly minimal sets and the associated notion of dimension. This notion of dimension allowed them to prove that an uncountably categorical theory which is not totally categorical has exactly $\aleph_0$-many models up to isomorphism. This is one of several related statements commonly known as the Baldwin–Lachlan Theorem.

In the 1970s and '80s, the ideas of Morley's theorem and the Baldwin–Lachlan theorem were generalized and extended, most notably by Saharon Shelah. One of Shelah's original motivations was to remove the countable language hypothesis in Morley's theorem, which he succeeded in doing.

Theorem 0.4 (Shelah)**.** *Let T be a complete L-theory. If T is κ-categorical for some uncountable cardinal $\kappa \geq |L|$, then T is λ-categorical for every uncountable cardinal $\lambda > |L|$.*

We will not prove Shelah's theorem in this class: the proof is rather technical. What we will study is *stability theory*, the general and powerful framework he developed to obtain this result and many many many others.

We will deal with countable languages almost exclusively – the exception is that we will sometimes form the language L_A obtained by adjoining to L one constant symbol for each element of a subset $A \subseteq M \models T$, and A may be uncountable.

Where questions of categoricity are concerned, we might as well assume our theory is consistent and complete: If T is incomplete, then it has models M and M' which are not even elementarily equivalent, much less isomorphic.

Finally, by Exercise 1 above, there is not much of interest to say about a complete theory with a finite model. In sum, we adopt the following convention.

Convention. Throughout these notes, L is a countable first-order language and T is a complete L-theory with infinite models.

1 A brief primer on set theory

In this section, I will develop the set theory background necessary for the rest of the notes. Since this is not a course in set theory, I will adopt a someone informal style: rather than starting from the axioms of ZFC, I will assume you know what a set is and prove the results we need on ordinal and cardinal numbers in the style of "ordinary mathematics".

1.1 Comparing cardinality

Before defining "cardinal", we will start by explaining how to compare cardinalities of sets. That is, we do not yet define the notation $|X|$, just the relations $|X| \leq |Y|$ and $|X| = |Y|$.

Definition 1.1. Let X and Y be sets. We define $|X| \leq |Y|$ if there exists an injective function $X \to Y$. We define $|X| = |Y|$ if there exists a bijective function $X \to Y$.

If there is a surjective function $X \to Y$, we can define an injective function $Y \to X$ by mapping each element of Y to one of its preimages[3] in X, and hence $|Y| \leq |X|$. The converse (if $|Y| \leq |X|$, then there is a surjective function $X \to Y$) is true except when $Y = \varnothing$ and $X \neq \varnothing$.

Exercise 2. Show that the relation $|X| \leq |Y|$ is a pre-order on the class of all sets. That is, it is reflexive and transitive. Show that the relation $|X| = |Y|$ is an equivalence relation on the class of all sets. That is, it is reflexive, transitive, and symmetric.

It is not obvious that $|X| = |Y|$ is the equivalence relation induced by the preorder $|X| \leq |Y|$. This is the Cantor–Schröder–Bernstein theorem.

Theorem 1.2. *If $|X| \leq |Y|$ and $|Y| \leq |X|$, then $|X| = |Y|$.*

Proof. Let $g\colon X \to Y$ and $f\colon Y \to X$ be injective functions witnessing $|X| \leq |Y|$ and $|Y| \leq |X|$. Let $X' = \mathrm{im}(g) \subseteq Y$. Then g witnesses that $|X| = |X'|$, so it suffices to prove that $|X'| = |Y|$. Moreover, $f' = g \circ f\colon Y \to Y$ is an injective function such that $\mathrm{im}(f') \subseteq X'$.

Replacing X with X' and f with f', we have reduced to showing that if $X \subseteq Y$ and $f\colon Y \to Y$ is an injective function with $\mathrm{im}(f) \subseteq X$, then $|X| = |Y|$.

We define sets Z_n by induction. Set $Z_0 = Y \setminus X$, and define $Z_{n+1} = f(Z_n)$ for all n. Let $Z = \bigcup_{n\in\mathbb{N}} Z_n$, and define

$$h(y) = \begin{cases} f(y) & \text{if } y \in Z \\ y & \text{if } y \notin Z. \end{cases}$$

I claim that $h\colon Y \to X$ is a bijection, which establishes $|X| = |Y|$.

[3] In general, this argument requires the axiom of choice to pick a preimage for each element of Y.

First we show that $\mathrm{im}(h) \subseteq X$. If $y \in Z$, then $h(y) = f(y) \in \mathrm{im}(f) \subseteq X$. If $y \notin Z$, then $y \notin Z_0$, so $y \notin Y \setminus X$, and hence $h(y) = y \in X$.

For surjectivity, let $x \in X$. If $x \notin Z$, then $x = h(x)$. If $x \in Z$, then $x \in Z_n$ for some n. Note that $n > 0$, since $x \notin Z_0 = Y \setminus X$. Then there is some $y \in Z_{n-1}$ such that $x = f(y)$. Since $y \in Z$, $h(y) = f(y) = x$.

For injectivity, suppose $h(y) = h(y')$. If $y, y' \in Z$, then $f(y) = f(y')$, so $y = y'$ since f is injective. If $y, y' \notin Z$, then $y = h(y) = h(y') = y'$. The remaining case is that $y \in Z$ and $y' \notin Z$. Then $f(y) = h(y) = h(y') = y'$. But if $y \in Z_n$, then $y' = f(y) \in Z_{n+1} \subseteq Z$, contradiction. So this case does not occur. □

Cantor's classic diagonalization argument shows that for every set, there is a set of strictly larger cardinality, namely its powerset: $\mathcal{P}(X) = \{Y \mid Y \subseteq X\}$.

Theorem 1.3 (Cantor)**.** *For all sets X, $|X| < |\mathcal{P}(X)|$.*

Proof. To show $|X| \leq |\mathcal{P}(X)|$, it suffices to produce an injective function $f\colon X \to \mathcal{P}(X)$. One such function is given by $f(x) = \{x\}$.

To show $|X| \neq |\mathcal{P}(X)|$, suppose for contradiction that there is a bijective function $g\colon X \to \mathcal{P}(X)$. Define $Y = \{x \in X \mid x \notin g(x)\} \in \mathcal{P}(X)$. Since g is surjective, there is some $y \in X$ such that $g(y) = Y$. Is $y \in Y$? If $y \in Y$, then by definition of Y, $y \notin g(y) = Y$. But if $y \notin Y = g(y)$, then by definition of Y, $y \in Y$. In either case, we have a contradiction. □

1.2 Ordinals

It will be useful to pick a canonical representative for each equivalence class of the equivalence relation $|X| = |Y|$. To do this, we will introduce the ordinal numbers. Ordinals have the additional advantage of providing a framework for induction and recursion for infinite sets.

The **ordinals** are a linearly ordered number system extending the natural numbers, characterized by the following properties:

- 0 is an ordinal.
- For every ordinal α, there is a next largest ordinal $\alpha + 1$ (this is called a **successor** ordinal).
- For every non-empty set of ordinals S with no greatest element, there is an ordinal $\sup S$ which is the least upper bound of S (this is called a **limit** ordinal).
- The ordinals are the smallest system with these closure properties.

We will take these properties as axiomatic, and avoid the details of constructing the ordinals in ZFC. The last clause essentially asserts that induction can be extended to the ordinals.

Theorem 1.4 (Transfinite induction)**.** *Let P be a property of ordinals. Suppose that:*

(1) Base case: 0 satisfies P.

(2) Successor step: If α satisfies P, then $\alpha + 1$ satisfies P.

(3) Limit step: If γ is a limit ordinal and every $\alpha < \gamma$ satisfies P, then γ satisfies P.

Then every ordinal satisfies P.

Proof. Consider the collection of all ordinals α such that all $\beta \leq \alpha$ satisfy P. This collection contains 0 and is closed under successors and limits, so it contains all the ordinals. $\square$

Exercise 3. Prove the principle of "strong induction" for the ordinals: Suppose that for every ordinal α, if all ordinals $\beta < \alpha$ satisfy P, then α satisfies P. Then every ordinal satisfies P.

Each natural number n is an ordinal (the nth successor of 0). The least limit ordinal is denoted ω. The ordinals continue:

$$\begin{aligned}
&0, 1, 2, \dots, \omega, \omega+1, \omega+2, \dots, \\
&\omega+\omega(=\omega\cdot 2), \omega\cdot 2+1, \omega\cdot 2+2, \dots \\
&\omega\cdot 3, \dots, \omega\cdot 4, \dots, \omega\cdot\omega(=\omega^2), \dots
\end{aligned}$$

Formally, when we construct the ordinals in set theory, an ordinal α is identified with the set of ordinals less than α. So $0 = \varnothing$, $1 = \{0\}$, $2 = \{0, 1\}$, $\omega = \{0, 1, 2, \dots\} = \mathbb{N}$, etc. Observe that with this convention, for any ordinals α and β are ordinals, $\alpha \in \beta$ if and only if $\alpha < \beta$, and $\alpha \subseteq \beta$ if and only if $\alpha \leq \beta$. We can also give a concrete description of the successor and supremum operations: $\alpha + 1 = \alpha \cup \{\alpha\}$, and $\sup S = \bigcup_{\alpha \in S} \alpha$.

Proposition 1.5. *For every set X, there is some ordinal α such that $|\alpha| \not\leq |X|$, i.e., there is no injective function $\alpha \to X$.*

Proof. Let $S(X)$ be the set of all ordinals which admit an injective function to X. It suffices to find an ordinal α which is not in $S(X)$. If $S(X)$ has a greatest element β, then we can take $\alpha = \beta + 1$. If $S(X)$ has no greatest element, then we can take $\alpha = \sup S$. $\square$

The careful reader may notice something troubling at this point: Why can't we apply the proof of Proposition 1.5, replacing $S(X)$ with the collection Ord of *all* ordinals, to obtain an ordinal which is not in Ord, and hence a contradiction? This is sometimes called the "Burali-Forti paradox". The resolution is that Ord is not a set. Instead, it is a *proper class*: a collection that we can talk about, but which is somehow "too big" to be a set. This highlights a subtlety in the proof of Proposition 1.5: How do we know that $S(X)$ is a set, when the class of all ordinals is not a set? To answer this question and give a more formal proof, we have to carefully apply the set existence properties established by the axioms of ZFC. I will omit a formal proof, since our goal in this section is to avoid delving into these kinds of details.

Definition 1.6. A linear order $(L, \leq)$ is a **well-order** if every non-empty subset of L has a least element.

Proposition 1.7. *The class of ordinals is a well-order.*

Proof. Let X be a set of ordinals with no least element. Let $P(\alpha)$ be the property "$\alpha \notin X$". If all ordinals $\beta < \alpha$ satisfy P, then $\beta \notin X$ for all $\beta < \alpha$. Since X has no least element, $\alpha \notin X$, so α satisfies P. By strong induction, every ordinal satisfies P, so X is empty. $\square$

Note that every subset of a well-order is a well-order. It follows that every ordinal α (viewed, remember, as the set of ordinals less than α) is a well-order.

Proposition 1.8. *Suppose $(L, \leq)$ is a well-ordered set. Then there is an ordinal α such that $(\alpha, \leq) \cong (L, \leq)$.*

Proof. Suppose for contradiction that there is no such ordinal α. We construct a family of functions f_β for all ordinals β, such that:

(1) f_β is an embedding $(\beta, \leq) \to (L, \leq)$.

(2) $\mathrm{im}(f_\beta)$ is a downwards-closed subset of L. That is, if $b \in \mathrm{im}(f_\beta)$ and $a \leq b$ in L, then $a \in \mathrm{im}(f_\beta)$.

(3) If $\beta < \beta'$, then $f_\beta \subseteq f_{\beta'}$.

In the base case, we let f_0 be the empty function $\varnothing \to L$.

For the successor step, given f_β, we define $f_{\beta+1}$. By our assumption, f_β is not an isomorphism, so f_β is not surjective. Thus $L \setminus \mathrm{im}(f_\beta)$ is non-empty. Let b the least element of $L \setminus \mathrm{im}(f_\beta)$, which exists because L is a well-order. Since $\mathrm{im}(f_\beta)$ is downwards-closed, b is greater than every element of $\mathrm{im}(f_\beta)$. Define $f_{\beta+1} = f_\beta \cup (\beta, b)$, the map which extends f_β by mapping β to b.

For the limit step, given $(f_\beta)_{\beta<\gamma}$ with γ a limit ordinal, define $f_\gamma = \bigcup_{\beta<\gamma} f_\beta$.

In each case, it is easy to verify conditions 1, 2, and 3. The result of this construction is that every ordinal admits an injective map into L, contradicting Proposition 1.5. $\square$

We can carry out essentially the same argument for unordered sets.

Theorem 1.9 (Zermelo). *Every set is in bijection with an ordinal.*

Proof. Let X be a set. Suppose for contradiction that there is no ordinal α such that X is in bijection with α. We construct an injective function $f_\beta \colon \beta \to X$ for all ordinals β, such that when $\beta < \beta'$, $f_\beta \subseteq f_{\beta'}$.

In the base case, we let f_0 be the empty function $\varnothing \to X$.

For the successor step, given f_β, we define $f_{\beta+1}$. By our assumption, f_β is not a bijection, so f_β is not surjective. Thus $X \setminus \mathrm{im}(f_\beta)$ is non-empty. Let b an arbitrary element of $X \setminus \mathrm{im}(f_\beta)$. Define $f_{\beta+1} = f_\beta \cup (\beta, b)$, the map which extends f_β by mapping β to b.

For the limit step, given $(f_\beta)_{\beta<\gamma}$ with γ a limit ordinal, define $f_\gamma = \bigcup_{\beta<\gamma} f_\beta$.

The result of this construction is that every ordinal admits an injective map into X, contradicting Proposition 1.5. $\square$

The key difference between the proofs of Proposition 1.8 and Theorem 1.9 is that the well-order on L gave us a canonical choice of where to map each ordinal, while in the proof of Zermelo's theorem, we had to make many arbitrary choices. The proof of Proposition 1.8 (and everything else we've done up to this point) goes through without the axiom of choice, while the proof of Theorem 1.9 requires the axiom of choice (as will much of what we prove from here on out).

Zermelo's theorem is often called the Well-Ordering Theorem (because it implies that every set can be well-ordered), and it turns out to be equivalent to the axiom of choice over ZF set theory. It makes possible the following definition, which gives us a much clearer picture of the possible cardinalities of sets.

Definition 1.10. For any set X, the **cardinality** of X, denoted $|X|$, is the least ordinal κ such that there is a bijective function $X \to \kappa$. A **cardinal** is an ordinal κ such that $|\kappa| = \kappa$ (i.e., an ordinal which is not in bijection with any smaller ordinal).

Exercise 4. Show that this notation agrees with that in Definition 1.1. That is, there is an injective function $X \to Y$ if and only if $|X| \leq |Y|$ (as ordinals), and there is a bijective function $X \to Y$ if and only if $|X| = |Y|$ (as ordinals).

Lemma 1.11. *Every infinite cardinal is a limit ordinal.*

Proof. It suffices to show that for every infinite successor ordinal $\alpha + 1$, there is a bijection between $\alpha + 1$ and α (since then $\alpha + 1$ is not a cardinal). Note that since $\alpha + 1$ is infinite, $\alpha \geq \omega$. Define:

$$f(\beta) = \begin{cases} \beta + 1 & \beta < \omega \\ \beta & \omega \leq \beta < \alpha \\ 0 & \beta = \alpha. \end{cases}$$

and verify that this gives a bijection $(\alpha + 1) \to \alpha$. □

Every finite ordinal is a cardinal: $0, 1, 2, \dots$, i.e. the natural numbers. The infinite cardinals, being a subclass of the ordinals, are well-ordered, and hence can be indexed by the ordinals:

- $\aleph_0 = \omega$ is the smallest infinite cardinal.
- Given $\kappa = \aleph_\alpha$, we define $\kappa^+ = \aleph_{\alpha+1}$ to be the least cardinal greater than $\aleph_\alpha$ (this is a **successor** cardinal).
- For a limit ordinal γ, we define $\aleph_\gamma = \sup\{\aleph_\alpha \mid \alpha < \gamma\}$ (this is a **limit** cardinal).

Exercise 5. Verify that the definition above makes sense:

(a) For any ordinal α, there is a least cardinal greater than $\aleph_\alpha$.

(b) For a limit ordinal γ, $\aleph_\gamma$, as defined above, is a cardinal.

(c) Every infinite cardinal is $\aleph_\alpha$ for some ordinal α.

(*Hint:* Prove by strong transfinite induction that for every ordinal β, if β is an infinite cardinal, then $\beta = \aleph_\alpha$ for some α.)

Note that, by convention, all successor cardinals and limit cardinals are uncountably infinite (i.e., 2 is not a successor cardinal, despite being the next cardinal after 1, and $\aleph_0$ is not a limit cardinal, despite being the supremum of the finite cardinals). Different sources may disagree on this point.

1.3 Cardinal arithmetic

Definition 1.12. Let κ and λ be cardinals.

- $\kappa + \lambda$ is the cardinality of the disjoint union

$$\kappa \sqcup \lambda = \{(0, \alpha) \mid \alpha \in \kappa\} \cup \{(1, \beta) \mid \beta \in \lambda\}.$$

- $\kappa \cdot \lambda$ is the cardinality of the Cartesian product

$$\kappa \times \lambda = \{(\alpha, \beta) \mid \alpha \in \kappa, \beta \in \lambda\}.$$

- κ^λ is the cardinality of the set of functions

$$\mathrm{Fun}(\lambda, \kappa) = \{f \mid f \colon \lambda \to \kappa \text{ is a function}\}.$$

For finite cardinals, these operations agree with the usual addition, multiplication, and exponentiation of natural numbers. For infinite cardinals, we will now show that addition and multiplication are very simple operations. In contrast, cardinal exponentiation can be extremely complicated: most facts about it are independent from ZFC set theory.

To understand cardinal multiplication, we introduce a canonical well-ordering of the product $\kappa \times \kappa = \{(\beta, \beta') \mid \beta, \beta' < \kappa\}$:

$$\begin{aligned}(\beta, \beta') < (\gamma, \gamma') \text{ iff } & \max(\beta, \beta') < \max(\gamma, \gamma') \\ & \text{or } \max(\beta, \beta') = \max(\gamma, \gamma') \text{ and } \beta < \gamma \\ & \text{or } \max(\beta, \beta') = \max(\gamma, \gamma') \text{ and } \beta = \gamma \text{ and } \beta' < \gamma'\end{aligned}$$

Exercise 6. Show that $<$, as defined above, well-orders $\kappa \times \kappa$.

Theorem 1.13. *For every infinite cardinal κ, $(\kappa \times \kappa, <) \cong (\kappa, <)$. As a consequence, $\kappa \cdot \kappa = \kappa$.*

Proof. Since every infinite cardinal is $\aleph_\alpha$ for some ordinal α, we can prove things about the infinite cardinals by transfinite induction (over their indexing ordinals). So we assume that $(\lambda \times \lambda, <) \cong (\lambda, <)$ for all infinite cardinals less than κ and prove the same for κ.

By Exercise 6, $(\kappa \times \kappa, <)$ is a well-ordered set, so by Proposition 1.8, there is an isomorphism $f \colon (\delta, <) \to (\kappa \times \kappa, <)$ for some ordinal δ. It suffices to show

$\delta = \kappa$. Note that $\kappa \leq \kappa \cdot \kappa$, since the function $\alpha \mapsto (\alpha, 0)$ is an injective map $\kappa \to \kappa \times \kappa$. So $\delta \geq |\delta| = |\kappa \times \kappa| \geq \kappa$.

Suppose for contradiction that $\delta > \kappa$, i.e., $\kappa \in \delta$. Let $(\beta, \beta') = f(\kappa)$. Then f restricts to an isomorphism from κ to $X = \{(\gamma, \gamma') \mid (\gamma, \gamma') < (\beta, \beta')\}$.

Now let $\alpha = \max(\beta, \beta') + 1$. Note that since $(\beta, \beta') \in \kappa \times \kappa$, $\max(\beta, \beta') < \kappa$. By Lemma 1.11, κ is a limit ordinal, so also $\alpha < \kappa$. In particular, $\lambda = |\alpha|$ is a cardinal less than κ.

We have $X \subseteq \alpha \times \alpha$, since if $(\gamma, \gamma') < (\beta, \beta')$, then $\max(\gamma, \gamma') \leq \max(\beta, \beta') < \alpha$, so $\gamma \in \alpha$ and $\gamma' \in \alpha$. Thus $\kappa = |X| \leq |\alpha \times \alpha| = \lambda \cdot \lambda$. If λ is finite, then $\lambda \cdot \lambda$ is finite. If λ is infinite, then $\lambda \cdot \lambda = \lambda$ by induction. In either case, we have a contradiction, since κ is an infinite cardinal greater than λ. □

Corollary 1.14. *Let κ and λ be non-zero cardinals such that at least one of κ and λ is infinite. Then:*

$$\kappa + \lambda = \max(\kappa, \lambda)$$
$$\kappa \cdot \lambda = \max(\kappa, \lambda)$$

Proof. Cardinal addition and multiplication are commutative, since there are bijections $\kappa \sqcup \lambda \to \lambda \sqcup \kappa$ and $\kappa \times \lambda \to \lambda \times \kappa$. So without loss of generality, assume $\kappa \leq \lambda$. In particular, $\kappa \neq 0$ and λ is infinite. Then:

$$\lambda \leq \kappa \cdot \lambda \leq \lambda \cdot \lambda = \lambda$$

and

$$\lambda \leq \kappa + \lambda \leq \lambda + \lambda = 2 \cdot \lambda = \lambda.$$ □

The following result is useful for estimating the sizes of infinite unions.

Corollary 1.15. *If $(X_i)_{i \in I}$ is an infinite family of sets, then the cardinality of $\bigcup_{i \in I} X_i$ is bounded above by the supremum of $|I|$ and $|X_i|$ for all $i \in I$.*

Proof. If $\kappa = \sup_{i \in I} |X_i|$, then

$$\left| \bigcup_{i \in I} X_i \right| \leq \left| \bigsqcup_{i \in I} \kappa \right| = |I \times \kappa| = \max(|I|, \kappa).$$ □

Exercise 7. Prove the following, for all cardinals κ, λ, κ', λ', μ:

(a) If $\lambda > 0$, then $\kappa \leq \kappa^\lambda$.

(b) If $\kappa > 1$, then $\lambda \leq \kappa^\lambda$.

(c) If $0 < \kappa \leq \kappa'$ and $\lambda \leq \lambda'$, then $\kappa^\lambda \leq {\kappa'}^{\lambda'}$.

(d) $\kappa^{\lambda + \mu} = \kappa^\lambda \cdot \kappa^\mu$.

(e) $(\kappa \cdot \lambda)^\mu = \kappa^\mu \cdot \lambda^\mu$.

(f) $(\kappa^\lambda)^\mu = \kappa^{\lambda \cdot \mu}$.

For any set X, there is a natural bijection between $\mathcal{P}(X)$ and $\mathrm{Fun}(X, 2)$, so $|\mathcal{P}(X)| = 2^{|X|}$. Thus Cantor's theorem implies $\kappa < 2^\kappa$ for all cardinals κ. It is natural to ask whether 2^κ is the successor κ^+ or some larger cardinal.

The **Continuum Hypothesis** (CH) is the statement that $2^{\aleph_0} = \aleph_1$. The **Generalized Continuum Hypothesis** (GCH) is the statement that $2^\kappa = \kappa^+$ for all infinite cardinals κ.

Fact 1.16. Both CH and GCH are independent of ZFC.

Definition 1.17. Let κ be an infinite cardinal. The **cofinality** of κ, $\mathrm{cf}(\kappa)$, is the least cardinal λ such that κ can be written as a union of λ-many sets of cardinality $< \kappa$:

$$\kappa = \bigcup_{\alpha \in \lambda} X_\alpha \quad \text{with } |X_\alpha| < \kappa \text{ for all } \alpha \in \lambda.$$

By Corollary 1.14, a finite union of sets of cardinality $< \kappa$ has cardinality $< \kappa$, so $\mathrm{cf}(\kappa)$ cannot be finite. On the other hand, we always have $\kappa = \bigcup_{\alpha<\kappa}\{\alpha\}$, a union of κ-many sets of size 1. So $\aleph_0 \leq \mathrm{cf}(\kappa) \leq \kappa$.

Definition 1.18. An infinite cardinal κ is **regular** if $\mathrm{cf}(\kappa) = \kappa$. Otherwise, κ is **singular**.

Example 1.19. $\aleph_0$ is regular, since a union of finitely many finite sets is finite.

The successor cardinal $\aleph_1$ is regular, since a union of countably many countable sets is countable (by Corollary 1.15).

The limit cardinal $\aleph_\omega$ is singular with $\mathrm{cf}(\aleph_\omega) = \aleph_0$, since $\aleph_\omega = \bigcup_{n\in\omega} \aleph_n$.

The next proposition generalizes the proof that $\aleph_1$ is regular.

Proposition 1.20. *Every successor cardinal is regular.*

Proof. Write our successor cardinal as κ^+, and let $\lambda = \mathrm{cf}(\kappa^+)$. Then we can write $\kappa^+ = \bigcup_{\alpha\in\lambda} X_\alpha$, with $|X_\alpha| < \kappa^+$ for all α. It follows that $|X_\alpha| \leq \kappa$ for all α, so $\sup_{\alpha\in\lambda} |X_\alpha| \leq \kappa$. But then by Corollary 1.15,

$$\kappa^+ = \left|\bigcup_{\alpha\in\lambda} X_\alpha\right| \leq \max(\lambda, \sup_{\alpha\in\lambda} |X_\alpha|) \leq \max(\lambda, \kappa),$$

so $\lambda = \kappa^+$. $\square$

Example 1.19 and Proposition 1.20 suggest the question: are there any regular limit cardinals? This turns out to be rather subtle from the metamathematical point of view.

Definition 1.21. A regular limit cardinal is called **weakly inaccessible**. A cardinal is a **strong limit cardinal** if for all $\lambda < \kappa$, we have $2^\lambda < \kappa$. A regular strong limit cardinal is called **(strongly) inaccessible**.

Fact 1.22. If ZFC is consistent, then ZFC does not prove that weakly inaccessible cardinals exist. Even stronger, ZFC does not prove that it is *consistent* that weakly inaccessible cardinals exist.

This is different from the status of CH and GCH. If ZFC is consistent, then ZFC proves neither the CH nor its negation. On the other hand, if ZFC is consistent, then ZFC does not prove that weakly inaccessible cardinals exist. However, it is possible that ZFC outright proves that there are no weakly inaccessible cardinals. Nevertheless, most set theorists believe that weakly and strongly inaccessible cardinals are consistent with ZFC.

We end with one of the few facts about cardinal exponentiation that is actually provable in ZFC. It is instructive to think about how the proof of Theorem 1.23 can be viewed as a generalization of Cantor's diagonalization argument.

Theorem 1.23. *For any infinite cardinal κ, $\kappa^{\mathrm{cf}(\kappa)} > \kappa$.*

Proof. Let $\lambda = \mathrm{cf}(\kappa)$. By Exercise 7, $\kappa \leq \kappa^{\lambda}$.

Now suppose for contradiction that there is a bijection $f\colon \kappa \to \mathrm{Fun}(\lambda, \kappa)$. Since $\lambda = \mathrm{cf}(\kappa)$, we can write $\kappa = \bigcup_{\alpha \in \lambda} X_\alpha$, with $|X_\alpha| < \kappa$ for all $\alpha \in \lambda$. Fixing $\alpha \in \lambda$, let $Y_\alpha = \{f(\beta)(\alpha) \mid \beta \in X_\alpha\} \subseteq \kappa$. Then $|Y_\alpha| \leq |X_\alpha| < \kappa$, so we can pick some $\gamma_\alpha \in \kappa$ such that $\gamma \notin Y_\alpha$. Defining $g(\alpha) = \gamma_\alpha$, we have a function $g\colon \lambda \to \kappa$ such that for all $\alpha \in \lambda$, and all $\beta \in X_\alpha$, $g(\alpha) \neq f(\beta)(\alpha)$.

Now since f is surjective, there is some $\beta^* \in \kappa$ such that $f(\beta^*) = g$. Since $\bigcup_{\alpha \in \lambda} X_\alpha = \kappa$, there is some $\alpha^* \in \lambda$ such that $\beta^* \in X_{\alpha^*}$. But then we have $g(\alpha^*) = f(\beta^*)(\alpha^*)$, contradicting our choice of g. $\square$

2 Types and saturated models

We now return to the context of model theory. Recall our convention: L is a countable first-order language and T is a complete L-theory with infinite models.

2.1 Types

Let's begin by setting some notation and terminology. A **context** is a tuple of distinct variables. To simplify notation, we often write variable contexts as a single letter, e.g. x, and we write $|x|$ for the cardinality of the set of variables in x. Contexts are usually finite, but they need not be. We say a formula or type is in context x if all of its free variables come from x. We write $\varphi(x)$ or $\Sigma(x)$ to indicate that the formula φ or the set of formulas Σ is in context x.

When I speak of a **set** A, I mean a subset of a model of T, i.e., $A \subseteq M \models T$. An **interpretation** of a context x in a set A is an assignment of an element of A to each variable in x (equivalently, a tuple from A indexed by x). We write A^x for the set of interpretations of x in A. For every $n \in \omega$, there is a canonical context $x_1, \ldots, x_n$. We write A^n for the set of interpretations of this context in A (equivalently, the tuples from A of length n).

For a set A, we define the language L_A obtained from A by adjoining a new constant symbol for each element of A. We define $T_A = \mathrm{Th}_{L_A}(M)$. We sometimes call the new constant symbols "parameters from A". Depending on whether we want to make the parameters explicit, a formula with parameters from A can be written as $\varphi(x)$, where $\varphi(x)$ is an L_A-formula, or as $\psi(x, a)$, where $\psi(x, y)$ is an L-formula and $a \in A^y$. We write $F_x(A)$ for the set of all L_A-formulas in context x.

Definition 2.1. Suppose M and N are models of T and $A \subseteq M$. A **partial elementary map** is a function $f \colon A \to N$ such that for all formulas $\varphi(x)$ and all $a \in A^x$, $M \models \varphi(a)$ if and only if $N \models \varphi(f(a))$.

Note that a model $N \models T_A$ is the same as a model $N \models T$, equipped with a partial elementary map $f \colon A \to N$.

A **partial type** over A in context x is a set of formulas in $F_x(A)$. A partial type $\Sigma(x)$ over A is **consistent** if $\Sigma(x) \cup T_A$ is satisfiable: there a model $N \models T_A$ and a tuple $n \in N^x$ such that $N \models \varphi(n)$ for all $\varphi \in \Sigma(x)$. A partial type $\Sigma(x)$ is **complete** if it is consistent and $\varphi(x) \in \Sigma(x)$ or $\neg\varphi(x) \in \Sigma(x)$ for all $\varphi(x) \in F_x(A)$. We write $S_n(A)$ for the space of complete types over A in context $x_1, \ldots, x_n$. More generally, for a context x, we write $S_x(A)$ for the space of complete types over A in context x.

If $p \in S_x(A)$, with $A \subseteq N$ and $f \colon A \to N$ is a partial elementary map, define the **pushforward of** p **along** f:

$$f_* p = \{\varphi(x, f(a)) \mid \varphi(x, a) \in p\}.$$

Note that $f_* p \cup T_{f(A)}$ is consistent, since this differs from $p \cup T_A$ only in the names used for the constant symbols. Thus $f_* p \in S_x(f(A))$.

Proposition 2.2. *Suppose A is a set and x is a context.*

(1) $|F_x(A)| = \max(\aleph_0, |A|, |x|)$.

(2) $|S_x(A)| \leq 2^{|F_x(A)|}$.

Proof. Let $\kappa = \max(\aleph_0, |A|, |x|)$. A formula in $F_x(A)$ is a finite sequence of symbols from an alphabet which includes the symbols in L (which we assume is countable), the constant symbols from A, the variables in x, and finitely many extra symbols such as logical connectives and parentheses. So the alphabet has size $\leq \kappa$.

Let $\mathrm{Seq}(n)$ be the set of sequences of length n from this alphabet. Then $|\mathrm{Seq}(n)| \leq \kappa^n = \kappa$, since κ is infinite. Now $F_x(A) \subseteq \bigcup_{n\in\omega} \mathrm{Seq}(n)$, so $|F_x(A)| \leq \max(\aleph_0, \kappa) = \kappa$ by Corollary 1.15.

Conversely, we have $\aleph_0 \leq |F_x(A)|$ (there are infinitely many formulas, e.g. $\top$, $\top \wedge \top$, $\top \wedge \top \wedge \top$, etc.), $|A| \leq |F_x(A)|$ (e.g. $a = a$ for each $a \in A$), and $|x| \leq |F_x(A)|$ (e.g. $x_i = x_i$ for each x_i in x).

A type is a set of formulas, so $|S_x(A)| \leq |\mathcal{P}(F_x(A))| = 2^{|F_x(A)|}$. $\square$

Note two important special cases of Proposition 2.2: If x and A are countable, then $|F_x(A)| = \aleph_0$ and $|S_x(A)| \leq 2^{\aleph_0}$. If x is countable and A is infinite, then $|F_x(A)| = |A|$ and $|S_x(A)| \leq 2^{|A|}$.

Lemma 2.3. *Suppose $p(x) \in S_x(A)$ where $A \subseteq M \models T$. Then there is an elementary extension $M \preceq M'$ with $|M'| = \max(|M|, |x|)$ such that $p(x)$ is realized in M'.*

Proof. Since $p(x) \in S_x(A)$, $p(x) \cup T_A$ is consistent. This gives a model N with a partial elementary map $f\colon A \to N$ and a tuple $n \in N^x$ realizing $f_*p(x)$. To improve this to an elementary extension of M, we need to show that that $p(x) \cup T_M$ is consistent.

Since $p(x)$ and T_M are closed under conjunction, by compactness it suffices to show that for any formula $\varphi(x) \in p(x)$ and any L_M-sentence $\psi \in T_M$, $\{\varphi(x), \psi\}$ is consistent. We can write $\varphi(x)$ as $\varphi(x, a)$ and ψ as $\psi(a, m)$, where a is the tuple of all parameters from A appearing in $\varphi(x)$ and ψ, and m is the tuple of all parameters from $M \setminus A$ appearing in ψ. Now $M \models \exists y\, \psi(a, y)$, so $N \models \exists y\, \psi(f(a), y)$. Let $n' \in N^y$ be a witness to the existential quantifier. Interpreting the variables x as n, the constant symbols a as $f(a)$, and the constant symbols m and n', N is a model of $\{\varphi(x), \psi\}$. Thus $p(x) \cup T_M$ is consistent.

Now a model of $p(x) \cup T_M$ consists of an elementary extension $M \preceq M'$ and $a \in (M')^x$ such that $M' \models p(a)$. By downward Löwenheim–Skolem, we can take an elementary substructure of M' containing M and a of cardinality $\max(|M|, |x|)$. $\square$

2.2 Saturated models

Saturation is a useful notion of "completeness" of a model M. In this section, we will establish the useful properties of saturated models of size κ: universality (for models of size $\leq \kappa$), homogeneity (for subsets of size $< \kappa$), and uniqueness.

Definition 2.4. Let κ be an infinite cardinal. A model $M \models T$ is **κ-saturated** if for all $A \subseteq M$ with $|A| < \kappa$, every type in $S_1(A)$ is realized in M.

Theorem 2.5. *Suppose $M \models T$ is κ-saturated. Given a model $N \models T$ with $|N| \leq \kappa$, a set $A \subseteq N$ with $|A| < \kappa$, and a partial elementary map $f\colon A \to M$, f extends to an elementary embedding $N \to M$.*

Proof. Let $\lambda = |N \setminus A|$, and enumerate $N \setminus A$ as $(n_\alpha)_{\alpha<\lambda}$. For all $\alpha \leq \lambda$, write $N_\alpha = A \cup \{n_\beta \mid \beta < \alpha\}$. Note that $|N_\alpha| \leq |A| + |\alpha| < \kappa$ for all $\alpha < \lambda$. We build a sequence of partial elementary maps $(f_\alpha)_{\alpha\leq\lambda}$ by transfinite recursion, so that $\mathrm{dom}(f_\alpha) = N_\alpha$.

Take $f_0 = f$ with domain $N_0 = A$. It is partial elementary by hypothesis.

Given f_α with $\alpha < \lambda$, let $p_\alpha = \mathrm{tp}(n_\alpha/N_\alpha)$. Then $(f_\alpha)_* p_\alpha \in S_1(f(N_\alpha))$. Since $|N_\alpha| < \kappa$, $|f(N_\alpha)| < \kappa$, so $(f_\alpha)_* p_\alpha$ is realized in M by some element m_α. Define $f_{\alpha+1}$ to be the map extending f_α by mapping n_α to m_α. Check that $f_{\alpha+1}$ is partial elementary.

When $\gamma \leq \lambda$ is a limit ordinal, define $f_\gamma = \bigcup_{\alpha<\gamma} f_\alpha$. Check that f_γ is partial elementary and has domain N_γ.

Now f_λ is a partial elementary map $N = N_\lambda \to M$, i.e., an elementary embedding $N \to M$. □

As a corollary, we find that the definition of κ-saturation, which only say that 1-types are realized, implies that types in arbitrary contexts (of size at most κ) are realized.

Corollary 2.6. *Suppose $M \models T$ is κ-saturated. For every set $A \subseteq M$ with $|A| < \kappa$ and every context x with $|x| \leq \kappa$, every type $p \in S_x(A)$ is realized in M.*

Proof. Since $p(x)$ is consistent, there is some model $N \models T$ together with a partial elementary map $f\colon A \to N$ and a tuple $n \in N^x$ realizing $f_* p(x)$. By Löwenheim–Skolem, we may assume $|N| = \max(|A|, |x|) \leq \kappa$. Let $g\colon f(A) \to M$ be the partial elementary map inverse to A. By Theorem 2.5, g extends to an elementary embedding $h\colon N \to M$. Then $h(n)$ realizes $h_* f_* p(x) = p(x)$, since f and h are inverses. □

Corollary 2.7 (Universality)**.** *Suppose $M \models T$ is κ-saturated. For every model $N \models T$ with $|N| \leq \kappa$, there is an elementary embedding $N \to M$.*

Proof. Take $A = \varnothing \subseteq N$. Since T is complete, the empty function $f\colon A \to N$ is partial elementary. By Theorem 2.5, f extends to an elementary embedding $N \to M$. □

Note that Corollary 2.7 implies that if M is κ-saturated, then $\kappa \leq |M|$, since any model of cardinality κ embeds in M. So the "most saturated" a model could be is $|M|$-saturated.

Definition 2.8. M is **saturated** if it is $|M|$-saturated.

In the next theorem, we use the same proof strategy as in Theorem 2.5, but instead of just constructing a map $N \to M$, we want to go "back-and-forth" to construct an isomorphism. For this reason, we need to assume M and N are saturated and of the same cardinality.

Theorem 2.9. *Suppose $M \models T$ and $N \models T$ are saturated with $|M| = |N| = \kappa$. Let $A \subseteq M$ be a subset with $|A| < \kappa$ and $f\colon A \to N$ a partial elementary map. Then f extends to an isomorphism $M \cong N$.*

Proof. Since $|A| = |f(A)| < \kappa$, we have $|M \setminus A| = \kappa$ and $|N \setminus f(A)| = \kappa$. Enumerate $M \setminus A$ as $(m_\alpha)_{\alpha<\kappa}$ and $N \setminus f(A)$ as $(n_\alpha)_{\alpha<\kappa}$. We build a sequence of partial elementary maps $(f_\alpha)_{\alpha\leq\kappa}$ extending f by transfinite recursion. Along the way, we define sequences $(m'_\alpha)_{\alpha<\kappa}$ and $(n'_\alpha)_{\alpha<\kappa}$. If we set $M_\alpha = A \cup \{m_\beta, m'_\beta \mid \beta < \alpha\}$ and $N_\alpha = f(A) \cup \{n_\beta, n'_\beta \mid \beta < \alpha\}$, we ensure that $\mathrm{dom}(f_\alpha) = M_\alpha$ and $\mathrm{ran}(f_\alpha) = N_\alpha$. In particular, for $\alpha < \kappa$, $|M_\alpha| \leq |A| + 2|\alpha| < \kappa$ and similarly $|N_\alpha| < \kappa$.

Take $f_0 = f$. This has domain $M_0 = A$ and range $N_0 = f(A)$. It is partial elementary by hypothesis.

Given f_α with $\alpha < \kappa$, let $p_\alpha = \mathrm{tp}(m_\alpha/M_\alpha)$. Then $(f_\alpha)_* p_\alpha \in S_1(N_\alpha)$. Since $|N_\alpha| < \kappa$, $(f_\alpha)_* p$ is realized in N by some element n'_α. Define g_α to be the map extending f_α by mapping m_α to n'_α. Check that g_α is partial elementary.

Write g_α^{-1} for the partial elementary map with domain $\mathrm{ran}(g_\alpha)$ which is inverse to g_α. Now let $q_\alpha = \mathrm{tp}(n_\alpha/\mathrm{ran}(g_\alpha))$. Then $(g_\alpha^{-1})_* q_\alpha \in S_1(\mathrm{dom}(g_\alpha))$. Since $|\mathrm{dom}(g_\alpha)| = |M_\alpha \cup \{m_\alpha\}| < \kappa$, $(g_\alpha^{-1})_* q_\alpha$ is realized in M by some element m'_α. Define $f_{\alpha+1}$ to be the map extending g_α by mapping m'_α to n_α. Check that $f_{\alpha+1}$ is partial elementary, and note $\mathrm{dom}(f_{\alpha+1}) = M_\alpha \cup \{m_\alpha, m'_\alpha\} = M_{\alpha+1}$ and $\mathrm{ran}(f_{\alpha+1}) = N_\alpha \cup \{n_\alpha, n'_\alpha\} = N_{\alpha+1}$.

When $\gamma \leq \lambda$ is a limit ordinal, define $f_\gamma = \bigcup_{\alpha<\gamma} f_\alpha$. Check that f_γ is partial elementary, $\mathrm{dom}(f_\gamma) = M_\gamma$, and $\mathrm{ran}(f_\gamma) = N_\gamma$.

Now f_κ is a surjective partial elementary map $M = M_\kappa \to N_\kappa = N$, i.e., an isomorphism $M \cong N$. □

Corollary 2.10 (Homogeneity)**.** *Suppose M is a saturated model of T, $A \subseteq M$ is a subset with $|A| < |M|$, and $f\colon A \to M$ is a partial elementary map. Then f extends to an automorphism of M.*

Proof. Take $N = M$ in Theorem 2.9. □

The significance of Corollary 2.10 is that inside a saturated model, "all structural properties" of a tuple are captured by its complete type.

Corollary 2.11 (Uniqueness)**.** *Any two saturated models of T of the same cardinality are isomorphic.*

Proof. Suppose M and N are saturated models of T with $|M| = |N|$. Let $A = \varnothing \subseteq M$. Since T is complete, the empty function $f\colon A \to N$ is partial elementary. By Theorem 2.9, f extends to an isomorphism $M \cong N$. $\square$

Morley's original proof of the categoricity theorem proceeds by showing that if T is κ-categorical for some uncountable κ, then every uncountable model of T is saturated. It follows from Corollary 2.11 that T is categorical in every uncountable cardinal. As noted in the introduction, we will be following a different proof strategy (though the fact that every uncountable model is saturated will be a consequence of our proof).

2.3 Existence of saturated models

The question remains: do any saturated models exist? The first thing we need to do is to realize lots of types at once.

Lemma 2.12. *Let $X \subseteq S_1(A)$ with $A \subseteq M \models T$. Then there is an elementary extension $M \preceq M'$ such that every type in X is realized in M'. Further, we can take $|M'| \leq \max(|M|, |X|)$.*

Proof. Let $\kappa = |X|$, and enumerate $X = (p_\alpha)_{\alpha<\kappa}$. Write $X_\alpha = \{p_\beta \mid \beta < \alpha\}$. We build an elementary chain of models $(M_\alpha)_{\alpha\leq\kappa}$ by transfinite recursion such that all of the types in X_α are realized in M_α, and $|M_\alpha| \leq \max(|M|, \kappa)$.

Take $M_0 = M$. Since $X_0 = \varnothing$, M_0 vacuously realizes all types in X_0. And $|M_0| \leq \max(|M|, \kappa)$.

Given M_α with $\alpha < \kappa$, by Lemma 2.3, we can find an elementary extension $M_\alpha \preceq M_{\alpha+1}$ such that p_α is realized in $M_{\alpha+1}$ and $|M_{\alpha+1}| = |M_\alpha| \leq \max(|M|, \kappa)$. Since the extension is elementary, all of the types in X_α are realized in $M_{\alpha+1}$ (by the same elements realizing them in M_α).

When $\gamma \leq \kappa$ is a limit ordinal, define $M_\gamma = \bigcup_{\alpha<\gamma} M_\alpha$. Then M_γ is an elementary extension of each M_α, so it realizes all the types in $X_\gamma = \bigcup_{\alpha<\gamma} X_\alpha$. And $|M_\gamma| \leq \max(|\gamma|, \sup_{\alpha<\gamma} |M_\alpha|) \leq \max(|M|, \kappa)$, since $|\gamma| \leq \kappa$.

Now M_κ is an elementary extension of M realizing all types in $X_\kappa = X$, and $|M_\kappa| \leq \max(|M|, \kappa) = \max(|M|, |X|)$. $\square$

Now let's try to construct κ-saturated model. We can use Lemma 2.12 to realize types, but upon passing to an elementary extension, we have new types over new parameter sets to realize. So we need to build another elementary chain. In order to "catch our tail", we need the length of this elementary chain to have length a regular cardinal $\geq \kappa$.

Theorem 2.13. *Let κ be an infinite cardinal. Then every model $M \models T$ has a κ-saturated elementary extension.*

Proof. First, we may assume that κ is regular. If κ is singular, we use the argument below to build a κ^+-saturated elementary extension M' of M (since κ^+ is regular, Proposition 1.20). In particular, M' is κ-saturated.

Now, for a regular cardinal κ, we build an elementary chain $(M_\alpha)_{\alpha\leq\kappa}$ by transfinite recursion. Let $M_0 = M$. Given M_α, let $M_{\alpha+1}$ be an elementary extension of M_α realizing all types in $S_1(M_\alpha)$ by Lemma 2.12. Given a limit ordinal $\gamma \leq \kappa$, let $M_\gamma = \bigcup_{\alpha<\gamma} M_\alpha$.

To show that M_κ is κ-saturated, let $A \subseteq M_\kappa$ with $|A| < \kappa$, and let $p \in S_1(A)$. First, I claim that there is some $\gamma < \kappa$ such that $A \subseteq M_\gamma$. We have $M_\kappa = \bigcup_{\alpha<\kappa} M_\alpha$. For each $a \in A$, let $f(a)$ be the least ordinal α such that $a \in M_\alpha$. Let $\gamma = \sup_{a\in A} f(a)$, so $A \subseteq M_\gamma$. Since κ is regular, $|A| < \kappa$, and $f(a) < \kappa$ for all a, $\gamma < \kappa$.

Now $p \in S_1(A)$ extends to a type $\widehat{p} \in S_1(M_\gamma)$ (e.g., the complete type of any realization of p in an elementary extension of M_γ). And $\widehat{p}$ is realized in $M_{\gamma+1} \preceq M_\kappa$, and hence in M_κ. $\square$

Unfortunately, the proof of Theorem 2.13 typically gives us a κ-saturated model which has size much larger than κ. Even if we start with $|M_0| = \aleph_0$, in the worst case we have $|M_{\alpha+1}| = |S_1(M_\alpha)| = 2^{|M_\alpha|}$, so it seems the proof requires us to iterate the powerset operation κ-many types. In order to get a κ-saturated model of cardinality κ, we either need to make some set-theoretic assumptions about κ or some model-theoretic assumptions about T.

Theorem 2.14. *Suppose κ is a strongly inaccessible cardinal. Then T has a saturated model of cardinality κ.*

Proof. Since κ is strongly inaccessible, κ is regular. We show that in the proof of Theorem 2.13, we can ensure that $|M_\alpha| < \kappa$ for all $\alpha < \kappa$.

We start with a countable M_0, so $|M_0| < \kappa$. Given $|M_\alpha| = \lambda < \kappa$, we can pick $M_{\alpha+1}$ so that $|M_{\alpha+1}| = |S_1(M_\alpha)| \leq 2^\lambda < \kappa$, since κ is a strong limit cardinal. For limit ordinals $\gamma < \kappa$, we have $|\bigcup_{\alpha<\gamma} M_\alpha| \leq \max(|\gamma|, \sup_{\alpha<\gamma} |M_\alpha|) < \kappa$, where $\sup_{\alpha<\gamma} |M_\alpha| < \kappa$ because κ is regular.

Finally, $|M_\kappa| \leq \max(|\kappa|, \sup_{\alpha<\kappa} |M_\alpha|) = \kappa$. Since M_κ is κ-saturated, in fact $|M_\kappa| = \kappa$. Thus M_κ is saturated. $\square$

Instead of assuming that κ is really big, we could instead assume that we just don't have that many types to realize. This leads naturally the notion of κ-stability.

Definition 2.15. Let κ be an infinite cardinal. T is **κ-stable** if for all sets A with $|A| \leq \kappa$ and all $n \in \omega$, we have $|S_n(A)| \leq \kappa$.

In the definition of κ-stability, we count n-types (for arbitrary n) over arbitrary sets A. Let us take a moment to note that it is equivalent to count 1-types over models. This is actually all we need for the construction of saturated models in Theorem 2.18 below.

For every $n \in \omega$ and every set A, there is a surjective map $r\colon S_{n+1}(A) \to S_n(A)$ given by $r(p) = \{\varphi \in p \mid \varphi \in F_n(A)\}$. This map restricts a type to its first n variables.

Lemma 2.16. *Let $q \in S_n(A)$, and let $(b_1, \ldots, b_n)$ be any realization of q (in any model containing A). With $r\colon S_{n+1}(A) \to S_n(A)$ defined as above, there is a bijection between $r^{-1}(\{q\}) = \{p \in S_{n+1}(A) \mid r(p) = q\}$ and $S_1(Ab_1 \ldots b_n)$.*

Proof. The bijection is given by $p(x_1, \ldots, x_n, x_{n+1}) \mapsto p(b_1, \ldots, b_n, x)$.

In one direction, if $r(p) = q$, then p contains q, and $p(b_1, \ldots, b_n, x)$ contains $q(b_1, \ldots, b_n) = T_{Ab_1 \ldots b_n}$, so $p(b_1, \ldots, b_n, x)$ is a consistent type over $Ab_1 \ldots b_n$. Since p is complete, also $p(b_1, \ldots, b_n, x)$ is complete, so it is in $S_1(Ab_1 \ldots b_n)$.

In the other direction, if $p(b_1, \ldots, b_n, x) \in S_1(Ab_1 \ldots b_n)$, then letting b_{n+1} be any realization of this type, we have $p(x_1, \ldots, x_n, x_{n+1}) = \mathrm{tp}(b_1, \ldots, b_n, b_{n+1}/A)$, so this type is in $S_{n+1}(A)$. □

Proposition 2.17. *T is κ-stable if and only if for all models $M \models T$ with $|M| \leq \kappa$, we have $|S_1(M)| \leq \kappa$.*

Proof. One direction is trivial. For the other direction, we first note that if $A \subseteq M \models T$ with $|A| \leq \kappa$, we have $|S_1(A)| \leq \kappa$. Indeed, by Löwenheim–Skolem, there is an elementary substructure $M' \preceq M$ with $A \subseteq M'$ and $|M'| \leq \kappa$. Then $|S_1(A)| \leq |S_1(M')| \leq \kappa$.

Next we prove by induction on n that $|S_n(A)| \leq \kappa$ whenever $|A| \leq \kappa$. In the base case, we always have $|S_0(A)| = 1$. For the inductive step, we defined above a surjective map $r\colon S_{n+1}(A) \to S_n(A)$. So $|S_{n+1}(A)| = \left|\bigcup_{q \in S_n(A)} r^{-1}(\{q\})\right|$. For each $q \in S_n(A)$, pick a realization $(b_1, \ldots, b_n)$ in some model containing A. By Lemma 2.16, $|r^{-1}(\{q\})| = |S_1(Ab_1 \ldots b_n)| \leq \kappa$, since $|Ab_1 \ldots b_n| \leq \kappa + n = \kappa$.

By induction, $|S_n(A)| \leq \kappa$. So we have written $S_{n+1}(A)$ as a union of at most κ-many sets of size at most κ, and it follows that $|S_{n+1}(A)| \leq \kappa$. □

Theorem 2.18. *Let κ be a regular cardinal and assume T is κ-stable. Then T has a saturated model of cardinality κ.*

Proof. We show that in the proof of Theorem 2.13, we can ensure that $|M_\alpha| \leq \kappa$ for all $\alpha < \kappa$.

We start with an arbitrary $M_0 \models T$ with $|M_0| = \kappa$. Given M_α, we can pick $M_{\alpha+1}$ so that $|M_{\alpha+1}| = |S_1(M_\alpha)| = \kappa$, by κ-stability. For limit ordinals $\gamma \leq \kappa$, we have $|\bigcup_{\alpha<\gamma} M_\alpha| \leq \max(|\gamma|, \sup_{\alpha<\gamma} |M_\alpha|) \leq \kappa$. So $|M_\gamma| = \kappa$.

In particular, $|M_\kappa| = \kappa$, and M_κ is κ-saturated, so M_κ is saturated. □

Victor Harnik improved Theorem 2.18 by removing the hypothesis that κ is regular. The proof is significantly more difficult, since the construction of Theorem 2.13 relies crucially on regularity.

We end with two more sufficient conditions for the existence of saturated models. The first is a set-theoretic hypothesis, the second is a model-theoretic property of T.

Exercise 8. Suppose GCH holds at λ in the sense that $2^\lambda = \lambda^+$. Show that T has a saturated model of cardinality $2^\lambda = \lambda^+$. In particular, if CH holds, then T has a saturated model of cardinality $2^{\aleph_0} = \aleph_1$.

Hint: Mimic the proof of Theorem 2.13 with the regular cardinal $\kappa = \lambda^+ = 2^\lambda$. Build an elementary chain of models all of which have size κ. Instead of realizing *all* types over M_α in $M_{\alpha+1}$, just realize each $p \in S_1(A)$ for each $A \subseteq M_\alpha$ with $|A| < \kappa$. Note that you have to do some cardinal arithmetic to show that there are only at most κ-many such types.

Note that Theorem 2.14 and Exercise 8 both require strong set-theoretic hypotheses: we cannot prove in ZFC that there are any strongly inaccessible cardinals or that GCH holds at any cardinals. It follows from a theorem of Hugh Woodin that there are theories T such that (assuming the consistency of a large cardinal axiom) it is consistent with ZFC that T has no saturated models at all.

Definition 2.19. A theory T is **small** if $|S_n(\varnothing)| \leq \aleph_0$ for all $n \in \omega$.

Exercise 9. Show that T is small if and only if for any *finite* set A, $|S_1(A)| \leq \aleph_0$.
Hint: Use Lemma 2.16.

Exercise 10 (if you have not seen this result before)**.** Show that T has a countable saturated model if and only if T is small.

Hint: Use Exercise 9. The strategy for building the saturated model is similar to the strategy from the hint from Exercise 8 (but the cardinal arithmetic is easier!).

The Ryll-Nardzewski Theorem says that T is $\aleph_0$-categorical if and only if $|S_n(\varnothing)|$ is finite for all $n \in \omega$. Thus every $\aleph_0$-categorical theory is small (and indeed the unique countable model is saturated).

As a consequence of Exercise 9, we also have that every $\aleph_0$-stable theory is small. Neither of the converses is true: there are small $\aleph_0$-stable theories which are not $\aleph_0$-categorical and small $\aleph_0$-categorical theories which are not $\aleph_0$-stable. We will see some examples of these in Section 3.1.

Interlude: The monster model

For many purposes, it is convenient to work inside a "large" saturated model $\mathcal{U} \models T$ ($\mathcal{U}$ stands for "universe"). Taking "small" to mean "cardinality strictly less than $|\mathcal{U}|$", the following are the most important properties ensured by saturation:

(1) (Universality) Every complete type over a small subset of $\mathcal{U}$ is realized in $\mathcal{U}$. So we never have to move to elementary extensions to realize types.

(2) (Homogeneity) If B is a small set and $a, a' \in \mathcal{U}^x$ with $\mathrm{tp}(b/A) = \mathrm{tp}(b'/A)$, then there is an automorphism $\sigma \in \mathrm{Aut}(\mathcal{U}/B)$ such that $\sigma(a) = a'$. Here $\mathrm{Aut}(\mathcal{U}/B)$ is the group of all automorphisms of $\mathcal{U}$ which fix B pointwise.

Convention. (a) We fix a saturated model $\mathcal{U} \models T$ called the **monster model**. The cardinality of $\mathcal{U}$ is "as large as we need it to be" (see below).

(b) We call a set or model **small** if its cardinality is strictly less than $|\mathcal{U}|$. Otherwise it is **large**.

(c) From now on, "model" means "small model". The only large model we will refer to is $\mathcal{U}$. We assume every model is an elementary substructure of $\mathcal{U}$, and every small set is a subset of $\mathcal{U}$.

(d) A "formula with parameters", when we do not specify a set that the parameters come from, means a formula with arbitrary parameters from $\mathcal{U}$. When $\varphi(x)$ is a formula (possibly with parameters) and $a \in \mathcal{U}^x$, we write $\models \varphi(a)$ instead of $\mathcal{U} \models \varphi(a)$.

Convention (c) is justified by Corollary 2.7, which implies that every small model is isomorphic to an elementary substructure of $\mathcal{U}$. By Löwenheim–Skolem, every set is a subset of a small model, hence can be identified with a subset of $\mathcal{U}$. Further, in any particular situation (i.e., given any data involving families of models, sets, types, etc.), we may assume all models and all sets are small, simply by working in a larger monster model if necessary. This is what we mean by assuming that $|\mathcal{U}|$ is "as large as we need it to be".

What about existence of $\mathcal{U}$? By the previous paragraph, we want not just some saturated model of T, but arbitrarily large saturated models of T.

As we will see, when T is a stable theory (to be defined in Section 3.3 below), T is κ-stable for arbitrarily large cardinals κ, so T has arbitrarily large saturated models by Theorem 2.18. So for stable theories, the existence of $\mathcal{U}$ is well-justified (and most of the time, in this class, T is stable). In unstable theories, there are several options:

(1) We could work in a set theory extending ZFC, e.g. by assuming the existence of a proper class of inaccessible cardinals or assuming GCH.

(2) We could work in a class theory (like NBG) that admits proper classes as first-class objects. Here we can build a saturated monster models whose domain is a proper class.

(3) We could work with “big” models or “special” models instead of saturated models. These exist in ZFC and have the essential properties of universality and homogeneity (but “small” now means much smaller than the cardinality of $\mathcal{U}$). The construction of big and special models is a bit more technical than the construction of saturated models.

(4) We could view the monster model convention as “laziness” and check that every proof using the monster model could be unwound to a proof using small models only. Usually the unwound proofs will involve a lot more bookkeeping, passing to elementary extensions, etc.

Let us make one observation about the monster model before moving on. If $\varphi(x)$ is a formula (possibly with parameters), we call the set

$$\varphi(\mathcal{U}) = \{a \in \mathcal{U}^x \mid \models \varphi(a)\}$$

a **definable set**. In the monster model, every definable set is either finite or large. Indeed, suppose for contradiction that $\varphi(\mathcal{U})$ is a small infinite set. Then the partial type $\{\varphi(x)\} \cup \{x \neq a \mid a \in \varphi(\mathcal{U})\}$ is over a small set of parameters, and is consistent by compactness, hence realized in $\mathcal{U}$, contradiction. It is useful to keep this distinction in mind: we typically consider small sets of parameters, while definable sets are large (or finite).

3 Stability

3.1 Examples and non-examples by counting types

Example 3.1. Consider the theory T_∞ of infinite sets. This theory is complete and has quantifier elimination. For $M \models T_\infty$, we would like to count types in $S_1(M)$. To characterize types in the singleton variable x, it suffices to consider atomic formulas of the form $x = m$. If a type contains $x = m$, it is the realized type $\mathrm{tp}(m/M)$. The only other type is the unique non-realized type containing $x \neq m$ for all $m \in M$. Thus $|S_1(M)| = |M| + 1 = |M|$. It follows that T_∞ is κ-stable for all infinite cardinals κ.

Example 3.2. Consider the theory VS_k of non-trivial k-vector spaces. The language is $\{+, -, 0, (c)_{c\in k}\}$, where each symbol c is a unary function symbol for multiplication by $c \in k$. This theory is complete and has quantifier elimination.

For $M \models \mathrm{VS}_k$, we would like to count types in $S_1(M)$. Terms over M in the singleton variable x are k-linear combinations of x and the elements of M. An atomic formula is an equality between terms. Since M is closed under linear combinations, we find that every atomic formula is equivalent to $cx = m$, with $c \in k$ and $m \in M$. When $c = 0$, this is equivalent to $0 = m$, which does not depend on x. When $c \neq 0$, we find that the atomic formula is equivalent to $x = m'$ for some $m' \in M$.

Now the same analysis as in Example 3.1 applies. If a type contains $x = m$, it is the realized type $\mathrm{tp}(m/M)$. The only other type is the unique non-realized type containing $x \neq m$ for all $m \in M$. Thus $|S_1(M)| = |M| + 1 = |M|$. It follows that T_∞ is κ-stable for all infinite cardinals κ.

Exercise 11. Show that the theory ACF_0 (algebraically closed fields of characteristic 0) is κ-stable for all infinite cardinals κ. You may use the fact that ACF_0 is complete and has quantifier elimination.

Example 3.3. Consider $T = \mathrm{Th}(\mathbb{Z}, +, -, 0, (D_p)_{p\text{ prime}})$. This is the complete theory of the integers as an abelian group with additional unary relation symbols D_p, where $\mathbb{Z} \models D_p(x)$ if and only if $p \mid x$.

It is a fact that T has quantifier elimination in this language. The D_p are necessary for quantifier elimination: for example, without them, a formula of the form

$$\exists y\, (\underbrace{y + \cdots + y}_{p\text{ times}} = x)$$

is not equivalent to a quantifier-free formula.

I claim that T is κ-stable if and only if $\kappa \geq 2^{\aleph_0}$.

For any set Q of primes, the partial type

$$\{D_p(x) \mid p \in Q\} \cup \{\neg D_p(x) \mid p \notin Q\}$$

is consistent by compactness. So there are already $2^{\aleph_0}$-many pairwise contradictory partial types over $\varnothing$. It follows that for any model M, $|S_1(M)| \geq 2^{\aleph_0}$, and hence T is not κ-stable when $\kappa < 2^{\aleph_0}$.

On the other hand, let M be a model with $|M| \leq \kappa$, where $\kappa \geq 2^{\aleph_0}$. For every prime p, there is a definable equivalence relation on M defined by $x \equiv_p y$ if and only if $D_p(x-y)$. T asserts that the equivalence relation $\equiv_p$ has exactly p classes (in the standard model, these are the classes of $0, 1, \ldots, (p-1)$). So m has representatives for the p distinct $\equiv_p$ classes: call them $m_0^p, \ldots, m_{p-1}^p$.

I claim that a type $q \in S_1(M)$ is either a realized type (and there are at most κ-many of these) or it is completely determined by which $\equiv_p$-class the singleton variable x is in for each prime p. That is, by which of the formulas $D_p(x - m_i^p)$ is in q. The number of such choices is at most

$$\prod_{p \text{ prime}} p \leq \prod_{p \text{ prime}} \aleph_0 = \aleph_0^{\aleph_0} = 2^{\aleph_0} \leq \kappa.$$

Why is $\aleph_0^{\aleph_0} = 2^{\aleph_0}$? Because $2^{\aleph_0} \leq \aleph_0^{\aleph_0} \leq (2^{\aleph_0})^{\aleph_0} = 2^{\aleph_0 \cdot \aleph_0} = 2^{\aleph_0}$.

Let's prove the claim. To count types over M, it suffices to consider atomic and negated atomic formulas, by quantifier elimination. When x is a singleton, every atomic formula in context x is equivalent to $nx = m$ or $D_p(nx - m)$ (equivalently, $nx \equiv_p m$) with $n \in \mathbb{Z}$ and $m \in M$.

If a type $q(x) \in S_1(M)$ contains the formula $nx = m$ with $n \neq 0$, then $x \in M$, since T asserts that for all y and all $n \neq 0$, there is at most one x such that $nx = y$. So every non-realized type in $S_1(M)$ contains all the formulas $nx \neq m$ for $n \neq 0$.

Now let's say we know, for each prime p, which of the formulas $x \equiv_p m_i^p$ is in the non-realized type $q(x)$. If $x \equiv_p m_i^p$, then $nx \equiv_p nm_i^p$. Now in M, $nm_i^p \equiv_p m_j^p$ for some $0 \leq j \leq p-1$. So the formula $D_p(nx - m)$ is in q if and only if $m \equiv_p m_j^p$ in M. Thus $q(x)$ is uniquely determined by the $\equiv_p$-class of x for all primes p.

Example 3.4. Let $L = \{E\}$, and let T be the theory of an equivalence relation with infinitely many infinite classes. T is complete and has quantifier elimination. Let $M \models T$ with $|M| \leq \kappa$. By quantifier elimination, to count types over M we only need to consider atomic formulas $x = m$ and xEm with $m \in M$. There are three kinds of types in $S_1(M)$:

(1) The realized types, which contain $x = m$ for some $m \in M$. There are at most κ-many of these.

(2) Types which are not realized but contain xEm for some $m \in M$. Since there are at most κ-many equivalence classes in M, there are at most κ-many of these.

(3) Types which contain $\neg xEm$ for all $m \in M$. There is exactly one such type.

It follows that $|S_1(M)| \leq \kappa$, so T is κ-stable for all infinite cardinals κ.

Exercise 12. Let $L = \{E_n \mid n \in \omega\}$, and let T be the theory of binary refining equivalence relations. By this, I mean that for all $n \in \omega$, E_n has 2^n equivalence classes, each of which is infinite, and each E_n-class is partitioned into exactly 2 E_{n+1}-classes.

(a) Show that T has quantifier elimination.

(b) Show that T is consistent and complete.

(c) Show that T is κ-stable if and only if $\kappa \geq 2^{\aleph_0}$.

Example 3.5. Let $L = \{E_n \mid n \in \omega\}$, and let T be the theory of cross-cutting equivalence relations, each of which has infinitely many infinite classes. By "cross-cutting", I mean that for all $n \in \omega$, if we pick one E_i-class C_i for all $0 \leq i \leq n$, the intersection $\bigcap_{i=0}^{n} C_i$ is non-empty. This can be expressed by the following axioms, one for each $n \in \omega$:

$$\forall x_0 \ldots \forall x_n \exists y \left(\bigwedge_{i=0}^{n} y E_n x_i \right)$$

A canonical example of a model for T is $\mathrm{Fun}(\omega, \omega)$ where $f E_n g$ if and only if $f(n) = g(n)$.

T is complete and has quantifier elimination. Let $M \models T$ with $|M| \leq \kappa$. By quantifier elimination, to count types over M we only need to consider atomic formulas $x = m$ and $x E_n m$ with $m \in M$ and $n \in \omega$. We know there are at most κ-many realized types. To give a non-realized type in the singleton variable x, we need to specify, for each $n \in \omega$, whether x is in one of the E_n-classes in M (at most κ-many choices), or whether x is in a new E_n-class (just one choice). It follows that $|S_1(M)| \leq \kappa + (\kappa + 1)^{\aleph_0} = \kappa^{\aleph_0}$.

On the other hand, if we take any model M of size κ such that each E_n has κ-many equivalence classes (e.g. an elementary substructure of $\mathrm{Fun}(\omega, \kappa)$ of size κ containing a complete set of representatives for each E_n), then by compactness each of the types described above is consistent, and we obtain $|S_1(M)| = \kappa^{\aleph_0}$. It follows that M is κ-stable if and only if $\kappa^{\aleph_0} = \kappa$.

Note that $\aleph_0^{\aleph_0} = 2^{\aleph_0} > \aleph_0$. Stronger, if κ is a cardinal of cofinality $\aleph_0$, then $\kappa^{\aleph_0} = \kappa^{\mathrm{cf}(\kappa)} > \kappa$ by Theorem 1.23. On the other hand, lots of cardinals κ satisfy $\kappa^{\aleph_0} = \kappa$: For any cardinal λ, let $\kappa = \lambda^{\aleph_0}$. Then $\kappa^{\aleph_0} = (\lambda^{\aleph_0})^{\aleph_0} = \lambda^{\aleph_0 \cdot \aleph_0} = \lambda^{\aleph_0} = \kappa$.

Exercise 13. Show that the theory T from Example 3.5 has quantifier elimination.

Example 3.6. Consider the theory DLO (dense linear orders without endpoints). This theory is complete and has quantifier elimination. It is also $\aleph_0$-categorical, hence small (by the Ryll-Nardzewski Theorem). But it is not $\aleph_0$-stable: we have $|S_1(\mathbb{Q})| = 2^{\aleph_0}$. To see this, note that $\mathbb{Q} \preceq \mathbb{R}$ and distinct real numbers have distinct 1-types over $\mathbb{Q}$: For any $r < r'$ in $\mathbb{R}$, there is some $q \in \mathbb{Q}$ with $r < q < r'$. Then $(x < q) \in \mathrm{tp}(r/\mathbb{Q})$ but $\neg(x < q) \in \mathrm{tp}(r'/\mathbb{Q})$.

It will follow from the results of the next section that DLO is not κ-stable for any infinite κ.

Example 3.7. Consider the complete theory of the random graph, T_{RG}. This theory is complete and has quantifier elimination. It is also $\aleph_0$-categorical, hence small (by the Ryll-Nardzewski Theorem). But it is not κ-stable for any infinite κ: Let M be a model with $|M| = \kappa$. For any subset $A \subseteq M$, the partial type

$$\{(xRa) \mid a \in A\} \cup \{\neg(xRa) \mid a \notin A\}$$

is consistent, by compactness and the extension properties of the random graph. This gives 2^κ-many distinct types in $S_1(M)$.

Having seen lots of examples and non-examples of κ-stability, we can ask: Where does instability come from? Can we understand in a more concrete way the phenomena that causes some theories to have lots of types and others to have few types?

One observation is that in DLO and T_{RG}, it suffices to look at instances of a single formula $\varphi(x;y)$ (namely $x > y$ in DLO and xRy in T_{RG}) to find $2^{\aleph_0}$-many types over a countable set. In both cases, we can explain this by observing that we can build a "complete binary decision tree" for instances $\varphi(x;m)$ with $m \in M$. With only countably many parameters (corresponding to the nodes of the tree), we can obtain continuum-many types (corresponding to the paths through the tree). We will make this more precise in the next section.

On the other hand, in Examples 3.3 and 3.5, we can find complete decision trees (corresponding to picking equivalence classes for different equivalence relations), but we have to use instances of different formulas at each level of the tree. This allows us to show that these theories are not $\aleph_0$-stable, but they still manage to be κ-stable for larger cardinals κ. This suggests that the instability here is a *global* phenomenon, in the sense that it requires us to think about all formulas in the language. *Locally* (i.e., concentrating on one formula at a time), we just see individual equivalence relations, and the behavior of these theories looks much more like that of Example 3.4, where we cannot find any complete decision trees.

All this suggests that it we should develop a language for studying T locally: one formula at a time. This is what we do in the next section.

3.2 Stable formulas

A **partitioned formula** is a formula $\varphi(x;y)$, where the variable context has been partitioned into two finite disjoint tuples x and y. We call x the **object** variables and y the **parameter** variables.

We write $\varphi^{\mathrm{opp}}(y;x)$ for the same formula as $\varphi(x;y)$, but partitioned so that y is the tuple of object variables and x is the tuple of parameter variables.

Let $\varphi(x;y)$ be a partitioned formula. A **φ-formula** over a set B is a formula $\varphi(x;b)$ or $\neg\varphi(x;b)$ with $b \in B^y$. We write $F^\varphi_x(B)$ for the set of all φ-formulas over B.

A **partial φ-type** over B is a set of φ-formulas in $F^\varphi_x(B)$. A partial φ-type $\Sigma(x)$ over B is **consistent** if $\Sigma(x) \cup T_B$ is satisfiable. It is **complete** if it is consistent and $\varphi(x,b) \in \Sigma(x)$ or $\neg\varphi(x,b) \in \Sigma(x)$ for all $b \in B^y$. We write $S^\varphi_x(B)$ for the space of complete φ-types over B. We write $\mathrm{tp}^\varphi(a/B)$ for the complete φ-type of a over B.

Definition 3.8. Let κ be an infinite cardinal and $\varphi(x;y)$ a partitioned formula. We say $\varphi(x;y)$ is **κ-stable** (with respect to T) if $|S_x^\varphi(B)| \leq \kappa$ for all sets B with $|B| \leq \kappa$.

A **(set-theoretic) tree** is a partially ordered set $(\mathcal{T}, \leq)$ such that for any $t \in \mathcal{T}$, the set $\downarrow t = \{s \in \mathcal{T} \mid s \leq t\}$ is well-ordered by $\leq$.

For any ordinal α and any set X, we abuse notation by writing $X^{<\alpha}$ for the set $\{f\colon \beta \to X \mid \beta < \alpha\} = \bigcup_{\beta<\alpha} \mathrm{Fun}(\beta, X)$. The set $X^{<\alpha}$ has a natural tree structure given by the subset relation $f \subseteq g$ (equivalently, $\mathrm{dom}(f) \leq \mathrm{dom}(g)$ and $g|_{\mathrm{dom}(f)} = f$). For any $f \in X^{<\alpha}$ with $\mathrm{dom}(f) = \beta$, the set $\downarrow f$ consists of $\{f|_\gamma \mid \gamma < \beta\}$, so it is well-ordered with order-type β.

In addition to functions, we sometimes think of elements of $X^{<\alpha}$ as sequences from X (indexed by ordinals $< \alpha$) or visually as elements of a $|X|$-branching tree of height α. Given a node $f \in X^{<\alpha}$ with $\mathrm{dom}(f) = \beta$, and given $x \in X$, we sometimes write fx for the "child" of f defined by $f \cup \{(\beta, x)\}$.

We also write X^α for the set $\mathrm{Fun}(\alpha, X)$. In the context of trees, we think of an element $g \in X^\alpha$ as describing a path through the tree $X^{<\alpha}$, visiting the nodes $g|_\beta$ for all $\beta < \alpha$, or as a "leaf" of the tree.

Definition 3.9. For an ordinal α, a **binary tree of height** α for $\varphi(x;y)$ consists of two families of tuples $(a_f)_{f\in 2^\alpha} \in \mathcal{U}^x$ and $(b_g)_{g \in 2^{<\alpha}} \in \mathcal{U}^y$, such that for all $f \in 2^\alpha$ and all $\beta < \alpha$ we have

$$\models \varphi(a_f; b_{f|_\beta}) \text{ iff } f(\beta) = 1.$$

Note that the existence of a binary tree of height α for $\varphi(x;y)$ is expressed by the following partial type $\Gamma_\alpha((x_f)_{f\in 2^\alpha}, (y_g)_{g\in 2^{<\alpha}})$:

$$\{\varphi(x_f; y_{f|\beta}) \mid \beta < \alpha, f(\beta) = 1\} \cup \{\neg\varphi(x_f; y_{f|\beta}) \mid \beta < \alpha, f(\beta) = 0\}.$$

For finite n, Γ_n is finite, so we identity Γ_n with the single formula obtained by taking the conjunction of these finitely many formulas.

Proposition 3.10 (Binary tree property)**.** *For a partitioned formula $\varphi(x;y)$, the following are equivalent:*

(1) For all ordinals α, φ admits a binary tree of height α.

(2) φ admits a binary tree of height ω.

(3) For all $n < \omega$, φ admits a binary tree of height n.

(4) For all $n < \omega$, φ admits a binary tree of height n in every *model of T.*

Proof. (1) $\Rightarrow$ (2) is trivial.

(2) $\Rightarrow$ (3): Suppose $(a_f)_{f\in 2^\omega}$ and $(b_g)_{g\in 2^{<\omega}}$ is a binary tree of height ω for φ. Fix $n \in \omega$. For each $h \in 2^n$, pick some $h^* \in 2^\omega$ extending h. Then $(a_{h^*})_{h\in 2^n}$ and $(b_g)_{g\in 2^{<n}}$ is a binary tree of height n for φ.

(3) ⇒ (4): For any model $M \models T$, M has a binary tree for φ of height n if and only if

$$M \models \exists (x_f)_{f \in 2^n} \exists (y_g)_{g \in 2^{<n}} \, \Gamma_n((x_f)_{f \in 2^n}, (y_g)_{g \in 2^{<n}}).$$

Since T is complete, if $\mathcal{U}$ satisfies this sentence, then every model satisfies this sentence.

(4) ⇒ (1): It suffices to show that the partial type Γ_α defined above is consistent with T. A finite subset $\Delta \subseteq \Gamma_\alpha$ mentions only finitely many of the variables x_f and y_g. By choosing n large enough, we can realize Δ in a binary tree for φ of height n. By compactness, Γ_α consistent. □

Exercise 14. In the proof of (4) ⇒ (1) above, I wrote "by choosing n large enough, we can realize Δ in a binary tree for φ of height n." Fill in the details.

Definition 3.11. We say that $\varphi(x; y)$ has the **binary tree property** (with respect to T) if the equivalent conditions of Proposition 3.10 are satisfied.

It is easy to see that if φ has the binary tree property, then (using a binary tree of height ω), φ is not $\aleph_0$-stable. With a little cardinal arithmetic, we can show φ is not κ-stable for any infinite κ.

Proposition 3.12. *Suppose the partitioned formula $\varphi(x; y)$ has the binary tree property with respect to T. Then for every infinite cardinal κ, $\varphi(x; y)$ is* not *κ-stable.*

Proof. Suppose $\varphi(x; y)$ has the binary tree property, and let κ be an infinite cardinal. Let λ be the least cardinal such that $2^\lambda > \kappa$. Such a cardinal exists since the class of such cardinals is non-empty; in particular, it contains κ, so $\lambda \leq \kappa$. By definition of λ, $2^\mu \leq \kappa$ for all $\mu < \lambda$, so also:

$$|2^{<\lambda}| = \left| \bigcup_{\alpha < \lambda} \mathrm{Fun}(\alpha, 2) \right| \leq \max(\lambda, \sup_{\alpha < \lambda} 2^{|\alpha|}) \leq \kappa.$$

Since $\varphi(x; y)$ has the binary tree property, we can find $(a_g)_{g \in 2^\lambda} \in \mathcal{U}^x$ and $(b_f)_{f \in 2^{<\lambda}} \in \mathcal{U}^y$ such that for all $g \in 2^\lambda$ and all $\beta < \lambda$ we have

$$M \models \varphi(a_g; b_{g|_\beta}) \text{ iff } g(\beta) = 1.$$

Let B be the set of all elements of $\mathcal{U}$ appearing in the tuples b_f for $f \in 2^{<\lambda}$. Then $|B| \leq \kappa$.

I claim that for distinct $g, h \in 2^\lambda$, $\mathrm{tp}(a_g/B) \neq \mathrm{tp}(a_h/B)$. Indeed, let $\alpha < \lambda$ be the least ordinal such that $g(\alpha) \neq h(\alpha)$. Without loss of generality, $g(\alpha) = 1$ and $h(\alpha) = 0$. Let $\ell = g|_\alpha = h|_\alpha$. Then $\varphi(x; b_\ell) \in \mathrm{tp}(a_g/B)$, but $\neg\varphi(x; b_\ell) \in \mathrm{tp}(a_h/B)$.

This shows that $|S_1(B)| \geq 2^\lambda > \kappa$. So φ is not κ-stable. □

We would now like to go the other direction and show that if $\varphi(x;y)$ does not have the binary tree property, then φ is κ-stable for all infinite κ. In order to count complete φ-types over B, we note that $p \in S_x^\varphi(B)$ is completely determined by the set $Y_p = \{b \in B^y \mid \varphi(x;b) \in p\}$. In the worst case, the number of such sets Y_p is $|\mathcal{P}(B^y)| = 2^{|B|}$. But what if every set Y_p were itself defined by a formula over B? The number of subsets of B^y which are definable over B is at most $|F_y(B)| = |B|$ (when B is infinite). This leads to the notion of definability of types.

Definition 3.13. Let $\varphi(x;y)$ be a partitioned formula, and let A and B be sets. A complete φ-type $p \in S_x^\varphi(B)$ is **definable over** A if there is some formula $\psi(y) \in F_y(A)$ such that for all $b \in B^y$, $\varphi(x;b) \in p$ if and only if $\models \psi(b)$. If $p \in S_x^\varphi(B)$ is definable over B, we just say that p is **definable**.

Example 3.14. Let $T = \mathrm{DLO}$, and let $\varphi(x;y)$ be $x > y$. For each type $p \in S_x^\varphi(\mathbb{Q})$, the set $Y_p = \{b \in \mathbb{Q} \mid (x > b) \in p\}$ is a downwards-closed subset of $\mathbb{Q}$. By compactness, every downwards-closed set corresponds to a consistent φ-type. We can classify them as follows:

(1) For each $a \in \mathbb{Q}$, we have the realized type p_a, with

$$Y_{p_a} = \{b \in \mathbb{Q} \mid b < a\}.$$

This type is definable by $y < a$.

(2) For each $a \in \mathbb{Q}$, we have the type p_{a^+} "infinitesimally above" a, with

$$Y_{p_{a^+}} = \{b \in \mathbb{Q} \mid b \leq a\}.$$

This type is definable by $y \leq a$.

(3) We have the types "at $\pm\infty$", p_∞ and $p_{-\infty}$, with

$$Y_{p_\infty} = \mathbb{Q}$$
$$Y_{p_{-\infty}} = \varnothing.$$

These types are definable by $\top$ and $\bot$, respectively.

(4) For each $r \in \mathbb{R}$, we have the type p_r, with

$$Y_{p_r} = \{b \in \mathbb{Q} \mid b < r\}.$$

These types are not definable. By quantifier elimination, a definable subset of $\mathbb{Q}$ with parameters in $\mathbb{Q}$ is a finite boolean combination of singletons $\{a\}$ and intervals (a, ∞) with $a \in \mathbb{Q}$. The set $(-\infty, r) \cap \mathbb{Q}$ cannot be written as such a boolean combination.

In order to show that types are definable, it is useful to introduce Shelah's local 2-rank. This is just one of many notions of rank used in model theory (we will see some more later in this course). Many notions of rank measure how much a definable set can be "split" into smaller definable pieces. In this case, the name local comes from the fact that we only use φ-formulas to do the "splitting", and the number 2 refers to the fact that the rank goes up when we can split a definable set into two pieces. The connection with the binary tree property is that repeated splitting a definable set into two pieces by φ-formulas is equivalent to building a binary tree for φ (see Exercise 15 below).

Definition 3.15. Let $\varphi(x; y)$ be a partitioned formula. For any formula with parameters $\theta(x)$, we define the **local 2-rank** of θ, $R_2^{\varphi}(\theta)$, recursively as follows:

(1) $R_2^{\varphi}(\theta) \geq 0$ if and only if θ is satisfiable.

(2) $R_2^{\varphi}(\theta) \geq n+1$ if and only if there is some $b \in \mathcal{U}^y$ such that

$$R_2^{\varphi}(\theta \wedge \varphi(x; b)) \geq n \text{ and } R_2^{\varphi}(\theta \wedge \neg\varphi(x; b)) \geq n.$$

If θ is not satisfiable, we set $R_2^{\varphi}(\theta) = -\infty$. If $R_2^{\varphi}(\theta) \geq n$ for all n, we set $R_2^{\varphi}(\theta) = \infty$. Otherwise, we set $R_2^{\varphi}(\theta)$ to be the maximal $n \in \omega$ such that $R_2^{\varphi}(\theta) \geq n$.

Exercise 15. Show that we can equivalently define the local 2-rank as follows: $R_2^{\varphi}(\theta) \geq n$ if and only if there is a binary tree for $\varphi(x; y)$ of height n, consisting of tuples $(a_f)_{f \in 2^n}$ and $(b_g)_{g \in 2^{<n}}$ such that $\models \theta(a_f)$ for all $f \in 2^n$.

Proposition 3.16. *Suppose φ does not have the binary tree property. Then for any set B, every type in $S_x^{\varphi}(B)$ is definable.*

Proof. Since φ does not have the binary tree property, there is some $N \in \omega$ such that φ does not admit a binary tree of height N. It follows from Exercise 15 that $R_2^{\varphi}(\theta(x)) < N$ for all $\theta(x)$.

Let $p \in S_x^{\varphi}(B)$. Define

$$\Theta = \left\{ \bigwedge_{i=1}^{n} \psi_i \mid n \in \omega, \psi_1, \ldots, \psi_n \in p \right\},$$

So Θ is the set of all finite conjunctions of φ-formulas in p. Choose a formula $\theta(x) \in \Theta$ of minimal R_2^{φ}-rank m.

For any $b \in B^y$, I claim that $\varphi(x; b) \in p$ if and only if $R_2^{\varphi}(\theta(x) \wedge \varphi(x; b)) \geq m$.

If $\varphi(x; b) \in p$, then $\theta(x) \wedge \varphi(x; b) \in \Theta$, so $R_2^{\varphi}(\theta(x) \wedge \varphi(x; b)) \geq m$ by minimality of m. On the other hand, if $\varphi(x; b) \notin p$, suppose for contradiction that $R_2^{\varphi}(\theta(x) \wedge \varphi(x; b)) \geq m$. Then $\theta(x) \wedge \neg\varphi(x; b) \in \Theta$, so $R_2^{\varphi}(\theta(x) \wedge \neg\varphi(x; b)) \geq m$ by minimality of m. Then by definition of the local 2-rank, $R_2^{\varphi}(\theta(x)) \geq m+1$, contradiction.

It remains to show that the condition on b that $R_2^{\varphi}(\theta(x) \wedge \varphi(x; b)) \geq m$ is definable. Consider the formula $\psi(y)$:

$$\exists (x_f)_{f \in 2^m} \exists (y_g)_{g \in 2^{<m}} \left(\Gamma_m((x_f)_{f \in 2^m}, (y_g)_{g \in 2^{<m}}) \wedge \bigwedge_{f \in 2^m} \theta(x_f) \wedge \varphi(x_f, y) \right).$$

The formula $\psi(y)$ expresses that there is a binary tree for φ of height m such that all the leaves $(a_f)_{f\in 2^m}$ satisfy $\theta(x)\wedge\varphi(x;y)$. By Exercise 15, this is equivalent to $R_2^\varphi(\theta(x)\wedge\varphi(x;y))\geq m$. The only parameters used in $\psi(y)$ are those parameters from B appearing in $\theta(x)$. So $\psi(y)$ defines p over B. $\square$

Theorem 3.17 (Stable formula theorem)**.** *Let $\varphi(x;y)$ be a partitioned formula. The following are equivalent:*

(1) φ does not have the binary tree property.

(2) φ is κ-stable for some infinite cardinal κ.

(3) φ is κ-stable for all infinite cardinals κ.

(4) Every complete φ-type over any set B is definable.

Proof. (3)⇒(2): Trivial.

(2)⇒(1): This is the contrapositive of Proposition 3.12.

(1)⇒(4): This is Proposition 3.16.

(4)⇒(3): Let κ be an infinite cardinal and B a set with $|B|\leq\kappa$. Every type $p\in S_x^\varphi(B)$ is definable over B by a formula in $F_y(B)$. Since $|F_y(B)|\leq\kappa$, $|S_x^\varphi(B)|\leq\kappa$. $\square$

Definition 3.18. We say that $\varphi(x;y)$ is **stable** (relative to T) if the equivalent conditions of Theorem 3.17 are satisfied.

3.3 Stable theories

Definition 3.19. T is **stable** if every partitioned formula $\varphi(x;y)$ is stable relative to T.

Theorem 3.20. *The following are equivalent:*

(1) T is stable.

(2) T is κ-stable for some infinite cardinal κ.

(3) T is κ-stable for all infinite cardinals κ such that $\kappa^{\aleph_0}=\kappa$.

Proof. (1)⇒(3): Let κ be such that $\kappa^{\aleph_0}=\kappa$. Let A be a set with $|A|\leq\kappa$. The map $S_x(A)\to\prod_{\varphi(x;y)}S_x^\varphi(A)$ given by $p\mapsto(p|_\varphi)_{\varphi(x;y)}$ is injective, since if $p\neq q$, then they differ on some formula over A. It suffices to consider parameter tuples $y=(y_1,\ldots,y_n)$ with $n\in\omega$, so the number of partitioned formulas $\varphi(x;y)$ in the product is $\aleph_0$. Each such formula is κ-stable, so by Theorem 3.17:

$$|S_x(A)|\leq\prod_{\varphi(x;y)}|S_x^\varphi(A)|\leq\prod_{\varphi(x;y)}\kappa=\kappa^{\aleph_0}=\kappa.$$

(3)⇒(2): For any cardinal $\lambda\geq 2$, let $\kappa=\lambda^{\aleph_0}$. Then κ is infinite and we have:

$$\kappa^{\aleph_0}=(\lambda^{\aleph_0})^{\aleph_0}=\lambda^{\aleph_0\cdot\aleph_0}=\lambda^{\aleph_0}=\kappa.$$

By (3), T is κ-stable.

(2)$\Rightarrow$(1): Let $\varphi(x;y)$ be a partitioned formula. Let κ be an infinite cardinal such that T is κ-stable, and let A be a set with $|A| \leq \kappa$. The restriction map $S_x(A) \to S_x^\varphi(A)$ is surjective, so we have $|S_x^\varphi(A)| \leq |S_x(A)| \leq \kappa$. Since A was arbitrary, $\varphi(x;y)$ is κ-stable, hence stable by Theorem 3.17. $\square$

The amazing thing about Theorem 3.20 is that conditions (2) and (3), which are about counting types over possibly uncountable parameter sets, and which seem like they could depend on facts about cardinal arithmetic, and hence be independent of ZFC, turn out to be equivalent to condition (1), which is in turn equivalent to the condition that no formula $\varphi(x;y)$ has the binary tree property relative to T. The binary tree property can be checked by asking whether T proves that $\varphi(x;y)$ has binary trees of every finite height. This is a pattern in stability theory: looking at phenomena that occur in large uncountable models can lead us to identify very concrete combinatorial equivalents of these phenomena, which end up being useful dividing lines in the complexity of theories.

The proof of (3)$\Rightarrow$(2) in Theorem 3.20 shows more: For all λ, we have $\lambda \leq \lambda^{\aleph_0}$, so if T is stable, then there are arbitrarily large cardinals κ such that T is κ-stable.

Corollary 3.21. *Suppose T is stable. Then every model of T has a saturated elementary extension.*

Proof. Suppose $M \models T$ and let $\lambda = |M|$. Then there is a cardinal $\kappa \geq \lambda$ with $\kappa^{\aleph_0} = \kappa$ (e.g. $\kappa = \lambda^{\aleph_0}$ works). Since T is κ-stable, T has a saturated model M' of cardinality κ by Theorem 2.18. By Corollary 2.7, there is an elementary embedding $M \to M'$. $\square$

For a set A (large or small), a **definable subset** of A^x is a set of the form $\varphi(A) = \{a \in A^x \mid \mathcal{U} \models \varphi(a)\}$ for some formula φ with parameters.

Definition 3.22. A set A (large or small) is **stably embedded** if every definable subset of A^x (for every context x) is definable over A. That is, for every formula $\varphi(x;b)$ with parameters $b \in \mathcal{U}^y$, there is a formula $\psi(x)$ over A such that $\varphi(A,b) = \psi(A)$.

Corollary 3.23. *Suppose T is stable. Then every set is stably embedded.*

Proof. Let A be a set and $\varphi(x;b)$ a formula with parameters $b \in \mathcal{U}^y$. Consider $p = \mathrm{tp}^{\varphi^{\mathrm{opp}}}(b/A)$. Since T is stable, $\varphi^{\mathrm{opp}}(y;x)$ is a stable formula, so p is definable over A. That is, there is a formula $\psi(x)$ over A such that $\psi(A) = \varphi^{\mathrm{opp}}(b;A) = \varphi(A;b)$. $\square$

Example 3.24. Recall that ACF_0 is a stable theory. Suppose we are interested in rational points $(a_1,\ldots,a_n) \in \mathbb{Q}^n$ such that $f(a_1,\ldots,a_n) = 0$, where $f(x_1,\ldots,x_n) \in \mathbb{C}[x_1,\ldots,x_n]$ is a polynomial in n variables over the complex numbers. Since $\mathbb{Q}$ is stably embedded in $\mathbb{C}$, there is a formula $\varphi(x_1,\ldots,x_n)$ over $\mathbb{Q}$ (which by quantifier elimination is a boolean combination of polynomial equations over $\mathbb{Q}$) defining (in $\mathbb{C}$) the set of rational solutions to $f(x_1,\ldots,x_n) = 0$.

4 The order property and indiscernibles

In this section, we will explore some additional equivalents to stability and introduce an important tool: indiscernibles. The section will culminate with our first result about κ-categorical theories for uncountable κ. We begin with a proof a Ramsey's Theorem, since many of the results of this section rely on this result of combinatorial set theory.

Given a set X and $k \in \omega$, we write $[X]^k$ for the set of subsets of X of size exactly k. Ramsey's Theorem is a statement about functions $c\colon [X]^k \to r$ with $r \in \omega$. We describe such a function as a "coloring" of each set in $[X]^k$ by one of r colors. A subset $H \subseteq X$ is **homogeneous** for c if $c|_{[H]^k}$ is a constant function, i.e., every set in $[H]^k$ gets the same color.

Theorem 4.1 (Ramsey)**.** *Let X be an infinite set. For any $r, k \in \omega$ and any function $c\colon [X]^k \to r$, there is an infinite $H \subseteq X$ which is homogeneous for c.*

Proof. We proceed by induction on k. The base case $k = 1$ is the "infinitary pigeonhole principle": if we partition an infinite set into finitely many pieces, one of the pieces must be infinite (since a finite union of finite sets is finite). So given $c\colon X \to r$, we can take H to be one of the sets $c^{-1}(\{i\})$ for $i \in r$.

For the inductive step, we are given $c\colon [X]^{k+1} \to r$. We define a chain of infinite sets $X_0 \supset X_1 \supset X_2 \supset \dots$ and elements $x_i \in X_i$ by recursion. Let $X_0 = X$. Given X_i, pick any $x_i \in X_i$ and define a function $c_i\colon [X_i \setminus \{x_i\}]^k \to r$ by $c_i(Z) = c(Z \cup \{x_i\})$. By the inductive hypothesis, there is an infinite set $X_{i+1} \subseteq X_i \setminus \{x_i\}$ which is homogeneous for c_i.

Let $Y = \{x_i \mid i \in \omega\}$. For each $i \in \omega$, let $k_i \in r$ be the constant value of c_i on X_{i+1}. Let $d\colon Y \to r$ be the function $d(x_i) = k_i$. By the infinitary pigeonhole principle, there is some infinite $H \subseteq Y$ which is homogeneous for d. That is, there is some $k^* \in r$ such that $k_i = k^*$ whenever $x_i \in H$. Now for any $Z \in [H]^k$, we can write $Z = \{x_{i_1}, \dots, x_{i_{k+1}}\}$ with $i_1 < \dots < i_{k+1}$. Since $Z\setminus\{x_{i_1}\} \subseteq X_{i_1+1} \subset X_{i_1}\setminus\{x_{i_1}\}$, we have $c(Z) = c_{i_1}(\{x_{i_2}, \dots, x_{i_{k+1}}\}) = k_{i_1} = k^*$. So H is homogeneous for c. $\square$

Exercise 16. Let $(X, \leq)$ be an infinite linearly ordered set. Show that there is either an infinite increasing sequence $x_0 < x_1 < x_2 < \dots$ or an infinite decreasing sequence $x_0 > x_1 > x_2 > \dots$ in X.

Hint: Choose a well-ordering $\preceq$ of X. For each pair $\{x, y\} \in [X]^2$ color $\{x, y\}$ according to whether $\preceq$ agrees with $\leq$. Apply Ramsey's Theorem.

4.1 The order property

Definition 4.2. Let $\varphi(x; y)$ be a partitioned formula, and let $(I, \leq)$ be a linear order. We say that $\varphi(x; y)$ has the **order property** indexed by I (with respect to T) if there exist $(a_i)_{i\in I} \in \mathcal{U}^x$ and $(b_j)_{j\in I} \in \mathcal{U}^y$ such that for all $i, j \in I$, we have:

$$\models \varphi(a_i; b_j) \text{ iff } i \leq j.$$

We say that $\varphi(x;y)$ has the **order property** (with respect to T) if it has the order property indexed by ω.

Example 4.3. Let T be DLO, and let $\varphi(x;y)$ be the formula $x \leq y$. Let $(a_i)_{i\in\omega}$ be any increasing sequence in $\mathcal{U}$. Then taking $b_i = a_i$ for all i, $(a_i)_{i\in\omega}$ and $(b_i)_{i\in\omega}$ witness the order property for φ.

Example 4.4. Let T be T_{RG}, and let $\varphi(x;y)$ be the formula xRy. Let $(a_i)_{i\in\omega}$ be any sequence of distinct elements in $\mathcal{U}$. For all $j \in \omega$, the partial type $\{a_iRy \mid i \leq j\} \cup \{\neg a_iRy \mid i > j\}$ is consistent by compactness and the extension property for random graphs. Let b_j realize this partial type. Then $(a_i)_{i\in\omega}$ and $(b_i)_{i\in\omega}$ witness the order property for φ.

Note that for any set $S \subseteq \omega$, the partial type $\{a_iRy \mid i \in S\} \cup \{\neg a_iRy \mid i \notin S\}$ is consistent, by the same argument. Thus we can find sequences $(a_i)_{i\in\omega}$ and $(b_S)_{S\subseteq\omega}$ witnessing a stronger property of φ, the **independence property**:

$$\models \varphi(a_i; b_S) \text{ iff } i \in S.$$

The order property is the special case of the independence property where we only consider sets S of the form $\{i \in \omega \mid i \leq j\}$.

Proposition 4.5 (Order property)**.** *For a partitioned formula $\varphi(x;y)$, the following are equivalent:*

(1) For every linear order $(I, \leq)$, φ has the order property indexed by I.

(2) φ has the order property (indexed by ω).

(3) For all $n < \omega$, φ has the order property indexed by n.

(4) For all $n < \omega$, φ has the order property indexed by n in every *model of T.*

The proof is just like the proof of Proposition 3.10.

Proposition 4.6. *Let $\varphi(x;y)$ be a partitioned formula.*

(1) In Definition 4.2, it is equivalent to require $\models \varphi(a_i; b_j)$ iff $i < j$.

(2) φ has the order property if and only if $\neg\varphi(x;y)$ has the order property.

(3) φ has the order property if and only if φ^{opp} has the order property.

(4) If $(\varphi_1 \vee \varphi_2)(x;y)$ has the order property, then $\varphi_1(x;y)$ has the order property or $\varphi_2(x;y)$ has the order property.

Proof. (1) If $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ satisfy $\models \varphi(a_i; b_j)$ iff $i \leq j$, then $(a_{i+1})_{i\in\omega}$ and $(b_j)_{j\in\omega}$ satisfy $\models \varphi(a_{i+1}; b_j)$ iff $i < j$.

Similarly, if $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ satisfy $\models \varphi(a_i; b_j)$ iff $i < j$, then $(a_i)_{i\in\omega}$ and $(b_{j+1})_{j\in\omega}$ satisfy $\models \varphi(a_i; b_{j+1})$ iff $i \leq j$.

(2) If $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ satisfy $\models \varphi(a_i; b_j)$ iff $i \leq j$, then $\models \neg\varphi(a_i; b_j)$ iff $i > j$. Thus $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ witness the order property for $\neg\varphi$ indexed by ω^* (ω with the reverse of the usual order) in the equivalent strict form of (1).

For the converse, observe that $\neg\neg\varphi \equiv \varphi$.

(3) If $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ satisfy $\models \varphi(a_i; b_j)$ iff $i \leq j$, then $\models \varphi^{\text{opp}}(b_j; a_i)$ iff $j \geq i$. Thus $(b_j)_{j\in\omega}$ and $(a_i)_{i\in\omega}$ witness the order property for $\varphi^{\text{opp}}(y; x)$ indexed by ω^*.

For the converse, observe that $(\varphi^{\text{opp}})^{\text{opp}} = \varphi$.

(4) Suppose $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ satisfy $\models (\varphi_1 \vee \varphi_2)(a_i; b_j)$ iff $i < j$. By Ramsey's Theorem (coloring $[\omega]^2$ by 2 colors, depending on whether or not φ_1 holds), there is an infinite subset $H \subseteq \omega$ and some $k \in \{1, 2\}$ such that for all $i < j$ in H, $\models \varphi_k(a_i, b_j)$. On the other hand, for all $i \geq j$ in H, $\not\models (\varphi_1 \vee \varphi_2)(a_i; b_j)$, so $\not\models \varphi_k(a_i; b_j)$. Thus $(a_i)_{i\in H}$ and $(b_j)_{j\in H}$ witness the order property for φ_k. $\square$

Theorem 4.7 (Stable formula theorem, continued)**.** *Let $\varphi(x; y)$ be a partitioned formula. The following are equivalent:*

(1) φ is stable.

(2) φ does not have the order property.

(3) For every model $M \models T$, every φ-type in $S^{\varphi}_x(M)$ is definable by a formula $\psi(y)$ which is a boolean combination of φ^{opp}-formulas.

Proof. (1)$\Rightarrow$(2): Assuming φ has the order property, we show that φ is unstable. By Proposition 4.5, we can find $(a_r)_{r\in\mathbb{R}}$ and $(b_s)_{s\in\mathbb{R}}$ such that $\models \varphi(a_r, b_s)$ if and only if $r \leq s$. Let B be the set of elements appearing in the tuples $(b_s)_{s\in\mathbb{Q}}$. Then B is countable, but when $r \neq r'$ in $\mathbb{R} \setminus \mathbb{Q}$, $\text{tp}^{\varphi}(a_r/B) \neq \text{tp}^{\varphi}(a_{r'}/B)$. So $|S^{\varphi}_x(B)| \geq 2^{\aleph_0}$, and φ is unstable.

(2)$\Rightarrow$(3): Assume that there is a model $M \models T$ and a type $p \in S^{\varphi}_x(M)$ which is not definable over M by any boolean combination of φ^{opp}-formulas. We show that φ has the order property. Let $a^* \in \mathcal{U}^x$ be a realization of p.

Note that for any finite subset $B \subseteq M$, there are only finitely many tuples in B^y, so a φ-type $q \in S^{\varphi}_x(B)$ is a finite set of formulas and can be identified with its conjunction $\theta_q(x)$. Now since p is consistent, $\mathcal{U} \models \exists x\, \theta_q(x)$, so $M \models \exists x\, \theta_q(x)$, and hence q is realized in M. This is where we use the assumption that M is a model.

I claim that for any finite sequence $a_0, \ldots, a_{n-1} \in M^x$, there are $b, b' \in M^y$ such that $\models \varphi(a^*, b)$, $\models \neg\varphi(a^*, b')$, and $\models \varphi(a_i, b) \leftrightarrow \varphi(a_i, b')$ for all $i < n$. If not, then for every set $S \subseteq n$, either all $b \in M^y$ satisfying

$$\bigwedge_{i\in S} \varphi(a_i, y) \wedge \bigwedge_{i\notin S} \neg\varphi(a_i, y)$$

satisfy $\varphi(a^*, y)$ (in which case we call S good), or all such b satisfy $\neg\varphi(a^*, y)$. Then $p = \mathrm{tp}^\varphi(a^*/M)$ is definable over M by

$$\bigvee_{S \text{ good}} \left(\bigwedge_{i \in S} \varphi(a_i, y) \wedge \bigwedge_{i \notin S} \neg\varphi(a_i, y) \right).$$

This is a boolean combination of φ^{opp}-formulas, contradicting the choice of p.

We now construct sequences $(b_i)_{i\in\omega}$ and $(b'_i)_{i\in\omega}$ in M^y, and $(a_i)_{i\in\omega}$ in M^x, by recursion. Suppose we are given $(b_i)_{i<n}$, $(b'_i)_{i<n}$, and $(a_i)_{i<n}$. By the claim, there are $b_n, b'_n \in M^y$ such that $\models \varphi(a^*, b_n)$ and $\models \neg\varphi(a^*, b'_n)$ and $\models \varphi(a_i, b_n) \leftrightarrow \varphi(a_i, b'_n)$ for all $i < n$. Let $a_n \in M^x$ realize $\mathrm{tp}^\varphi(a^*/b_0b'_0 \ldots b_nb'_n)$.

The result of our construction is that that whenever $i \geq j$, we have

$$\models \varphi(a_i, b_j) \text{ and } \models \neg\varphi(a_i, b'_j),$$

but whenever $i < j$, we have

$$\models \varphi(a_i, b_j) \text{ if and only if } \models \varphi(a_i, b'_j).$$

By Ramsey's Theorem (coloring pairs from ω by 2 colors), there is an infinite set $H \subseteq \omega$ such that either:

(1) For all $i < j \in H$, $\models \varphi(a_i, b_j)$ and $\models \varphi(a_i, b'_j)$. In this case, the sequences $(a_i)_{i\in H}$ and $(b'_i)_{i\in H}$ witness the order property for φ.

(2) For all $i < j \in H$, $\models \neg\varphi(a_i, b_j)$ and $\models \neg\varphi(a_i, b'_j)$. In this case, the sequences $(a_i)_{i\in H}$ and $(b_i)_{i\in H}$ witness the order property for φ.

(3)$\Rightarrow$(1): We already know that stability is equivalent to the condition that every complete φ-type over a set B is definable. We only need to show that it suffices to check this over models.

Let B be an arbitrary set with $|B| \leq \kappa$. Let M be a model containing B with $|M| \leq \kappa$. Since every type in $S^\varphi_x(M)$ is definable over M, we have

$$|S^\varphi_x(B)| \leq |S^\varphi_x(M)| \leq |F^\varphi_x(M)| = |M| \leq \kappa.$$

Thus φ is stable. $\square$

With a very slightly more complicated argument, Theorem 4.7(3) can be improved to say that every complete φ-type over a model is definable by a *positive* Boolean combination of φ^{opp}-formulas (not using $\neg$).

One consequence of the order property characterization of stability is that it is easier to prove closure properties of the class of stable formulas using Proposition 4.6.

Corollary 4.8. *If φ is stable, so is φ^{opp}. Every boolean combination of stable formulas is stable.*

Proof. The first statement is a direct application of Proposition 4.6(3). For the second statement, note that stability is unaffected by adding new variables to the context of a formula, so we may assume that all formulas in our boolean combination have the same partitioned context $(x; y)$. Now by DeMorgan's Law $\varphi \wedge \psi \equiv \neg(\neg\varphi \vee \neg\varphi)$, we may assume that our boolean combination is built up using only negations and disjunctions. By Proposition 4.6(2), the negation of a stable formula is stable, and by Proposition 4.6(4), the disjunction of two stable formulas (in the same context) is stable. $\square$

4.2 Indiscernibles

Definition 4.9. Let $(I, \leq)$ be a linear order, let $\mathcal{I} = (a_i)_{i\in I}$ be a sequence from $\mathcal{U}^x$ indexed by I, and let B be a set. We say that $\mathcal{I}$ is an sequence of **order indiscernibles** over B (or a B**-indiscernible sequence**) if for all $n \in \omega$, all $i_1 < \cdots < i_n$ and $j_1 < \ldots j_n$ in I, and all L_B-formulas $\varphi(x_1, \ldots, x_n)$ (where each x_i is a tuple of length $|x|$), we have

$$\models \varphi(a_{i_1}, \ldots, a_{i_n}) \text{ iff } \models \varphi(a_{j_1}, \ldots, a_{j_n}).$$

We say that $\mathcal{I}$ is a sequence of **set indiscernibles** over B (or a B**-indiscernible set**) if the same is true for any distinct $i_1, \ldots, i_n$ and distinct $j_1, \ldots, j_n$ in I (not necessarily in increasing order).

When no set B is mentioned, we assume $B = \varnothing$.

Example 4.10. Let $T = \mathrm{DLO}$. By quantifier elimination, every strictly increasing sequence $(a_i)_{i\in\omega}$ of elements of $\mathcal{U}$ is an indiscernible sequence (over $\varnothing$). Such a sequence is not an indiscernible set: letting $\varphi(x_1, x_2)$ be $x_1 < x_2$, we have $\models \varphi(a_0, a_1)$ but $\models \neg\varphi(a_1, a_0)$.

On the other hand, with $T = \mathrm{Th}(\mathbb{Z}; <)$ the strictly increasing sequence $0, 1, 2, \ldots$ is not indiscernible: letting $\varphi(x_1, x_2)$ be $\exists y\,(x_1 < y \wedge y < x_2)$, we have $\models \neg\varphi(0, 1)$, but $\models \varphi(0, 2)$.

Example 4.11. Let T be the theory of an equivalence relation with infinitely many infinite classes. The following are all indiscernible sets:

- Any constant sequence $(a_i)_{i\in\omega}$, with $a_i = a$ for some $a \in \mathcal{U}$.
- Any sequence of inequivalent elements $(a_i)_{i\in\omega}$ with $\neg(a_i E a_j)$ for all $i \neq j$.
- Any sequence of distinct equivalent elements $(a_i)_{i\in\omega}$ with $(a_i E a_j)$ and $a_i \neq a_j$ for all $i \neq j$.

Exercise 17. With T as in Example 4.11, show that every indiscernible sequence in $\mathcal{U}^1$ fits into one of those three cases. Now give a similar classification of indiscernible sequences in $\mathcal{U}^n$ for arbitrary n.

Example 4.12. In a vector space, any sequence of linearly independent elements is an indiscernible set.

Exercise 18. With $T = \mathrm{DLO}$, show that there are no non-constant indiscernible sets. That is, if $\mathcal{I} = (a_i)_{i\in\omega}$ in $\mathcal{U}^x$ is an indiscernible set (where x is an arbitrary finite context), then $\mathcal{I}$ is constant.

Exercise 19. In the random graph, find:

(a) A non-constant indiscernible set.

(b) An indiscernible sequence which is not an indiscernible set.

Definition 4.13. Let $(I, \leq)$ be an infinite linear order, let $\mathcal{I} = (a_i)_{i\in I}$ be an I-indexed sequence from $\mathcal{U}^x$ (not necessarily indiscernible), and let B be a set. The **Ehrenfeucht–Mostowski type** of $\mathcal{I}$ over B is a set of L_B-formulas in contexts $x_1, \ldots, x_n$, where $n \in \omega$ and each x_i is a tuple of length $|x|$:

$$\mathrm{EM}(\mathcal{I}/B) = \{\varphi(x_1, \ldots, x_n) \mid \text{for all } i_1 < \cdots < i_n \in I, \models \varphi(a_{i_1}, \ldots, a_{i_n})\}.$$

A sequence $\mathcal{J} = (a'_i)_{j\in J}$ from $\mathcal{U}^x$, indexed by a linear order $(J, \leq)$, **satisfies** $\mathrm{EM}(\mathcal{I}/B)$ if for all $\varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B)$, we have $\models \varphi(a'_{j_1}, \ldots, a'_{j_n})$ for all $j_1 < \cdots < j_n \in J$. We also say that $\mathcal{J}$ is **locally based** on $\mathcal{I}$ over B.

The sequence $\mathcal{I}$ is indiscernible over B if and only if $\mathrm{EM}(\mathcal{I}/B)$ is complete in the sense that for any L_B-formula $\varphi(x_1, \ldots, x_n)$, either $\varphi \in \mathrm{EM}(\mathcal{I}/B)$ or $\neg\varphi \in \mathrm{EM}(\mathcal{I}/B)$.

Using Ramsey's Theorem, it is always possible to take a sequence in a model of T and find an indiscernible sequence locally based on it.

Lemma 4.14 ("Standard Lemma")**.** *Let $\mathcal{I} = (a_i)_{i\in\omega}$ be a sequence from $\mathcal{U}^x$, and let B be a set. Then there is a B-indiscernible sequence $\mathcal{J} = (c_j)_{j\in\omega}$ satisfying* $\mathrm{EM}(\mathcal{I}/B)$.

Proof. Consider the partial type q (in context $(y_j)_{j\in\omega}$, where each y_j is a tuple of length $|x|$) consisting of formulas:

(a) $\varphi(y_{j_1}, \ldots, y_{j_n})$, where $j_1 < \cdots < j_n$ in ω and $\varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B)$.

(b) $\varphi(y_{j_1}, \ldots, y_{j_n}) \leftrightarrow \varphi(y_{j'_1}, \ldots, y_{j'_n})$, where $j_1 < \cdots < j_n$ and $j'_1 < \cdots < j'_n$ in ω and $\varphi(x_1, \ldots, x_n)$ is an L_B-formula.

It suffices to show that q is consistent, since the formulas of type (a) ensure that $\mathcal{J}$ satisfies $\mathrm{EM}(\mathcal{I}/B)$ and the formulas of type (b) ensure that $\mathcal{J}$ is indiscernible over B.

For any finite set of L_B-formulas Δ, let q_Δ be the same partial type, but with the formulas of type (b) restricted to those L_B-formulas appearing in Δ. A finite subset of q is contained in q_Δ for some finite set Δ, so by compactness it suffices to show that q_Δ is consistent. Further, by adding dummy variables, we may assume that each formula in Δ has the same context $x_1, \ldots, x_n$, where each x_i is a tuple of length $|x|$. Indeed, if $\psi(x_1, \ldots, x_m)$ is an L_B-formula with $m < n$, let $\psi'(x_1, \ldots, x_n)$ be the same formula with dummy variables $x_{m+1}, \ldots, x_n$ added.

Now for any $j_1 < \cdots < j_m$ and $j'_1 < \cdots < j'_m$ in ω, we can extend both sequences to $j_1 < \cdots < j_n$ and $j'_1 < \cdots < j'_n$ in ω, and (b) for ψ' implies

$$\psi(y_{j_1}, \ldots, y_{j_m}) \leftrightarrow \psi'(y_{j_1}, \ldots, y_{j_n}) \leftrightarrow \psi'(y_{j'_1}, \ldots, y_{j'_n}) \leftrightarrow \psi(y_{j'_1}, \ldots, y_{j'_m}).$$

Define a coloring $c\colon [\omega]^n \to \mathcal{P}(\Delta)$ by

$$c(\{i_1, \ldots, i_n\}) = \{\varphi(x_1, \ldots, x_n) \in \Delta \mid \models \varphi(a_{i_1}, \ldots, a_{i_n})\},$$

where we always enumerate $\{i_1, \ldots, i_n\}$ so that $i_1 < \cdots < i_n$.

By Ramsey's Theorem, there is an infinite subset $H \subseteq \omega$ which is homogeneous for c. Enumerate H in increasing order as $(a_{i_j})_{j\in\omega}$. This sequence satisfies q_Δ, which completes the proof. $\square$

Lemma 4.15 (Stretching indiscernibles)**.** *Let $\mathcal{I} = (a_i)_{i\in I}$ be an infinite B-indiscernible sequence in $\mathcal{U}^x$. For any linear order $(J, \leq)$, there is a J-indexed B-indiscernible sequence $\mathcal{J} = (a'_j)_{j\in J}$ satisfying* $\mathrm{EM}(\mathcal{I}/B)$.

Proof. Consider the partial type:

$$q = \{\varphi(y_{j_1}, \ldots, y_{j_n}) \mid \varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B), j_1 < \cdots < j_n \in J\}.$$

A finite subset of q mentions only finitely many variables $y_{j_1}, \ldots, y_{j_m}$. Assuming $j_1 < \cdots < j_m$ in J, this finite subset is satisfied by any $a_{i_1}, \ldots, a_{i_m}$ in $\mathcal{I}$ with $i_1 < \cdots < i_m$. By compactness, q is consistent. Since $\mathrm{EM}(\mathcal{I}/B)$ is complete, any realization of q is B-indiscernible. $\square$

Given a formula $\varphi(x_1, \ldots, x_n)$ and a permutation $\sigma \in S_n$, write φ^σ for the formula obtained from φ by substituting $x_{\sigma(i)}$ for x_i everywhere for all $1 \leq i \leq n$. That is, $\models \varphi^\sigma(a_1, \ldots, a_n)$ if and only if $\models \varphi(a_{\sigma(1)}, \ldots, a_{\sigma(n)})$.

We say $\mathrm{EM}(\mathcal{I}/B)$ is **symmetric** if whenever $\varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B)$ and $\sigma \in S_n$ is a permutation, also $\varphi^\sigma \in \mathrm{EM}(\mathcal{I}/B)$.

Lemma 4.16. *Let $\mathcal{I} = (a_i)_{i\in I}$ be a B-indiscernible sequence. $\mathcal{I}$ is a B-indiscernible set if and only if* $\mathrm{EM}(\mathcal{I}/B)$ *is symmetric.*

Proof. Suppose $\mathcal{I}$ is a B-indiscernible set. Let $\varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B)$ and $\sigma \in S_n$. To show $\varphi^\sigma \in \mathrm{EM}(\mathcal{I}/B)$, let $i_1 < \cdots < i_n$. Then $\models \varphi(a_{i_1}, \ldots, a_{i_n})$, and since $\mathcal{I}$ is a B-indiscernible set, also $\models \varphi(a_{i_{\sigma(1)}}, \ldots, a_{i_{\sigma(n)}})$, so $\models \varphi^\sigma(a_{i_1}, \ldots, a_{i_n})$, as desired.

Now suppose $\mathrm{EM}(\mathcal{I}/B)$ is symmetric. Let $\varphi(x_1, \ldots, x_n)$ be an L_B-formula, and let $i_1, \ldots, i_n$ and $j_1, \ldots, j_n$ be sequences of distinct elements in I. There are some permutations μ and ν with $i_{\mu(1)} < \cdots < i_{\mu(n)}$ and $j_{\nu(1)} < \cdots < j_{\nu(n)}$. Now if $\models \varphi(a_{i_1}, \ldots, a_{i_n})$, then $\models \varphi^{\mu^{-1}}(a_{i_{\mu(1)}}, \ldots, a_{i_{\mu(n)}})$, so $\varphi^{\mu^{-1}} \in \mathrm{EM}(\mathcal{I}/B)$. By symmetry, $\varphi^{\nu^{-1}} = (\varphi^{\mu^{-1}})^{\nu^{-1}\circ\mu} \in \mathrm{EM}(\mathcal{I}/B)$. So $\models \varphi^{\nu^{-1}}(a_{j_{\nu(1)}}, \ldots, a_{j_{\nu(n)}})$, and $\models \varphi(a_{j_1}, \ldots, a_{j_n})$. $\square$

The "absence of order" characterizing stable theories can be detected at the level of indiscernible sequences. We need one more easy lemma: our original definition of the order property applied only to formulas without parameters, but formulas with parameters can also witness the order property.

Lemma 4.17. *Suppose $\varphi(x;y)$ is a formula with parameters from C with the order property. Then some formula $\psi(x;y,z)$ with no parameters has the order property.*

Proof. Suppose $(a_i)_{i\in\omega}$ and $(b_j)_{j\in\omega}$ witness the order property for $\varphi(x;y)$. Make the parameters from C explicit by writing $\varphi(x;y) = \psi(x;y,c)$. Then $(a_i)_{i\in\omega}$ and $(b_jc)_{j\in\omega}$ witness the order property for $\psi(x;y,z)$: $\models \psi(a_i;b_j,c)$ if and only if $\varphi(a_i;b_j)$ if and only if $i \leq j$. $\square$

Theorem 4.18. *T is stable if and only if every indiscernible sequence is an indiscernible set.*

Proof. First, assume T is not stable. Then some formula $\varphi(x;y)$ has the order property, witnessed by $(a_i)_{i\in\omega}$ in $\mathcal{U}^x$ and $(b_j)_{j\in\omega}$ in $\mathcal{U}^y$. Let $\mathcal{I} = (a_nb_n)_{n\in\omega}$. Consider the formula $\theta(x_1,y_1,x_2,y_2) : \varphi(x_1;y_2)$. Letting τ be the transposition swapping 1 and 2, the formula θ^τ is $\varphi(x_2;y_1)$. For all $i<j$, $\models \varphi(a_i;b_j)$, but $\models \neg\varphi(a_j;b_i)$, so $\theta \in \mathrm{EM}(\mathcal{I})$ but $\neg\theta^\tau \in \mathrm{EM}(\mathcal{I})$.

By the Standard Lemma, there is an indiscernible sequence $\mathcal{J}$ realizing $\mathrm{EM}(\mathcal{I})$. Then $\mathrm{EM}(\mathcal{J})$ is not symmetric, so by Lemma 4.16, $\mathcal{J}$ is not an indiscernible set.

Conversely, suppose there is some B-indiscernible sequence $\mathcal{I} = (a_i)_{i\in I}$ in $\mathcal{U}^x$ which is not a B-indiscernible set. By Lemma 4.16, there is some formula $\varphi(x_1,\dots,x_n) \in \mathrm{EM}(\mathcal{I}/B)$ and some $\sigma \in S_n$ such that $\varphi^\sigma \notin \mathrm{EM}(\mathcal{I}/B)$. By stretching, we may assume that $\mathcal{I}$ is indexed by $\omega+\omega$.

Write τ_k for the transposition in S_n which swaps k and $k+1$ and fixes all other elements. The set $\{\tau_k \mid 1 \leq k < n\}$ is a generating set for S_n. Writing σ as a product of these transpositions, we can transform φ to φ^σ in a sequence of steps. It follows that there is a formula $\psi \in \mathrm{EM}(\mathcal{I}/B)$ and a transposition τ_k such that $\psi^{\tau_k} \notin \mathrm{EM}(\mathcal{I}/B)$.

Thus, for any sequence $i_1 < \cdots < i_k < i_{k+1} < \cdots < i_n$ from $\omega+\omega$, we have

$$\models \psi(a_{i_1},\dots,a_{i_k},a_{i_{k+1}},\dots,a_{i_n}) \text{ and } \models \neg\psi(a_{i_1},\dots,a_{i_{k+1}},a_k,\dots,a_{i_n}).$$

Let $\theta(x;y)$ be the formula $\psi(a_1,\dots,a_{k-1},x,y,a_{\omega+k+2},\dots,a_{\omega+n})$. This is a formula with parameters from $Ba_1\dots a_{k-1}a_{\omega+k+2},\dots,a_{\omega+n}$. The sequences $(a_{k+i})_{i\in\omega}$ and $(a_{k+j})_{j\in\omega}$ witness the order property, since $\models \theta(a_{k+i};a_{k+j})$ if and only if $i \leq j$ or $i<j$: which case we are in depends on whether

$$\psi(x_1,\dots,x_{k-1},x_k,x_k,x_{k+1},\dots,x_{n-1}) \in \mathrm{EM}(\mathcal{I}/B).$$

In either case, $\theta(x;y)$ has the order property, by Proposition 4.6. By Lemma 4.17, some formula without parameters has the order property, so T is unstable. $\square$

Exercise 20. Show that a formula $\varphi(x;y)$ is stable if and only if there exists $k\in\omega$ such that for any indiscernible sequence $(a_i)_{i\in I}$ in $\mathcal{U}^x$ and any $b \in \mathcal{U}^y$, we have $|\{a_i \mid i \in I \text{ and } \models \varphi(a_i;b)\}| \leq k$ or $|\{a_i \mid i\in I \text{ and } \models \neg\varphi(a_i;b)\}| \leq k$.

4.3 Ehrenfeucht–Mostowski models

In this section, we introduce a new method for building models of T, using indiscernible sequences. The resulting models, called Ehrenfeucht–Mostowski models, have two classic applications. The original application, due to Ehrenfeucht and Mostowski, was to build models with the maximal number of automorphisms. The second, due to Morley, was to build uncountable models realizing few types over countable sets. The latter result will lead to our first theorem on uncountably categorical theories: if T is κ-categorical for some uncountable κ, then T is $\aleph_0$-stable.

Definition 4.19. Let $\varphi(x; y_1, \dots, y_n)$ be a partitioned formula, where each variable x and y_i is a singleton. A **Skolem function** for φ is an n-ary function f such that

$$T \models \forall y_1, \dots, y_n\, ((\exists x\, \varphi(x; y_1, \dots, y_n)) \to \varphi(f(y_1, \dots, y_n); y_1, \dots, y_n)).$$

We include the case $n = 0$, in which case f is a 0-ary function symbol, i.e., a constant symbol.

We define the language L_{Sk} to be L together with a new n-ary function symbol f_φ for each L-formula $\varphi(x; y_1, \dots, y_n)$. The **Skolemization** of T is the L_{Sk}-theory T_{Sk} obtained by adding to T the sentences asserting that each f_φ is a Skolem function for φ. Note that T_{Sk} is not a complete theory.

The key property of T_{Sk} is the following: if $M \models T_{\mathrm{Sk}}$ and $A \subseteq M$, let $N = \langle A \rangle_{\mathrm{Sk}}$, the substructure of M generated by A. Then by the Tarski–Vaught test, $N|_L \preceq M|_L$. In particular, $N \models T$. Similarly, if $M, N \models T_{\mathrm{Sk}}$ and $h\colon N \to M$ is a homomorphism of L_{Sk}-structures, then $h\colon N|_L \to M|_L$ is an elementary embedding of L-structures.

Exercise 21. If necessary, review the Tarski–Vaught test and prove the assertions in the previous paragraph.

We can expand $\mathcal{U}$ to a model $\mathcal{U}_{\mathrm{Sk}} \models T_{\mathrm{Sk}}$, which we call a **Skolemization** of $\mathcal{U}$: for each formula $\varphi(x; y_1, \dots, y_n)$ and each tuple $(b_1, \dots, b_n) \in \mathcal{U}^n$, define $f_\varphi(b_1, \dots, b_n)$ to be an arbitrary element of $\varphi(\mathcal{U}; b_1, \dots, b_n)$, if this set is nonempty. Otherwise, define $f_\varphi(b_1, \dots, b_n)$ to be an arbitrary element of $\mathcal{U}$.

Fix a Skolemization $\mathcal{U}_{\mathrm{Sk}}$ of $\mathcal{U}$. Let $(a_i)_{i \in \omega}$ be an indiscernible sequence in $\mathcal{U}^1_{\mathrm{Sk}}$ such that $a_i \neq a_j$ when $i \neq j$ (which exists by the Standard Lemma). Let $\Sigma = \mathrm{EM}((a_i)_{i \in \omega})$.

Now for any linear order $(I, \le)$, by stretching there is an indiscernible sequence $\mathcal{I} = (a_i)_{i \in I}$ indexed by I realizing Σ. We define the **EM-model** with spine I and EM-type Σ by $M^\Sigma_I = \langle \mathcal{I} \rangle_{\mathrm{Sk}}|_L \preceq \mathcal{U}$. Note that for any element $b \in M^\Sigma_I$, there is some L_{Sk}-term $t(x_1, \dots, x_n)$ and some $a_{i_1}, \dots, a_{i_n}$ in $\mathcal{I}$ with $i_1 < \dots < i_n$ (without loss of generality) such that $b = t^{\mathcal{U}_{\mathrm{Sk}}}(a_{i_1}, \dots, a_{i_n})$. It follows that when I is infinite, $|M^\Sigma_I| = |I|$.

Exercise 22. Let $(I, \le)$ and $(J, \le)$ be linear orders, with associated EM-models M^Σ_I and M^Σ_J. Let $e\colon (I, \le) \to (J, \le)$ be an embedding of linear orders. Show that e induces an elementary embedding $\widehat{e}\colon M^\Sigma_I \to M^\Sigma_J$.

Verify that the assignment $(I, \leq) \mapsto M_I^\Sigma$ and $e \mapsto \widehat{e}$ defines a functor from the category of linear orders and order embeddings to the category of models of T and elementary embeddings. (This functor is not very "canonical": it depends on the choice of Skolemization $\mathcal{U}_{\mathrm{Sk}}$ and on the EM-type Σ.)

For any model $M \models T$, if $|M| = \kappa$, then the number of automorphisms of M is at most $\kappa^\kappa = 2^\kappa$. Using EM-models, we can show that this upper bound is always attained.

Exercise 23. Let κ be an infinite cardinal. Show that there is a model $M \models T$ with $|M| = \kappa$ and $|\mathrm{Aut}(M)| = 2^\kappa$.

As we have seen, it is easy to realize types, and hence to build κ-saturated models which realize lots of types. It is much harder to realize very few types! EM-models are one way to do this.

Theorem 4.20. *Let $(I, \leq)$ be a well-ordered set, and let M_I^Σ be an EM-model with spine I. Then for any countable set $B \subseteq M_I^\Sigma$, at most countably many types in $S_1(B)$ are realized in M_I^Σ.*

Proof. Let $B \subseteq M_I^\Sigma$ be a countable set. Let $\mathcal{I} = (a_i)_{i \in I}$ be the L_{Sk}-indiscernible sequence such that $M_I^\Sigma = \langle \mathcal{I} \rangle_{\mathrm{Sk}}|_L$, and write $M = \langle \mathcal{I} \rangle_{\mathrm{Sk}}$ (so M is an L_{Sk}-structure and $M_I^\Sigma = M|_L$). We will show the that at most countably many L_{Sk}-types from $S_1^{L_{\mathrm{Sk}}}(B)$ are realized in M. This suffices, since if two elements have the same L_{Sk}-type over B, they have the same L-type over B.

For each $b \in B$, pick some finite sequence $a_{i_1}, \ldots, a_{i_n}$ from $\mathcal{I}$ such that $b = t(a_{i_1}, \ldots, a_{i_n})$ for some L_{Sk}-term $t(x_1, \ldots, x_n)$. Let $J \subseteq I$ be the set of all indices of elements obtained in this way, and let $\mathcal{J} = (a_j)_{j \in J}$. Then J is countable and $B \subseteq \langle \mathcal{J} \rangle_{\mathrm{Sk}}$. It suffices to show that at most countably many L_{Sk}-types over $\mathcal{J}$ are realized in M, since if two elements have the same L_{Sk}-type over $\mathcal{J}$, they have the same L_{Sk}-type over B.

Now for each element $c \in M$, pick some L_{Sk}-term $t(x_1, \ldots, x_n)$ and some $a_{i_1}, \ldots, a_{i_n}$ in $\mathcal{I}$ with $i_1 < \cdots < i_n$, such that $c = t(a_{i_1}, \ldots, a_{i_n})$. We define the **signature** of c to consist of:

(1) The L_{Sk}-term $t(x_1, \ldots, x_n)$.

(2) For each $1 \leq k \leq n$, one of the following pieces of data:

  (a) If $i_k \in J$: the string "$= i_k$".
  (b) If i_k is greater than every element of J: the symbol "∞".
  (c) If $i_k \notin J$ but is bounded above by an element of j: the string "$< j$", where $j \in J$ is least such that $i_k < j$ (here we use the well-ordering).

Since L_{Sk} is a countable language and J is countable, there are only countably many possible signatures. Thus it suffices to show that if c and c' have the same signature, then $\mathrm{tp}_{L_{\mathrm{Sk}}}(c/\mathcal{J}) = \mathrm{tp}_{L_{\mathrm{Sk}}}(c'/\mathcal{J})$.

Let $\varphi(z, a_{j_1}, \ldots, a_{j_m}) \in \mathrm{tp}_{L_{\mathrm{Sk}}}(c/\mathcal{J})$, where $j_1 < \cdots < j_m \in J$. Since c and c' have the same signature, we have $c = t(a_{i_1}, \ldots, a_{i_n})$ and $c' = t(a_{i'_1}, \ldots, a_{i'_n})$

for a fixed L_{Sk}-term t. Additionally, for each $1 \leq k \leq n$ and each $1 \leq \ell \leq m$, $i_k \leq j_\ell$ if and only if $i'_k \leq j_\ell$, and $i_k \geq j_\ell$ if and only if $i'_k \geq j_\ell$.

Since $\varphi(z, a_{j_1}, \ldots, a_{j_m}) \in \mathrm{tp}_{L_{\mathrm{Sk}}}(c/\mathcal{J})$,

$$\models \varphi(t(a_{i_1}, \ldots, a_{i_n}), a_{j_1}, \ldots, a_{j_m}).$$

By indiscernibility of $\mathcal{I}$,

$$\models \varphi(t(a_{i'_1}, \ldots, a_{i'_n}), a_{j_1}, \ldots, a_{j_m}),$$

and hence $\varphi(z, a_{j_1}, \ldots, a_{j_m}) \in \mathrm{tp}_{L_{\mathrm{Sk}}}(c'/\mathcal{J})$. This completes the proof. $\square$

Corollary 4.21. *Let κ be an uncountable cardinal. If T is κ-categorical, then T is $\aleph_0$-stable.*

Proof. Suppose for contradiction that T is not $\aleph_0$-stable. Then there is some countable set B with $|S_1(B)| > \aleph_0$. Pick some set $X \subseteq S_1(B)$ with $|X| = \aleph_1$. By Lemma 2.12, there is a model M containing B and realizing all the types in X, and furthermore $|M| = \aleph_1$. Since $\aleph_1 \leq \kappa$, M has an elementary extension $M \preceq M'$ with $|M'| = \kappa$. Then $B \subseteq M'$ and M' realizes uncountably many types in $S_1(B)$.

Let M^Σ_κ be an EM-model with spine κ. Since $|M^\Sigma_\kappa| = \kappa$ and T is κ-categorical, there is an isomorphism $i\colon M' \cong M^\Sigma_\kappa$. Then $i(B)$ is a countable subset of M^Σ_κ, and M^Σ_κ realizes uncountably many types in $S_1(i(B))$. This contradicts Theorem 4.20. $\square$

5 Totally transcendental theories

5.1 $\aleph_0$-stability

Recall that T is stable if and only if there is no binary tree of φ-formulas (with parameters) for any partitioned formula $\varphi(x; y)$. You may recall from a previous course in model theory that T is small (Definition 2.19) if and only if there is no binary tree of formulas *without* parameters.

We now consider a common strengthening of these notions: a totally transcendental theory is one in which there is no binary tree of formulas *with* parameters.

Definition 5.1. Let φ and ψ be formulas with parameters in the same variable context x.

- We say φ **implies** ψ if $\models \forall x\, (\varphi \to \psi)$. Equivalently, $\varphi(\mathcal{U}) \subseteq \psi(\mathcal{U})$.
- We say φ and ψ are **contradictory** if $\models \forall x\, \neg(\varphi \wedge \psi)$. Equivalently, $\varphi(\mathcal{U}) \cap \psi(\mathcal{U}) = \varnothing$.

Definition 5.2. T is **totally transcendental**[4] if there is no binary tree of formulas with parameters $\{\varphi_g(x) \mid g \in 2^{<\omega}\}$ in a common finite variable context x, such that for all $g \in 2^{<\omega}$:

- φ_g is consistent.
- φ_{g0} and φ_{g1} each imply φ_g.
- φ_{g0} and φ_{g1} are contradictory.

One way to think about the relationship between totally transcendental and stable theories is that the *local* behavior of stable formulas is true *globally* for totally transcendental theories. And just as a formula is stable if and only if is κ-stable for all infinite κ, we have the following theorem:

Theorem 5.3. *The following are equivalent:*

(1) T is totally transcendental.

(2) T is $\aleph_0$-stable.

(3) T is κ-stable for all infinite cardinals κ.

[4]The terminology is due to Morley, who defined a notion of rank he called "transcendental rank". As we will see in Theorem 5.16 below, a theory is totally transcendental if and only if every formula has an ordinal valued transcendental rank. The transcendental rank was later renamed "Morley rank", but the name "totally transcendental" stuck. In Theorem 5.3 below, we will see that T is totally transcendental if and only if it is $\aleph_0$-stable. For this reason, totally transcendental theories are often called "ω-stable theories" (it is traditional to write ω-stable instead of $\aleph_0$-stable). But this equivalence only holds when the language is countable, and totally transcendental is the more robust notion. So I will stick to this terminology, even though we always work with a countable language in this class.

Proof. (3)⇒(2): Trivial.

(2)⇒(1): Suppose T is not totally transcendental. Let $\{\varphi_g(x) \mid g \in 2^{<\omega}\}$ be a binary tree of formulas, and let B be the set of all parameters appearing in these formulas. Then B is countable. For each $f \in 2^\omega$, consider the partial type $p_f = \{\varphi_{f|_n}(x) \mid n \in \omega\}$. I claim that p_f is a consistent partial type over B. By compactness, we only need to show that $\{\varphi_{f|_n}(x) \mid n \leq N\}$ is consistent for $N \in \omega$. But $\varphi_{f|_N}(x)$ is consistent by hypothesis, and $\varphi_{f|_N}(x)$ implies $\varphi_{f|_n}(x)$ for all $n \leq N$, so p_f is consistent.

For each $f \in 2^\omega$, let $q_f \in S_x(B)$ be a complete type extending p_f. Now if $f \neq f'$ in 2^ω, let n be least such that $f(n) \neq f'(n)$. Then q_f contains $\varphi_{f|_{n+1}}(x)$ and $q_{f'}$ contains $\varphi_{f'|_{n+1}}(x)$ and these two formulas are contradictory, so $q_f \neq q_{f'}$. Thus $|S_x(B)| = 2^{\aleph_0}$, and T is not $\aleph_0$-stable.

(1)⇒(3): Suppose there is some infinite cardinal κ such that T is not κ-stable. Let B be a set with $|B| \leq \kappa$ and $|S_1(B)| > \kappa$. For a formula $\varphi \in F_1(B)$, write $[\varphi] = \{p \in S_1(B) \mid \varphi \in p\}$. We say φ is *wide* if $|[\varphi]| > \kappa$. Otherwise, we say φ is *thin*. We say a type is *thin* if it contains a thin formula. Write $H \subseteq S_1(B)$ for the set of thin types. Then $H = \bigcup_{\varphi \text{ thin}}[\varphi]$, so $|H| \leq \kappa$.

I claim that if φ is wide, then there is a formula $\psi \in F_1(B)$ such that $\varphi \wedge \psi$ and $\varphi \wedge \neg\psi$ are both wide. Indeed, since $|[\varphi]| > \kappa$, $|[\varphi] \setminus H| > \kappa > 1$. Pick two distinct types $p, q \in [\varphi] \setminus H$. Since $p \neq q$, there is a formula $\psi \in F_1(B)$ such that $\psi \in p$ and $\neg\psi \in q$. Since p and q are not thin, $\varphi \wedge \psi \in p$ is wide and $\varphi \wedge \neg\psi \in q$ is wide.

Now we build a binary tree of wide formulas by induction. Let $\varphi_\varnothing$ be $\top$, which is wide since $[\top] = S_1(B)$. Given a wide formula φ_g with $g \in 2^{<\omega}$, let ψ be a wide formula with that $\varphi_g \wedge \psi$ and $\varphi_g \wedge \neg\psi$ are both wide. Set $\varphi_{g0} = \varphi_g \wedge \psi$, and set $\varphi_{g1} = \varphi_g \wedge \neg\psi$. This completes the proof. $\square$

We could have used the proof strategy of (1)⇒(3) in the proof of the Stable Formula Theorem to prove that a formula without the binary tree is κ-stable for all infinite κ. However, we preferred to go through definability of types, because of the independent importance of this notion.

Also, in the Stable Formula Theorem, we were able to prove that if φ is κ-stable for *some* infinite cardinal κ, then φ does not have the binary tree property. Here we are only able to prove that if T is $\aleph_0$-stable, then T is totally transcendental. This is because a binary tree of formulas cannot be blown up to arbitrary ordinal height by compactness, unlike an instance of the binary tree property for a single formula.

5.2 Prime and atomic models

Recall that for any set A, we topologize the type space $S_x(A)$ by taking as a basis sets of the form $[\varphi] = \{p \in S_x(A) \mid \varphi \in p\}$ for $\varphi \in F_x(A)$. Since $[\varphi]$ is the complement of $[\neg\varphi]$ each basic open set is in fact clopen. In this topology, $S_x(A)$ is a compact Hausdorff space.

Definition 5.4. A type $p \in S_x(A)$ is **isolated** if it is an isolated point in the space $S_x(A)$. Equivalently, if there is a formula $\varphi \in F_x(A)$ such that for all $\psi \in p$, $\models \forall x\,(\varphi \to \psi)$, so that $[\varphi] = \{p\}$.

We say **isolated types are dense over** A if for all finite contexts x, the set of isolated types is a dense set in the space $S_x(A)$. Equivalently, for every formula $\varphi \in F_x(A)$, there is some isolated type containing φ.

Exercise 24. Let a and b be finite tuples. Show that $\mathrm{tp}(ab/A)$ is isolated if and only if $\mathrm{tp}(a/A)$ is isolated and $\mathrm{tp}(b/Aa)$ is isolated.

Theorem 5.5. *Suppose T is totally transcendental. Then for every set A, isolated types are dense over A.*

Proof. Suppose there is a set A and a context x such that isolated types are not dense in $S_x(A)$. Then there is a formula $\varphi(x) \in F_x(A)$ such that no type in $[\varphi]$ is isolated.

We show that T is not totally transcendental by building a binary tree of formulas over A by induction. Let $\varphi_\varnothing$ be φ. Given φ_g with $g \in 2^{<\omega}$, any type containing φ_g contains φ, and hence is not isolated. It follows that $[\varphi_g]$ contains at least two distinct types $p \neq q$ (otherwise φ_g would be inconsistent or would isolate the unique type in $[\varphi_g]$). So there is some formula $\psi \in F_x(A)$ such that $\psi \in p$ and $\neg\psi \in q$. Let φ_{g0} be $\varphi_g \wedge \psi$, and let φ_{g1} be $\varphi_g \wedge \neg\psi$. $\square$

Note that the same argument shows that if T is small, then isolated types are dense over $\varnothing$. A classical consequence for the model theory of countable structures is the existence of prime and atomic models (over $\varnothing$).

Definition 5.6. Let M be a model and $A \subseteq M$. We say M is **prime over** A if every partial elementary map $f\colon A \to N \models T$ extends to an elementary embedding $M \to N$. We say $M \models T$ is a **prime model** if M is prime over $\varnothing$.

Note that in the definition of prime over A, the elementary embedding $M \to N$ is not required to be unique.

Definition 5.7. Let M be a model and $A \subseteq M$. We say M is **atomic over** A if every type over A realized in M is isolated. We say a model M is **atomic** if M is atomic over $\varnothing$.

Observe that M is prime over A if and only if M is a prime model of T_A, and M is atomic over A if and only if M is an atomic model of T_A.

You may recall the following fact from a previous course in model theory:

Fact 5.8. A countable model of T is prime if and only if it is atomic. Moreover, such a model exists if and only if isolated types are dense over $\varnothing$, in which case it is unique up to isomorphism.

By way of motivation for what comes next, let me sketch the proof of the equivalence between prime and atomic models.

In one direction, if M is not atomic, there is some tuple from M which realizes a non-isolated type p. By the omitting types theorem, there is a countable model $N \models T$ which does not realize p. But then there can be no elementary embedding $M \to N$, so M is not prime.

In the other direction, if M is countable and atomic, we can enumerate M as $(a_n)_{n\in\omega}$, and $\mathrm{tp}(a_n/a_0 \dots a_{n-1})$ is isolated for all n (by Exercise 24). For any other $N \models T$, we can define an elementary embedding $M \to N$ one step at a time, using the fact that isolated types are realized in every model. So M is prime.

The uniqueness of the countable prime/atomic model can then be established similarly, using a back-and-forth argument.

What about prime and atomic models over arbitrary sets A? When A is countable, we can just apply Fact 5.8 to the theory T_A. But when A is uncountable, two crucial points in the arguments above break down. First, the omitting types theorem only works when the language is countable. Second, if M is atomic of cardinality $\kappa = |A|$, we can enumerate $M = (a_\alpha)_{\alpha<\kappa}$, but there is no guarantee that $\mathrm{tp}(a_\alpha/(a_\beta)_{\beta<\alpha})$ is isolated when α is infinite.

The solution to this issue is to work with models which do have a nice enumeration.

Definition 5.9. A **construction sequence** over a set A is a sequence $(b_\beta)_{\beta<\alpha}$ indexed by an ordinal α such that for all $\gamma < \alpha$, $\mathrm{tp}(b_\gamma/A(b_\beta)_{\beta<\gamma})$ is isolated.

Let M be a model and $A \subseteq M$. We say M is **constructible** over A if there is a construction sequence $(m_\beta)_{\beta<\alpha}$ which enumerates M.

Theorem 5.10. *If a model M is constructible over A, then it is prime over A and atomic over A.*

Proof. Suppose M is enumerated by the construction sequence $(m_\beta)_{\beta<\alpha}$.

To show M is prime, let $f\colon A \to N \models T$ be a partial elementary map. We define a sequence of partial elementary maps $(f_\beta)_{\beta\leq\alpha}$ by recursion, such that $\mathrm{dom}(f_\beta) = A \cup \{m_\gamma \mid \gamma < \beta\}$.

Let $f_0 = f$. When γ is a limit ordinal, let $f_\gamma = \bigcup_{\beta<\gamma} f_\beta$. Given f_β, consider $p = \mathrm{tp}(m_\beta/A(m_\gamma)_{\gamma<\beta})$. By hypothesis, p is isolated by a formula φ. Then $(f_\beta)_*p$ is isolated by $(f_\beta)_*\varphi$. Since $M \models \exists x\, \varphi$ and f_β is partial elementary, $N \models \exists x\, (f_\beta)_*\varphi$, so we can pick some $n_\beta \in N$ realizing $(f_\beta)_*p$. Let $f_{\beta+1} = f_\beta \cup \{(m_\beta, n_\beta)\}$.

To show M is atomic, we argue by induction on β that for any tuple $m_{\beta_1}, \dots, m_{\beta_n}$, with $\max(\beta_1, \dots, \beta_n) = \beta$, $\mathrm{tp}(m_{\beta_1}, \dots, m_{\beta_n}/A)$ is isolated.

First note that if the $\beta_1, \dots, \beta_n$ are not distinct, say if $\beta_n = \beta_{n-1}$, then $\mathrm{tp}(m_{\beta_n}/Am_{\beta_1} \dots m_{\beta_{n-1}})$ is isolated (by $x = \beta_{n-1}$), so by Exercise 24, it suffices to show that $\mathrm{tp}(m_{\beta_1} \dots m_{\beta_{n-1}}/A)$ is isolated. Thus, without loss of generality, we may assume $\beta_1 < \cdots < \beta_n = \beta$.

Now by hypothesis, $\mathrm{tp}(m_\beta/A(m_\gamma)_{\gamma<\beta})$ is isolated, say by $\varphi(x, m_{\gamma_1}, \dots, m_{\gamma_k})$. Consider the tuple $m^* = (m_{\beta_1}, \dots, m_{\beta_{n-1}}, m_{\gamma_1}, \dots, m_{\gamma_k})$. We have

$$\max(\beta_1, \dots, \beta_{n-1}, \gamma_1, \dots, \gamma_k) < \beta,$$

so by induction $\mathrm{tp}(m^*/A)$ is isolated, and φ isolates $\mathrm{tp}(m_\beta/Am^*)$, so by Exercise 24, $\mathrm{tp}(m^* m_\beta/A)$ is isolated. But our original tuple is a subtuple of $m^* m_\beta$, so its type over A is isolated, again by Exercise 24. $\square$

Theorem 5.11. *Suppose that for every set B, isolated types are dense over B. Then for any set A, T has a constructible model over A.*

Proof. We build a construction sequence by transfinite recursion. First enumerate A as $(b_\beta)_{\beta<\alpha}$. Suppose for $\gamma \geq \alpha$ we have a construction sequence $(b_\beta)_{\beta<\gamma}$. Let $B = \{b_\beta \mid \beta < \gamma\}$. If B is an elementary substructure of $\mathcal{U}$, then B is a constructible model over A, and we are done.

Otherwise, by the contrapositive of the Tarski–Vaught test, there is some formula $\varphi(x)$ over B such that $\mathcal{U} \models \exists x\, \varphi(x)$, but there is no realization of φ in B. Since isolated types are dense over B, there is a type $p \in S_x(B)$ containing φ and isolated by a formula $\psi \in F_x(B)$. Since $\models \exists x\, \psi(x)$, we can define b_γ to be some element of $\mathcal{U}$ satisfying ψ, so that $\mathrm{tp}(b_\gamma/(b_\beta)_{\beta<\gamma}) = p$ is isolated.

This construction eventually stops with a constructible model over A. $\square$

Corollary 5.12. *Suppose T is totally transcendental. For every set B, T has prime models over B, and every prime model over B is atomic over B.*

Proof. By Theorem 5.5, isolated types are dense over B. By Theorem 5.11, T has a constructible model M over B. By Theorem 5.10, M is prime and atomic over B.

Now suppose $N \models T$ is prime over B. Then the inclusion $B \to M$ extends to an elementary embedding $N \to M$. Then every type in $S_x(B)$ realized in N is realized in M. Since M is atomic over B, N is atomic over B. $\square$

In fact, more is true: constructible models over B are unique up to isomorphism, and if T is totally transcendental, then every prime model over B is constructible over B. So totally transcendental theories admit canonical prime models over arbitrary sets. These results are more difficult. Proofs can be found in Marker or Tent and Ziegler.

Exercise 25. Consider the theory T_E of an equivalence relation with infinitely many infinite classes.

(1) Let A be an uncountable set of elements in a single equivalence class. Describe a constructible model over A, and define a construction sequence enumerating it.

(2) Let A be an uncountable set elements, no two of which are equivalent. Describe a constructible model over A, and define a construction sequence enumerating it.

Not every theory with constructible models over all sets is totally transcendental, or even stable, as the following example shows.

Exercise 26. Let $T = \mathrm{DLO}$. Show that isolated types are dense over all sets (and hence T has constructible models over all sets). It may help to first prove that it suffices to show that isolated types are dense in $S_1(A)$ for all sets A.

5.3 Morley rank

Just as in the case of a stable formula, in totally transcendental theories it will be useful to define a rank which measures how much a definable set can be "cut up" into smaller definable sets. We could define a global 2-rank in this case, analogously to the local 2-rank. But it turns out to be more useful if we "cut up" our definable sets into infinitely many pieces. In particular, this will allow us to prove that the rank of a union of two definable sets is the maximum of their ranks, which is a desirable property for notion of "dimension".

The resulting notion of rank, the Morley rank, takes values in $\mathrm{Ord} \cup \{\pm\infty\}$. We need to consider infinite ordinal values here because (unlike the case of the local 2-rank) we cannot use compactness to blow up trees of arbitrary finite height to trees of arbitrary ordinal height.

Definition 5.13. For any formula with parameters $\varphi(x)$, we define the **Morley rank** of φ, $\mathrm{MR}(\varphi)$, recursively as follows:

(1) $\mathrm{MR}(\varphi) \geq 0$ if and only if φ is satisfiable.

(2) $\mathrm{MR}(\varphi) \geq \alpha + 1$ if and only if there is a family of formulas with parameters $(\psi_i(x))_{i\in\omega}$ such that:

- $\mathrm{MR}(\psi_i) \geq \alpha$ for all i.
- ψ_i implies φ for all i.
- ψ_i and ψ_j are contradictory for all $i \neq j$.

(3) $\mathrm{MR}(\varphi) \geq \gamma$ when γ is a limit ordinal if and only if $\mathrm{MR}(\varphi) \geq \beta$ for all $\beta < \gamma$.

If φ is not satisfiable, we set $\mathrm{MR}(\varphi) = -\infty$. If $\mathrm{MR}(\varphi) \geq \alpha$ for all ordinals α, we set $\mathrm{MR}(\varphi) = \infty$. Otherwise, we set $\mathrm{MR}(\varphi)$ to be the maximal ordinal α such that $\mathrm{MR} \geq \alpha$.

Exercise 27. Suppose φ and ψ are formulas with parameters in the same context x. If φ implies ψ, show that $\mathrm{MR}(\varphi) \leq \mathrm{MR}(\psi)$. (*Hint:* Show by induction on α that if $\mathrm{MR}(\varphi) \geq \alpha$, then $\mathrm{MR}(\psi) \geq \alpha$.) Observe that it follows that equivalent formulas have equal Morley rank.

As a consequence of Exercise 27, the Morley rank is a property of the definable set D, not the particular formula with parameters defining D. We have:

- $\mathrm{MR}(D) = 0$ if and only if D is finite.
- $\mathrm{MR}(D) = 1$ if and only if D is infinite, but there is no way to cut up D into infinitely many disjoint infinite definable sets.
- $\mathrm{MR}(D) = 2$ if and only if D can be cut up into infinitely many disjoint infinite definable sets, but not in such a way that each of these sets can be cut up into infinitely many disjoint infinite definable sets.

- Note that $\mathrm{MR}(D) \geq \omega$ does not mean that D can be cut up into a countably branching tree of definable sets of countable height. Rather, it means that for any finite n, D can be cut up into a countably branching tree of definable sets of height n (but different n may require different trees).

Exercise 28. Let T be the theory of non-empty acyclic graphs such that every vertex has infinitely many neighbors (i.e., infinitely branching forests). Show that T is complete and totally transcendental, and when x is a single variable, $\mathrm{MR}(x = x) = \omega$.

Example 5.14. With $T = \mathrm{DLO}$, for any $a < b$, we have $\mathrm{MR}(a < x < b) = \infty$. We prove by induction that $\mathrm{MR}(a < x < b) \geq \alpha$ for all ordinals α. Since $<$ is dense, $a < x < b$ is satisfiable, so $\mathrm{MR}(a < x < b) \geq 0$. When γ is a limit ordinal, we have $\mathrm{MR}(a < x < b) \geq \beta$ for all $\beta < \gamma$ by induction, so $\mathrm{MR}(a < x < b) \geq \gamma$. Finally, to show $\mathrm{MR}(a < x < b) \geq \alpha + 1$, pick $a = a_0 < b_0 < a_1 < b_1 < \cdots < b$. By induction, $\mathrm{MR}(a_i < x < b_i) \geq \alpha$ for all $i \in \omega$, and these formulas imply $a < x < b$ and are pairwise contradictory.

Exercise 29. If $\varphi(x; b)$ is a formula with parameters $b \in \mathcal{U}^y$, show that $\mathrm{MR}(\varphi)$ only depends on the formula $\varphi(x; y)$ and the type $\mathrm{tp}(b) \in S_y(\varnothing)$. That is, if $\mathrm{tp}(b) = \mathrm{tp}(b')$, then $\mathrm{MR}(\varphi(x; b)) = \mathrm{MR}(\varphi(x; b'))$.

Lemma 5.15. *There is an ordinal γ_T such that for every formula with parameters θ, if $\mathrm{MR}(\theta) \geq \gamma_T$, then $\mathrm{MR}(\theta) = \infty$.*

Proof. Since there are only countably many formulas $\varphi(x; y)$ and at most $2^{\aleph_0}$-many types in $S_y(\varnothing)$, by Exercise 29 only $2^{\aleph_0}$-many ordinals are in the range of the function MR. Let γ_T be the least ordinal greater than the every ordinal in the range. $\square$

With a little more work, one can show that the ordinal γ_T is always $\leq \aleph_1$ (in a countable language). So the Morley rank of any formula is a countable ordinal. But we will not need this fact.

Theorem 5.16. *T is totally transcendental if and only if no formula with parameters has Morley rank ∞.*

Proof. First, assume T is not totally transcendental. Let $\{\varphi_g(x) \mid g \in 2^{<\omega}\}$ be a binary tree of formulas. I claim, by induction on α, that for all $g \in 2^{<\omega}$, $\mathrm{MR}(\varphi_g) \geq \alpha$. In particular, $\mathrm{MR}(\varphi_\varnothing) = \infty$.

In the base case, each φ_g is satisfiable by assumption, so $\mathrm{MR}(\varphi_g) \geq 0$.

When γ is a limit ordinal, assume $\mathrm{MR}(\varphi_g) \geq \beta$ for all $g \in 2^{<\omega}$ and all $\beta < \gamma$. Then $\mathrm{MR}(\varphi_g) \geq \gamma$ for all $g \in 2^{<\omega}$ by definition.

For the successor step, assume $\mathrm{MR}(\varphi_g) \geq \alpha$ for all $g \in 2^{<\omega}$. Fix some $h \in 2^{<\omega}$. Define $\psi_0 = \varphi_{h1}$, $\psi_1 = \varphi_{h01}$, $\psi_2 = \varphi_{h001}$, and in general $\psi_i = \varphi_{h0^i1}$. For all i, ψ_i implies φ_h and $\mathrm{MR}(\psi_i) \geq \alpha$ by the inductive hypothesis. Also, when $i < j$, $\psi_j = \varphi_{h0^j1}$ implies φ_{h0^i0}, so ψ_i and ψ_j are contradictory. Thus $\mathrm{MR}(\varphi_h) \geq \alpha + 1$.

Conversely, assume φ is a formula with parameters with $\mathrm{MR}(\varphi) = \infty$. We build a binary tree of formulas with Morley rank ∞ by induction. Let $\varphi_\varnothing$ be φ.

Suppose we are given a formula φ_g of Morley rank ∞ with $g \in 2^{<\omega}$. By Lemma 5.15, there is an ordinal γ_T such that if $\mathrm{MR}(\theta) \geq \gamma_T$, then $\mathrm{MR}(\theta) = \infty$. Since $\mathrm{MR}(\varphi_g) \geq \gamma_T + 1$, there is a family of formulas with parameters $(\psi_i(x))_{i\in\omega}$ such that:

- $\mathrm{MR}(\psi_i) \geq \gamma_T$ for all i.
- ψ_i implies φ_g for all i.
- ψ_i and ψ_j are contradictory for all $i \neq j$.

Let $\varphi_{g0} = \psi_0$, and let $\varphi_{g1} = \psi_1$. Then φ_{g0} and φ_{g1} imply φ_g and are contradictory. We have $\mathrm{MR}(\psi_0) \geq \gamma_T$ and $\mathrm{MR}(\psi_1) \geq \gamma_T$, so $\mathrm{MR}(\varphi_{g0}) = \mathrm{MR}(\varphi_{g1}) = \infty$, and we can continue the induction. □

Lemma 5.17. *Let φ and ψ be formulas with parameters in the same context x. Then:*

$$\mathrm{MR}(\varphi \vee \psi) = \max(\mathrm{MR}(\varphi), \mathrm{MR}(\psi)).$$

Proof. Since φ implies $\varphi \vee \psi$, $\mathrm{MR}(\varphi) \leq \mathrm{MR}(\varphi \vee \psi)$ by Exercise 27. The same is true for ψ, so $\max(\mathrm{MR}(\varphi), \mathrm{MR}(\psi)) \leq \mathrm{MR}(\varphi \vee \psi)$.

Conversely, it suffices to show by induction on α that for all formulas φ and ψ, if $\mathrm{MR}(\varphi \vee \psi) \geq \alpha$, then $\max(\mathrm{MR}(\varphi), \mathrm{MR}(\psi)) \geq \alpha$.

In the base case, if $\varphi \vee \psi$ is satisfiable, then φ is satisfiable or ψ is satisfiable.

If γ is a limit ordinal and $\mathrm{MR}(\varphi \vee \psi) \geq \gamma$, then $\mathrm{MR}(\varphi \vee \psi) \geq \beta$ for all $\beta < \gamma$. By induction, $\max(\mathrm{MR}(\varphi), \mathrm{MR}(\psi)) \geq \beta$ for all $\beta < \gamma$. It follows that $\max(\mathrm{MR}(\varphi), \mathrm{MR}(\psi)) \geq \gamma$.

If $\mathrm{MR}(\varphi \vee \psi) \geq \alpha + 1$, this is witnessed by a family of formulas $(\theta_i)_{i\in\omega}$. Since θ_i implies $\varphi \vee \psi$, we have that θ_i is equivalent to $\theta_i \wedge (\varphi \vee \psi)$, which is equivalent to $(\theta_i \wedge \varphi) \vee (\theta_i \wedge \psi)$. Now for all i, $\mathrm{MR}(\theta_i) \geq \alpha$, so by induction, $\max(\mathrm{MR}(\theta_i \wedge \varphi), \mathrm{MR}(\theta_i \wedge \psi)) \geq \alpha$. Then there is an infinite set $H \subseteq \omega$ such that for all $i \in H$, $\mathrm{MR}(\theta_i \wedge \varphi) \geq \alpha$, or for all $i \in H$, $\mathrm{MR}(\theta_i \wedge \psi) \geq \alpha$. In the first case, the formulas $(\theta_i \wedge \varphi)_{i\in H}$ witness $\mathrm{MR}(\varphi) \geq \alpha$. In the second case, the formulas $(\theta_i \wedge \psi)_{i\in H}$ witness $\mathrm{MR}(\psi) \geq \alpha$. □

For any ordinal α, we say that formulas (with parameters) φ and ψ are **α-equivalent**, written $\varphi \sim_\alpha \psi$ if $\mathrm{MR}((\varphi \wedge \neg\psi) \vee (\neg\varphi \wedge \psi)) < \alpha$.

For example: $\varphi \sim_0 \psi$ if and only if $\varphi(\mathcal{U}) = \psi(\mathcal{U})$. $\varphi \sim_1 \psi$ if and only if the symmetric difference of $\varphi(\mathcal{U})$ and $\psi(\mathcal{U})$ is finite. $\varphi \sim_2 \psi$ if and only if the symmetric difference of $\varphi(\mathcal{U})$ and $\psi(\mathcal{U})$ has Morley rank ≤ 1, etc.

Exercise 30. Show that $\sim_\alpha$ is an equivalence relation.

We say a formula (with parameters) φ with is **irreducible** if for all formulas ψ, either $\mathrm{MR}(\varphi \wedge \psi) < \mathrm{MR}(\varphi)$ or $\mathrm{MR}(\varphi \wedge \neg\psi) < \mathrm{MR}(\varphi)$.

A family of formulas (with parameters) $\psi_1, \ldots, \psi_n$ are **irreducible components** of φ if they are pairwise contradictory, φ is equivalent to $\bigvee_{i=1}^n \psi_i$, and for each i, ψ_i is irreducible with $\mathrm{MR}(\psi_i) = \mathrm{MR}(\varphi)$.

Theorem 5.18. *Suppose* $\mathrm{MR}(\varphi) = \alpha$ *is an ordinal. Then* φ *has a decomposition into irreducible components. Moreover, this decomposition is unique up to* α*-equivalence of the components.*

Proof. For contradiction, suppose φ does not have a decomposition into irreducible components. In particular, φ is not itself irreducible, so there is some formula θ such that $\mathrm{MR}(\varphi \wedge \theta) = \alpha$ and $\mathrm{MR}(\varphi \wedge \neg\theta) = \alpha$. If both $\varphi \wedge \theta$ and $\varphi \wedge \neg\theta$ had decompositions into irreducible components, then φ would as well. Without loss of generality, $\varphi_1 = \varphi \wedge \theta$ has no decomposition into irreducible components. Set $\psi_0 = \varphi \wedge \neg\theta$.

Now repeat the argument with φ_1, obtaining formulas φ_2 and ψ_1, each of Morley rank α, such that φ_2 has no decomposition into irreducible components. In this way, we construct a family $(\psi_n)_{n\in\omega}$ witnessing that $\mathrm{MR}(\varphi) \geq \alpha + 1$, contradiction.

For uniqueness, suppose $\psi_1, \dots, \psi_n$ and $\theta_1, \dots, \theta_m$ are both decompositions of φ into irreducible components. For all i, ψ_i is equivalent to $\bigvee_{j=1}^{m}(\psi_i \wedge \theta_j)$. Since $\mathrm{MR}(\psi_i) = \alpha$, there is at least one j such that $\mathrm{MR}(\psi_i \wedge \theta_j) = \alpha$, and since ψ_i is irreducible, there is at most one such j. This shows $n \leq m$.

Now since ψ_i is irreducible, $\mathrm{MR}(\psi_i \wedge \neg\theta_j) < \alpha$, and since θ_j is irreducible, $\mathrm{MR}(\neg\psi_i \wedge \theta_j) < \alpha$. By Lemma 5.17, $\psi_i \sim_\alpha \theta_j$. By the symmetric argument, $n = m$ and each θ_j is α-equivalent to exactly one ψ_i. □

Definition 5.19. Suppose φ is a formula whose Morley rank is an ordinal (i.e., not $-\infty$ or ∞). The **Morley degree** of φ, $\mathrm{MD}(\varphi)$, is the number of irreducible components of φ. Note that $\mathrm{MD}(\varphi) \in \omega \setminus \{0\}$.

Exercise 31. Suppose $\mathrm{MR}(\varphi) = \alpha$. Show that $\mathrm{MD}(\varphi)$ is the maximal $d \in \omega$ such that there are formulas (with parameters) $\psi_1, \dots, \psi_d$ in the same variable context as φ, such that:

- $\mathrm{MR}(\psi_i) \geq \alpha$ for all i.
- ψ_i implies φ for all i.
- ψ_i and ψ_j are contradictory for all $i \neq j$.

Note that by definition of Morley rank, we cannot have an infinite family of such formuls ψ_i. But the existence of a maximal d is not obvious.

Lemma 5.20. *Let* φ *and* ψ *be formulas with parameters in the same context* x*. Assume* φ *and* ψ *are contradictory, and both have ordinal Morley rank.*

(1) If $\mathrm{MR}(\varphi) < \mathrm{MR}(\psi)$*, then* $\mathrm{MD}(\varphi \vee \psi) = \mathrm{MD}(\psi)$*.*

(2) If $\mathrm{MR}(\varphi) = \mathrm{MR}(\psi)$*, then* $\mathrm{MD}(\varphi \vee \psi) = \mathrm{MD}(\varphi) + \mathrm{MD}(\psi)$*.*

Proof. Let $\theta_1, \dots, \theta_d$ be a decomposition of ψ into irreducible components. If $\mathrm{MR}(\varphi) < \mathrm{MR}(\psi)$, then $\theta_1' = \theta_1 \vee \varphi$ remains irreducible with Morley rank $\mathrm{MR}(\psi)$, so $\theta_1', \theta_2, \dots, \theta_d$ is a decomposition of $\varphi \vee \psi$ into irreducible components.

If $\mathrm{MR}(\varphi) = \mathrm{MR}(\psi)$, let $\chi_1, \dots, \chi_{d'}$ be a decomposition of φ into irreducible components. Then $\theta_1, \dots, \theta_d, \chi_1, \dots, \chi_{d'}$ is a decomposition of $\varphi \vee \psi$ into irreducible components. □

Finally, we would like to extend Morley rank and degree to complete types.

Definition 5.21. For a complete type $p \in S_x(A)$, we define

$$\begin{aligned}\mathrm{MR}(p) &= \min\{\mathrm{MR}(\varphi) \mid \varphi \in p\}\\ \mathrm{MD}(p) &= \min\{\mathrm{MD}(\varphi) \mid \varphi \in p \text{ and } \mathrm{MR}(\varphi) = \mathrm{MR}(p)\}.\end{aligned}$$

If $\mathrm{MR}(\varphi) = \infty$ for all $\varphi \in p$, we set $\mathrm{MR}(p) = \infty$ and leave $\mathrm{MD}(p)$ undefined. For $a \in \mathcal{U}^x$, we write $\mathrm{MR}(a/A)$ and $\mathrm{MD}(a/A)$ for $\mathrm{MR}(\mathrm{tp}(a/A))$ and $\mathrm{MD}(\mathrm{tp}(a/A))$.

A **global** type is a type in $S_x(\mathcal{U})$.

Theorem 5.22. *Let $\theta(x)$ be a consistent formula with parameters. There is a global type $p \in S_x(\mathcal{U})$ containing θ with $\mathrm{MR}(p) = \mathrm{MR}(\theta)$. If $\mathrm{MR}(\theta)$ is an ordinal, then the number of such types is exactly $\mathrm{MD}(\theta)$.*

Proof. Consider the partial type

$$\{\theta\} \cup \{\neg\varphi \mid \varphi \in F_x(\mathcal{U}) \text{ and } \mathrm{MR}(\varphi) < \mathrm{MR}(\theta)\}.$$

It suffices to show that this partial type is consistent, since no complete extension in $S_x(\mathcal{U})$ contains a formula of Morley rank smaller than $\mathrm{MR}(\theta)$.

If it is inconsistent, then by compactness there are finitely many formulas $\varphi_1, \dots, \varphi_n$ with $\mathrm{MR}(\varphi_i) < \mathrm{MR}(\theta)$ for all i such that $\{\theta, \neg\varphi_1, \dots, \neg\varphi_n\}$ is inconsistent. But then θ implies $\bigvee_{i=1}^n \varphi_i$. By Exercise 27, $\mathrm{MR}(\theta) \leq \mathrm{MR}(\bigvee_{i=1}^n \varphi_i)$. But by Lemma 5.17, $\mathrm{MR}(\bigvee_{i=1}^n \varphi_i) = \max(\mathrm{MR}(\varphi_1), \dots, \mathrm{MR}(\varphi_n)) < \mathrm{MR}(\theta)$, contradiction.

Now assume $\mathrm{MR}(\theta) = \alpha$ is an ordinal. Let $\psi_1, \dots, \psi_n$ be a decomposition of θ into irreducible components. For all $1 \leq i \leq n$, $\mathrm{MR}(\psi_i) = \alpha$, so there is a type in $S_x(\mathcal{U})$ containing ψ_i of Morley rank α. Since the ψ_i are contradictory, there are at least n such types.

In the other direction, note that any type in $S_x(\mathcal{U})$ containing φ contains ψ_i for some i. Now suppose that for some i, there are distinct types $p \neq q$ in $S_x(\mathcal{U})$, both containing ψ_i and both having Morley rank α. Let χ be a formula with $\chi \in p$ and $\neg\chi \in q$. Then $\mathrm{MR}(\psi_i \wedge \chi) \geq \alpha$ and $\mathrm{MR}(\psi_i \wedge \neg\chi) \geq \alpha$, contradicting irreducibility of ψ_i. This shows that there are at most n types of Morley rank α in $S_x(\mathcal{U})$ containing φ: one for each irreducible component. □

Lemma 5.23. *Let $p \in S_x(A)$ be a type such that $\mathrm{MR}(p)$ is an ordinal. Let $\varphi \in p$ be a type of minimal Morley rank and degree, so $\mathrm{MR}(\varphi) = \mathrm{MR}(p) = \alpha$ and $\mathrm{MD}(\varphi) = \mathrm{MD}(p) = d$. For any set $B \supseteq A$, if $q \in S_x(B)$ contains φ and $\mathrm{MR}(q) = \alpha$, then $q|_A = p$.*

Proof. Suppose for contradiction that $q|_A \neq p$. Then there is some $\psi \in F_x(A)$ such that $\psi \in q$ and $\psi \notin p$. Since $\varphi \wedge \psi \in q$, $\mathrm{MR}(\varphi \wedge \psi) = \alpha$. Since $\varphi \wedge \neg\psi \in p$, $\mathrm{MR}(\varphi \wedge \neg\psi) = \alpha$ and $\mathrm{MD}(\varphi \wedge \neg\psi) = d$. By Lemma 5.20,

$$\mathrm{MD}(\varphi) = \mathrm{MD}(\varphi \wedge \neg\psi) + \mathrm{MD}(\varphi \wedge \psi) \geq d + 1,$$

which contradicts $\mathrm{MR}(\varphi) = d$. □

Note that if $A \subseteq B$ and $p \in S_x(A)$ extends to a type $q \in S_x(B)$, then $\mathrm{MR}(q) \leq \mathrm{MR}(p)$. We say that q is a **generic extension** of p if $\mathrm{MR}(q) = \mathrm{MR}(p)$.

Corollary 5.24. *Suppose $A \subseteq B$ are sets. Every type $p \in S_x(A)$ with ordinal Morley rank has a generic extension to a type in $S_x(B)$.*

Proof. Let $\varphi \in p$ be a formula of minimal Morley rank and degree. By Theorem 5.22, there is a global type $q \in S_x(\mathcal{U})$ with $\mathrm{MR}(q) = \mathrm{MR}(\varphi) = \mathrm{MR}(p)$. By Lemma 5.23, $p \subseteq q$. Then also $p \subseteq q|_B \subseteq q$, so $q|_B \in S_x(B)$ is a generic extension of p. ☐

Corollary 5.25. *Suppose $A \subseteq B$ are sets and $p \in S_x(A)$ has ordinal Morley rank. If $G_p \subseteq S_x(B)$ is the set of generic extensions of p, then*

$$\sum_{q \in G_p} \mathrm{MD}(q) = \mathrm{MD}(p).$$

In particular, G_p is finite, and if $\mathrm{MD}(p) = 1$, then p has a unique generic extension.

Proof. First note that by Lemma 5.23 and Theorem 5.22, the number of global generic extensions of p is $\mathrm{MD}(p)$. Similarly, each $q \in G_p$ has $\mathrm{MD}(q)$ global generic extensions, each of which is a global generic extension of p. The result follows. ☐

The next exercise outlines an alternative topological way to define Morley rank and degree. This method is actually closer to Morley's original definition and can provide a useful intuition.

For any topological space X, the **Cantor-Bendixson derivative** of X is defined to be the subspace $X' = X \setminus \{x \in X \mid x \text{ is isolated in } X\}$. Define by transfinite recursion:

$$\begin{aligned} X^0 &= X \\ X^{\alpha+1} &= (X^\alpha)' \\ X^\gamma &= \bigcap_{\beta < \gamma} X^\beta \quad \text{when } \gamma \text{ is a limit ordinal.} \end{aligned}$$

The **Cantor-Bendixson rank** of a point $x \in X$ is the maximal ordinal α such that $x \in X^\alpha$, or ∞ if $x \in X^\alpha$ for all ordinals α.

So a point has Cantor-Bendixson rank 0 if it is isolated, Cantor-Bendixson rank 1 if it is isolated after removing the isolated points, etc.

Exercise 32. For $p \in S_x(\mathcal{U})$ define $\mathrm{CB}(p)$ to be the Cantor-Bendixson rank of p in the space $S_x(\mathcal{U})$. For a formula $\varphi \in F_x(\mathcal{U})$, define:

$$\mathrm{CB}(\varphi) = \sup\{\mathrm{CB}(p) \mid p \in [\varphi]\} \in \mathrm{Ord} \cup \{\pm\infty\}.$$

(a) If $\mathrm{CB}(\varphi) \neq -\infty$, show that there is some $p \in [\varphi]$ with $\mathrm{CB}(p) = \mathrm{CB}(\varphi)$ (i.e., the supremum in the definition of $\mathrm{CB}(\varphi)$ is attained). Use compactness.

(b) Show that $\mathrm{CB}(\varphi) = \mathrm{MR}(\varphi)$, and when $\mathrm{CB}(\varphi)$ is an ordinal, the number of points in $[\varphi]$ of maximal Cantor-Bendixson rank is $\mathrm{MD}(\varphi)$.

(c) For a type $p \in S_x(A)$, show that $\mathrm{MR}(p) = \sup\{\mathrm{CB}(q) \mid p \subseteq q \in S_x(\mathcal{U})\}$ and when $\mathrm{MR}(p)$ is an ordinal, $\mathrm{MD}(p)$ is the number of global extensions of p of maximal Cantor-Bendixson rank.

Example 5.26. Let T be the theory of an equivalence relation with infinitely many classes, each of which is infinite. In the context of a single variable x, $\mathrm{MR}(\top) = 2$ and $\mathrm{MD}(\top) = 1$. For any a, we have $\mathrm{MR}(xEa) = 1$ and $\mathrm{MD}(xEa) = 1$.

The type space $S_x(\mathcal{U})$ consists of: (0) The realized types, which are all isolated and have rank 0. (1) The types which contain xEa for some a but are not realized in $\mathcal{U}$ – such a type is isolated by xEa after removing all the isolated types, and has rank 1. (2) The unique type of rank 2 which contains $\neg xEa$ for all a. This is the generic global extension of the unique type in $S_x(\varnothing)$.

Example 5.27. Let T be the theory of an equivalence relation with 2 classes, each of which is infinite. In the context of a single variable x, $\mathrm{MR}(\top) = 1$ and $\mathrm{MD}(\top) = 2$. For any a, we have $\mathrm{MR}(xEa) = 1$ and $\mathrm{MD}(xEa) = 1$ (the two classes are the irreducible components of $\top$).

The type space $S_x(\mathcal{U})$ consists of: (0) The realized types, which are all isolated and have rank 0. (1) The two types which say x is in one of the two equivalence classes, but which are not realized – such a type is isolated by xEa after removing all the isolated types, and has rank 1. These two types are both generic global extension of the unique type in $S_x(\varnothing)$.

Example 5.28. Let T be the theory of an equivalence relation with infinitely many classes, each of which has size 2. In the context of a single variable x, $\mathrm{MR}(\top) = 1$ and $\mathrm{MD}(\top) = 1$. For any a, we have $\mathrm{MR}(xEa) = 0$ and $\mathrm{MD}(xEa) = 2$. The irreducible components of xEa are $x = a$ and $x = a'$, where a' is the unique element equivalent to a but not equal to a.

The type space $S_x(\mathcal{U})$ consists of: (0) The realized types, which are all isolated and have rank 0. (1) The unique type which contains $\neg xEa$ for all a. Note that if a type contains xEa, then it must contain $x = a$ or $x = a'$, where a and a' are the two elements equivalent to a.

6 Strongly minimal sets

6.1 Minimality and strong minimality

Definition 6.1. A formula (with parameters) φ is **algebraic** if $\varphi(\mathcal{U})$ is finite. A complete type $p \in S_x(A)$ is **algebraic** if it contains an algebraic formula.

We have seen that φ is algebraic if and only if $\mathrm{MR}(\varphi) = 0$ or $\mathrm{MR}(\varphi) = -\infty$ (i.e., $\varphi(\mathcal{U}) = \varnothing$). In the case $\mathrm{MR}(\varphi) = 0$, the irreducible components of φ are singletons, so $\mathrm{MD}(\varphi) = |\varphi(\mathcal{U})|$.

An algebraic type is one of Morley rank 0. Note that an algebraic type $p(x)$ is isolated by an algebraic formula, namely any formula in $p(x)$ of minimal degree.

The next simplest definable sets are those of Morley rank and degree 1.

Definition 6.2. A formula (with parameters) φ is **strongly minimal** if it is not algebraic, but for every other formula (with parameters) ψ, either $\varphi \wedge \psi$ or $\varphi \wedge \neg\psi$ is algebraic. Equivalently, φ is irreducible of Morley rank 1, so $\mathrm{MR}(\varphi) = 1$ and $\mathrm{MD}(\varphi) = 1$.

Definition 6.3. We say that T is **strongly minimal** if the formula $\top$ in context x (where x is a singleton) is strongly minimal.

Example 6.4. T_∞ is a strongly minimal theory. An atomic formula in the single variable x is equivalent to $\top$ or $\bot$ or $x = a$, hence defines a finite or cofinite set. By QE, every formula is equivalent to a Boolean combination of atomic formulas, which again defines a finite or cofinite set.

Example 6.5. VS_k is a strongly minimal theory, for exactly the same reason as T_∞.

Example 6.6. ACF_p, where p is prime or 0, is a strongly minimal theory. Just as in the previous two examples, by QE it suffices to show that every atomic formula in the single variable x defines a finite or cofinite set. Such a formula is equivalent to $\top$ or $\bot$ or $p(x) = 0$, where p is some polynomial with coefficients from $\mathcal{U}$. Since a polynomial in one variable has only finitely many roots in a field, we are done.

For another example of a strongly minimal formula in ACF_p, consider the formula $y = x^2$, call it φ. There is a definable bijection $f\colon \mathcal{U} \to \varphi(\mathcal{U})$, by $f(a) = (a, a^2)$. If there were some infinite and coinfinite definable subset $E \subseteq \varphi(\mathcal{U})$, then $f^{-1}(E)$ would be an infinite and coinfinite definable subset of $\mathcal{U}$, contradicting strong minimality of ACF_p.

On the other hand, the formula $y^2 = x^2$ is not strongly minimal, because the subset defined by $y = -x$ is infinite and coinfinite.

Example 6.7. Let T_S be the complete theory of $(\mathbb{N}; 0, S)$. T_S has quantifier elimination. An atomic formula in the single variable x is equivalent to $\top$, $\bot$, $x = S^n(a)$, or $S^n(x) = a$ for some $n \in \omega$, hence defines a finite or cofinite set. By QE, T_S is strongly minimal.

Example 6.8. The theory T_E is not strongly minimal, since every equivalence class is a definable infinite coinfinite set. But for any a, the formula xEa defining the equivalence class of a is strongly minimal. When restricted to realizations of xEa, every atomic formula in the single variable x is equivalent to $\top$, $\bot$, or $x = b$ for some bEa. By QE, xEa is strongly minimal. We say that the induced structure on xEa is that of a pure set.

Example 6.9. Let $T = \mathrm{Th}(\bigoplus_{n\in\omega} \mathbb{Z}/4\mathbb{Z})$. Consider the formula $x + x = 0$. This defines the set V of all elements of order ≤ 2, which is a vector space over $\mathbb{F}_2$. This is not enough to immediately conclude that V is strongly minimal, since the theory does not have QE and it seems we have to consider subsets of V defined by formulas with parameters outside of V. However, T is stable (as is every complete theory of an abelian group), so V is stably embedded. It follows that the induced structure on V is that of an $\mathbb{F}_2$-vector space, so V is strongly minimal.

Definition 6.10. Let $M \models T$. An infinite M-definable set $D \subseteq M$ is **minimal** if every M-definable subset of D is finite or cofinite. An infinite M-definable set $D \subseteq M$ is **strongly minimal** if the formula φ defining D is strongly minimal, i.e., if $\varphi(\mathcal{U})$ is minimal.

Example 6.11. Consider the structure $(\mathbb{N}; 0, S, \leq)$. That the definable set $\mathbb{N}$ is minimal in this structure, but not strongly minimal (so $\mathrm{Th}(\mathbb{N}; 0, S, \leq)$ is not strongly minimal).

The complete theory of this structure has quantifier elimination. An atomic formula in the single variable x is equivalent to $\top$, $\bot$, $x = S^n(a)$ or $S^n(x) = a$ for some $n \in \omega$, $x \leq a$, or $a \leq x$. When $a \in \mathbb{N}$, each of these sets is finite or cofinite, so $\mathbb{N}$ is minimal. However, when a is an element of $\mathcal{U} \setminus \mathbb{N}$, the set defined by $x \leq a$ is infinite and coinfinite, so $\mathbb{N}$ is not strongly minimal.

Example 6.12. Consider the structure (M, E), where E is an equivalence relation on M with one class of size n for each $n \in \omega \setminus \{0\}$ (and no other classes). The definable set M is minimal in this structure, but not strongly minimal (so $\mathrm{Th}(M, E)$ is not strongly minimal).

The complete theory of this structure has quantifier elimination in the expanded language where we add a unary predicate P_n naming the unique classs of size n for each n. An atomic formula with parameters from M in the single variable x in this expanded language is equivalent to $\top$, $\bot$, $x = a$, or $P_n(x)$ for some n (since xEa is equivalent to $P_n(x)$ when the class of a has n elements). Each of these sets is finite or cofinite, so M is minimal. However, when a is an element of $\mathcal{U} \setminus M$, the equivalence class of a is infinite, so the set defined by xEa is infinite and coinfinite.

Proposition 6.13. *Let $D \subseteq M^x \models T$ be a minimal set. Then D is strongly minimal if and only if for every partitioned formula $\psi(x; y)$, there exists $k \in \omega$ such that for all $b \in M^y$, $|\psi(D; b)| \leq k$ or $|D \setminus \psi(D; b)| \leq k$.*

Proof. Let D be defined by $\varphi(x; c)$. Suppose D is not strongly minimal. Then there is some formula $\psi(x; b)$ with $b \in \mathcal{U}^y$ such that $\varphi(x; c) \wedge \psi(x; b)$ and $\varphi(x; c) \wedge$

$\neg\psi(x;b)$ both define infinite subsets of $\mathcal{U}^x$. Let $k \in \omega$. Then

$$\models \exists y\,(\exists^{>k}x\,(\varphi(x;c) \wedge \psi(x;y)) \wedge \exists^{>k}x\,(\varphi(x;c) \wedge \neg\psi(x;y))).$$

Since $M \preceq \mathcal{U}$, this formula is also true in M. Letting $b' \in M^y$ be any witness to the existential quantifier, we find that $|\psi(D;b)| > k$ and $|D \setminus \psi(D;b)| > k$.

Conversely, suppose that there is some formula $\psi(x;y)$ such that for all $k \in \omega$, there exists $b \in M^y$ such that $|\psi(D;b)| > k$ and $|D \setminus \psi(D;b)| > k$. Consider the partial type:

$$\{(\exists^{>k}x\,(\varphi(x;c) \wedge \psi(x;y)) \wedge \exists^{>k}x\,(\varphi(x;c) \wedge \neg\psi(x;y))) \mid k \in \omega\}.$$

Our hypothesis (and compactness) implies that the partial type is consistent. Letting $b \in \mathcal{U}^y$ be a realization, we find that $\varphi(x;c) \wedge \psi(x;b)$ and $\varphi(x;c) \wedge \neg\psi(x;b)$ both define infinite subsets of $\mathcal{U}^x$, so $\varphi(x;c)$ is not strongly minimal. $\square$

Definition 6.14. A type $p \in S_x(A)$ is **strongly minimal** if $\mathrm{MR}(p) = 1$ and $\mathrm{MD}(p) = 1$.

If φ is a strongly minimal formula over A, then there is a unique strongly minimal type $p \in S_x(A)$ containing φ. Moreover, by Corollary 5.25, for any $A \subseteq B$, p has a unique non-algebraic extension to a type in $S_x(B)$.

Theorem 6.15. *Every strongly minimal theory T is totally transcendental.*

Proof. We show T is $\aleph_0$-stable (which suffices by Theorem 5.3). Let A be a set with $|A| = \aleph_0$. Since $x = x$ is strongly minimal, every type in $S_x(A)$ has Morley rank ≤ 1. There is a unique type of rank 1, namely the unique non-algebraic extension of the strongly minimal type in $S_x(\varnothing)$. Every other type has rank 0, and hence is algebraic and isolated by some algebraic formula in $F_x(A)$. Since $|F_x(A)| = \aleph_0$, $|S_x(A)| \leq \aleph_0 + 1 = \aleph_0$. $\square$

Later, we will prove a stronger result: Every strongly minimal theory is κ-categorical for every uncountable cardinal κ.

6.2 Algebraic closure

Definition 6.16. For a set A, we define

$$\begin{aligned}\mathrm{acl}(A) &= \{b \in \mathcal{U} \mid b \text{ satisfies an algebraic formula over } A\}\\ &= \{b \in \mathcal{U} \mid \mathrm{tp}(b/A) \text{ is algebraic}\}.\end{aligned}$$

Since the number of algebraic formulas over A is at most $\max(\aleph_0, |A|)$, and each is satisfied by only finitely many elements of $\mathcal{U}$, $|\mathrm{acl}(A)| \leq \max(\aleph_0, |A|)$.

Example 6.17. Let's look at acl in each of our standard strongly minimal theories:

1. In T_∞, $\mathrm{acl}(A) = A$.

2. In VS_k, $\operatorname{acl}(A) = \operatorname{Span}(A)$.

3. In ACF_p, with p prime or 0, $\operatorname{acl}(A)$ is the (field-theoretic) algebraic closure of the subfield generated by A. This is the reason for the name.

An important property of acl is that it does not depend on the ambient model.

Proposition 6.18. *Suppose $A \subseteq M \preceq \mathcal{U}$. Then* $\operatorname{acl}(A) \subseteq M$.

Proof. Suppose $b \in \operatorname{acl}(A)$. Then b satisfies an algebraic L_A-formula $\varphi(y)$. Say $\varphi(y)$ is satisfied by k elements of $\mathcal{U}$. Then

$$\models \exists^{=k} y\, \varphi(y),$$

so also $M \models \exists^{=k} y\, \varphi(y)$. Since $M \preceq \mathcal{U}$, the k elements of M satisfying φ also satisfy φ in $\mathcal{U}$. It follows that all k elements of $\mathcal{U}$ satisfying φ, including b, are already in M. $\square$

Proposition 6.19. *For any small set A,*

$$\operatorname{acl}(A) = \bigcap_{A \subseteq M \preceq \mathcal{U}} M.$$

Proof. By Proposition 6.18, for all models M such that $A \subseteq M \preceq \mathcal{U}$, $\operatorname{acl}(A) \subseteq M$. So

$$\operatorname{acl}(A) \subseteq \bigcap_{A \subseteq M \preceq \mathcal{U}} M.$$

For the reverse inclusion, suppose $b \in \mathcal{U} \setminus \operatorname{acl}(A)$. We would like to find a model M such that $A \subseteq M \preceq \mathcal{U}$ but $b \notin M$.

Let M' be any small elementary substructure of $\mathcal{U}$ containing A. Consider the partial type $\Sigma = p \cup \{x \neq m \mid m \in M'\}$. If Σ were inconsistent, then by compactness there would be some L_A-formula $\varphi(x) \in p$ and some $m_1, \ldots, m_n \in M'$ such that $\varphi(x) \models \bigvee_{i=1}^n x = m_i$. But then $\varphi(x)$ would be algebraic, contradiction. Let $q \in S_x(M')$ be a complete type over M' extending Σ, and let $b' \in \mathcal{U}^y$ realize q. Then $b' \notin M'$ and $\operatorname{tp}(b'/A) = \operatorname{tp}(b/A)$.

By homogeneity of $\mathcal{U}$, let $\sigma \in \operatorname{Aut}(\mathcal{U}/A)$ be an automorphism fixing A with $\sigma(b') = b$. Let $M = \sigma(M')$. Then $b \notin M$ and $A \subseteq M \preceq \mathcal{U}$, as desired. $\square$

Proposition 6.20. *Suppose A and B are small sets in $\mathcal{U}$ and $f \colon A \to B$ is a partial elementary bijection. Then f extends to a partial elementary bijection $\widehat{f} \colon \operatorname{acl}(A) \to \operatorname{acl}(B)$.*

Proof. Let $\sigma \in \operatorname{Aut}(\mathcal{U})$ be an automorphism extending f, and let $\widehat{f} = \sigma|_{\operatorname{acl}(A)}$. It remains to show that the image of $\widehat{f}$ is $\operatorname{acl}(B)$. But $c \in \operatorname{acl}(B)$ iff c satisfies an algebraic formula $\varphi(x; b)$ over B iff $\sigma^{-1}(c)$ satisfies an algebraic formula $\varphi(x; \sigma^{-1}(b))$ over A iff $c = \widehat{f}(\sigma^{-1}(c)) \in \operatorname{im}(\widehat{f})$. $\square$

Definition 6.21. Let X be a set. A **closure operator** on X is a function $\mathrm{cl}\colon \mathcal{P}(X) \to \mathcal{P}(X)$ such that for all $A, B \subseteq X$:

(1) (Reflexive) $A \subseteq \mathrm{cl}(A)$.

(2) (Monotone) If $A \subseteq B$, then $\mathrm{cl}(A) \subseteq \mathrm{cl}(B)$.

(3) (Idempotent) $\mathrm{cl}(\mathrm{cl}(A)) = \mathrm{cl}(A)$.

A closure operator is **finitary** if additionally:

(4) $\mathrm{cl}(A) = \bigcup_{F \subseteq_{\mathrm{fin}} A} \mathrm{cl}(F)$.

If cl is a closure operator on a set X, we say that $C \subseteq X$ is **closed** if $\mathrm{cl}(C) = C$.

Proposition 6.22. *For any model M,* acl *is a finitary closure operator on M.*

Proof. Reflexive: Suppose $a \in A$. Then a satisfies the algebraic formula $x = a$ over A, so $a \in \mathrm{acl}(A)$.

Monotone: Suppose $A \subseteq B$. Then every algebraic formula over A is also an algebraic formula over B, so $\mathrm{acl}(A) \subseteq \mathrm{acl}(B)$.

Idempotent: Since $A \subseteq \mathrm{acl}(A)$ we have $\mathrm{acl}(A) \subseteq \mathrm{acl}(\mathrm{acl}(A))$ by monotonicity. So it suffices to show $\mathrm{acl}(\mathrm{acl}(A)) \subseteq \mathrm{acl}(A)$. Let $c \in \mathrm{acl}(\mathrm{acl}(A))$. Then c satisfies an algebraic formula $\varphi(z, b_1, \ldots, b_n)$, where $b_1, \ldots, b_n \in \mathrm{acl}(A)$. Let $k = |\varphi(\mathcal{U}, b_1, \ldots, b_n)|$. Now since each $b_i \in \mathrm{acl}(A)$, it satisfies an algebraic formula $\psi_i(y_i)$ over A. It follows that c satisfies the formula $\theta(z)$:

$$\exists y_1 \ldots y_n \left(\varphi(z, y_1, \ldots, y_n) \wedge \exists^{\leq k} w\, \varphi(w, y_1, \ldots, y_n) \wedge \bigwedge \psi_i(y_i)\right).$$

It remains to show that $\theta(z)$ is algebraic. Since each ψ_i is algebraic, there are only finitely many tuples $b'_1, \ldots, b'_n$ witnessing the existential quantifiers. And each such witnessing tuple, there are at most k elements c' satisfying $\varphi(c', b'_1, \ldots, b'_n)$. So $\theta(z)$ is algebraic.

Finitary: In one direction, for all finite $F \subseteq A$, $\mathrm{acl}(F) \subseteq \mathrm{acl}(A)$ by monotonicity, so

$$\bigcup_{F \subseteq_{\mathrm{fin}} A} \mathrm{acl}(F) \subseteq \mathrm{acl}(A).$$

Conversely, suppose $b \in \mathrm{acl}(A)$. Then b satisfies some algebraic formula $\varphi(y)$ over A. Let F be the finitely many parameters from A appearing in $\varphi(y)$. Then $b \in \mathrm{acl}(F)$. So

$$\mathrm{acl}(A) \subseteq \bigcup_{F \subseteq_{\mathrm{fin}} A} \mathrm{acl}(F).$$ □

Exercise 33. Let cl be a closure operator on a set X.

(a) For $C \subseteq X$, define $\mathrm{cl}_C(A) = \mathrm{cl}(A \cup C)$, the **localization** of cl by C. Show that cl_C is a closure operator on X, which is finitary if cl is finitary.

(b) For $D \subseteq X$, define $\mathrm{cl}^D(A) = \mathrm{cl}(A) \cap D$, the **restriction** of cl to D. Show that cl^D is a closure operator on D, which is finitary if cl is finitary.

6.3 Pregeometries and dimension

Definition 6.23. A **pregeometry** (or **finitary matroid**) is a set X equipped with a finitary closure operator cl satisfying **exchange**: For all $C \subseteq X$ and $a, b \in X$, if $a \in \mathrm{cl}(C \cup \{b\})$ but $a \notin \mathrm{cl}(C)$, then $b \in \mathrm{cl}(C \cup \{a\})$.

Let $D \subseteq M \models T$ be a strongly minimal set, and fix a set $C \subseteq M$ over which D is definable. For $A \subseteq D$, we define $\mathrm{cl}(A) = \mathrm{acl}(A \cup C) \cap D$. Note that this is the restriction to D of the localization of acl by C, so by Exercise 33, it is a finitary closure operator on D.

Theorem 6.24. *Let D be a C-definable strongly minimal set. Then (D, cl) is a pregeometry.*

Proof. Let $\theta(x)$ be the L_C-formula defining D. Since cl is always a finitary closure operator on D, it suffices to show that cl satisfies exchange. Suppose $a \in \mathrm{cl}(A \cup \{b\})$ and $a \notin \mathrm{cl}(A)$, and assume for contradiction that $b \notin \mathrm{cl}(A \cup \{a\})$. Let $\varphi(x, b)$ be a formula over $C \cup A \cup \{b\}$ which implies $\theta(x)$ and is algebraic, say satisfied by k elements, one of which is a. Let $\psi(y)$ be the formula $\exists^{\leq k} x\, \varphi(x, y)$, and note that $\models \psi(b)$.

Since $b \notin \mathrm{cl}(A \cup \{a\})$, the formula $\varphi(a, y) \wedge \psi(y)$ is not algebraic, so its complement in D, $\theta(y) \wedge \neg(\varphi(a, y) \wedge \psi(y))$ is algebraic, say satisfied by ℓ elements. Let $\chi(x)$ be the formula $\exists^{\leq \ell} y\, (\theta(y) \wedge \neg(\varphi(x, y) \wedge \psi(y)))$, and note that $\models \chi(a)$.

Since $a \notin \mathrm{cl}(A)$, $\chi(x)$ is not algebraic. So it defines a cofinite subset of D, and we can pick $k+1$ distinct elements satisfying it, say $a_1, \dots, a_{k+1}$. For each i, there are at most ℓ elements of D satisfying $\neg(\varphi(a_i, y) \wedge \psi(y))$, so we can pick some $b' \in D$ such that $\varphi(a_i, b')$ for all $1 \leq i \leq k+1$ and also $\psi(b')$. But this contradicts the definition of $\psi(y)$. $\square$

Exercise 34. A **geometry** is a pregometry (X, cl) such that $\mathrm{cl}(\varnothing) = \varnothing$ and $\mathrm{cl}(\{a\}) = \{a\}$ for all $a \in X$.

(a) Let (X, cl) be a pregometry. Define $x \sim y$ if and only if $x \in \mathrm{cl}(\{y\})$. Show that $\sim$ defines an equivalence relation on $X \setminus \mathrm{cl}(\varnothing)$.

(b) Define $X' = (X \setminus \mathrm{cl}(\varnothing))/{\sim}$. Show that cl induces a natural closure operator cl' on X' and that (X', cl') is a geometry.

The geometry (X', cl') is called the **associated geometry** to (X, cl).

Example 6.25. (1) Let $M \models T_\infty$. Then M is a strongly minimal set, defined over $\varnothing$, and we have $\mathrm{cl}(A) = \mathrm{acl}(A) = A$ for all sets A. This is already a geometry, called the **trivial geometry**.

(2) Consider the strongly minimal theory T_S from Example 6.7. A general model $M \models T_S$ looks like $(\mathbb{N}, 0, S)$, together with some number of chains isomorphic to $(\mathbb{Z}, S)$. In this strongly minimal set, we have $\mathrm{cl}(\varnothing) = \mathbb{N}$, while for any $a \notin \mathrm{cl}(\varnothing)$, $\mathrm{cl}(\{a\})$ is the chain containing a. The associated geometry removes $\mathbb{N}$ and collapses each chain down to a single point. There are no algebraic dependencies between points in distinct chains, so this associated geometry is again the trivial geometry.

(3) Let $V \models \mathrm{VS}_k$ be a vector space over k. We have $\mathrm{cl}(\varnothing) = \{0\}$ and $\mathrm{cl}(\{v\}) = \mathrm{Span}(v)$ is the line through the origin containing the vector v. The associated geometry has as its points the set of lines through the origin in V. This geometry is denoted $\mathbb{P}(V)$, the projective space associated to V.

Zilber famously conjectured that for every strongly minimal set D, the geometry associated to (D, cl) is either:

(1) The trivial geometry ($\mathrm{cl}(A) = A$ for all A).

(2) The geometry of an affine or projective space over a division ring.

(3) The geometry of algebraic independence in an algebraically closed field.

This conjecture was refuted by Hrushovski, who used a modification of the Fraïssé limit construction to produce a strongly minimal set whose associated geometry does not fit into this trichotomy.

Definition 6.26. Let (X, cl) be a pregeometry. A subset $B \subseteq X$ is:

- **independent** if for all $b \in B$, $b \notin \mathrm{cl}(B \setminus \{b\})$.
- **a generating set** if $\mathrm{cl}(B) = X$.
- **a basis** if it is an independent generating set.

Suppose $B \subseteq B'$. If B' is independent, B is independent, since for $b \in B$, $b \in \mathrm{cl}(B \setminus \{b\})$ implies $b \in \mathrm{cl}(B' \setminus \{b\})$. If B is a generating set, then B' is a generating set, since $X = \mathrm{cl}(B) \subseteq \mathrm{cl}(B') \subseteq X$.

Lemma 6.27. *Suppose I is an independent set in the pregeometry X. For any $a \in X \setminus \mathrm{cl}(I)$, $I \cup \{a\}$ is independent.*

Proof. Let $I' = I \cup \{a\}$. Suppose for contradiction that I' is not independent. Then there is some $b \in I'$ such that $b \in \mathrm{cl}(I' \setminus \{b\})$. If $b = a$, then $a \in \mathrm{cl}(I)$, contradiction. If $b \neq a$, let $J = I \setminus \{b\}$, so $b \in \mathrm{cl}(J \cup \{a\})$. Since I is independent, $b \notin \mathrm{cl}(J)$, so by exchange $a \in \mathrm{cl}(J \cup \{b\}) = \mathrm{cl}(I)$, contradiction. □

Theorem 6.28. *Suppose B is an independent set in the pregeometry X. Then B can be extended to a basis for X.*

Proof. Apply Zorn's Lemma to the poset of independent subsets of X containing B. If $(B_i)_{i \in I}$ is a chain of independent subsets of X, we must show that $B^* = \bigcup_{i \in I} B_i$ is independent. Let $b \in B^*$. If $b \in \mathrm{cl}(B^* \setminus \{b\})$, then by local finiteness $b \in \mathrm{cl}(\{b_1, \ldots, b_n\})$ for some $b_1, \ldots, b_n \in B^* \setminus \{b\}$. Then there is some $i \in I$ such that $b, b_1, \ldots, b_n \in B_i$, so $b \in \mathrm{cl}(B_i \setminus \{b\})$, contradicting the fact that B_i is independent.

Let M be a maximal independent subset of X containing B. By Lemma 6.27, if there is some $x \in X \setminus \mathrm{cl}(M)$, then $M \cup \{x\}$ is independent, contradicting maximality. So $\mathrm{cl}(M) = X$, and hence M is a basis for X. □

Theorem 6.29. *Let X be a pregeometry. If $A \subseteq X$ is independent and $C \subseteq X$ is a generating set, then $|A| \leq |C|$.*

Proof. Enumerate $A = (a_\alpha)_{\alpha<\kappa}$ (note $\kappa = |A|$ may be finite!). By transfinite induction, we define elements $(c_\alpha)_{\alpha<\kappa}$ from C and prove that that for all $\beta \leq \kappa$, $A_\beta = \{c_\alpha \mid \alpha < \beta\} \cup \{a_\alpha \mid \beta \leq \alpha < \kappa\}$ is independent. Note that $A_{\beta+1} = (A_\beta \setminus \{a_\beta\}) \cup c_\beta$. So we define c_β at the successor step $\gamma = \beta + 1$. At the end of the induction, we have distinct elements $(c_\alpha)_{\alpha<\kappa}$ from C, so $|A| = \kappa \leq |C|$.

Base case: When $\beta = 0$, $A_0 = A$ is independent.

Limit step: When β is a limit, suppose for contradiction that A_β is not independent. By local finiteness, some finite set is dependent, say

$$c_{\alpha_1}, \ldots, c_{\alpha_n}, a_{\alpha_{n+1}}, \ldots, a_{\alpha_m}$$

with $\alpha_1 < \cdots < \alpha_n < \beta \leq \alpha_{n+1} < \cdots < \alpha_m$. But these elements are all in A_{α_n+1} with $\alpha_n + 1 < \beta$, which is independent by induction.

Successor step: When $\beta = \alpha + 1$, we must define c_α. I first claim that $C \not\subseteq \mathrm{cl}(A_\alpha \setminus \{a_\alpha\})$. If not, we would have $a_\alpha \in X = \mathrm{cl}(C) \subseteq \mathrm{cl}(A_\alpha \setminus \{a_\alpha\})$, contradicting independence of A_α. So we can pick any $c_\alpha \in C \setminus \mathrm{cl}(A_\alpha \setminus \{a_\alpha\})$. Note that since $A_\alpha \setminus \{a_\alpha\} \subseteq A_\alpha$, the set $A_\alpha \setminus \{a_\alpha\}$ is independent. By Lemma 6.27, $A_{\alpha+1} = (A_\alpha \setminus \{a_\alpha\}) \cup \{c_\alpha\}$ is independent. □

Corollary 6.30. *If B and B' are bases for the pregeometry X, then $|B| = |B'|$.*

Proof. Since B is independent and B' is a generating set, $|B| \leq |B'|$. Since B' is independent and B is a generating set, $|B'| \leq |B|$. So $|B| = |B'|$. □

Note that by Theorem 6.28, every pregeometry X has a basis (by extending the independent set $\varnothing$ to a basis), and by Corollary 6.30, every basis has the same cardinality. We define $\dim(X)$, the **dimension** of X, to be the cardinality of any basis.

With a well-defined notion of dimension in hand, we return to the context of strongly minimal sets. The first thing to notice is that pregeometries arising from strongly minimal sets are homogeneous: any two independent tuples of the same length have the same type.

Proposition 6.31. *Let $D \subseteq M \models T$ be a strongly minimal set, defined over C. For every n, there is a type $p_n(x_1, \ldots, x_n) \in S_n(C)$ which is satisfied by any independent tuple of size n from D. Any independent set of elements of D is a C-indiscernible set.*

Proof. We proceed by induction on n. When $n = 1$, a single element b_1 is independent if $b_1 \notin \mathrm{cl}(\varnothing) = \mathrm{acl}(C)$. So $\mathrm{tp}(b_1/C)$ is the unique non-algebraic type in $S_1(C)$.

Now consider independent tuples $b_1, \ldots, b_n$ and $b'_1, \ldots, b'_n$. By the inductive hypothesis, $\mathrm{tp}(b_1, \ldots, b_{n-1}/C) = \mathrm{tp}(b'_1, \ldots, b'_{n-1}/C) = p_{n-1}$, so the map $f(b_i) = b'_i$ is a partial elementary map. It remains to show that

$$\mathrm{tp}(b'_n/Cb'_1 \ldots b'_{n-1}) = f_*\mathrm{tp}(b_n/Cb_1 \ldots b_{n-1})$$

so that $\operatorname{tp}(b_1, \ldots, b_n/C) = \operatorname{tp}(b'_1, \ldots, b'_n/C) = p_n$.

Now $b_n \notin \operatorname{cl}(b_1, \ldots, b_{n-1}) = \operatorname{acl}(Cb_1 \ldots b_{n-1})$, so $\operatorname{tp}(b_n/Cb_1 \ldots b_{n-1})$ is the unique non-algebraic type in $S_1(Cb_1 \ldots b_{n-1})$. Similarly, $\operatorname{tp}(b'_n/Cb'_1 \ldots b'_{n-1})$ is the unique non-algebraic type in $S_1(Cb'_1 \ldots b'_{n-1})$. But since f is partial elementary, an L_C-formula $\varphi(x, b_1, \ldots, b_n)$ is algebraic if and only if $\varphi(x, b'_1, \ldots, b'_n)$ is algebraic. So $f_*\colon S_1(Cb_1 \ldots b_{n-1}) \to S_1(Cb'_1 \ldots b'_{n-1})$ maps the unique non-algebraic type to the unique non-algebraic type.

Let $I \subseteq D$ be an independent set. Since any two tuples of distinct elements from I have the same type over C, I is a C-indiscernible set. $\square$

The next theorem illustrates the idea of classifying models by the dimension of a strongly minimal set.

Theorem 6.32. *If T is a strongly minimal theory, then T is κ-categorical for every uncountable cardinal κ.*

Proof. Let κ be an uncountable model, and let M and N be two models of T of cardinality κ. Since T is strongly minimal, M itself is a strongly minimal set. Let $B_M \subseteq M$ be a basis. Then $M = \operatorname{acl}(B_M)$, so $\kappa = |M| = |\operatorname{acl}(B_M)| = \max(\aleph_0, |B_M|)$, and thus $|B_M| = \kappa$ (here we use that $\kappa > \aleph_0$). Similarly, let $B_N \subseteq N$ be a basis. By the same argument, $|B_N| = \kappa$.

Let $f\colon B_M \to B_N$ be any bijection. Then f is partial elementary, since any tuple $b_1, \ldots, b_n$ of distinct elements from B_M satisfies the same type p_n (from Proposition 6.31) as $f(b_1), \ldots, f(b_n)$. By Proposition 6.20, f extends to a partial elementary bijection $\widehat{f}\colon \operatorname{acl}(B_M) \to \operatorname{acl}(B_N)$. But $\operatorname{acl}(B_M) = M$ and $\operatorname{acl}(B_N) = N$, and a partial elementary bijection between two models is an isomorphism. $\square$

7 The theorems of Morley and Baldwin–Lachlan

In Theorem 6.32, we saw that every strongly minimal theory T is κ-categorical for every uncountable cardinal κ, because models of T can be classified up to isomorphism by their dimension (which is equal to their cardinality in the uncountable case). Fixing an uncountable κ, we would like to extend this classification result to arbitrary κ-categorical theories, not all of which are strongly minimal. In the non-strongly-minimal case, there are two major obstacles. I will now give a brief overview of these obstacles and a sketch of how we will overcome them.

First, since T is totally transcendental, there must be a strongly minimal formula $\varphi(x; c)$, possibly with parameters $c \in \mathcal{U}^y$. If a model $M \models T$ contains c (or even any realization of $\text{tp}(c/\varnothing)$), then $\varphi(M; c)$ is a strongly minimal set, and we can assign a "φ-dimension" to M. But if $\text{tp}(c/\varnothing)$ is not isolated, there will be models of T which omit this type. We will show that when T is κ-categorical, there is a strongly minimal formula with parameters in the prime model of T. Since the prime model embeds elementarily in every model of T, we can assume that every model contains these parameters and use this fixed strongly minimal formula $\varphi(x)$ to assign a φ-dimension to every model of T.

Second, suppose we have two models M and N of the same φ-dimension. Writing $D_M = \varphi(M; c)$ and $D_N = \varphi(N; c)$ for the strongly minimal sets, we obtain a partial elementary bijection $f\colon D_M \to D_N$, which can be extended to a partial elementary bijection $\text{acl}(D_M) \to \text{acl}(D_N)$. If every model of T is contained in the algebraic closure of its strongly minimal set, we are done (in this case, we say T is **almost strongly minimal**). But in general, we need to work harder. An example of an uncountably categorical theory which is not almost strongly minimal is the theory of the abelian group $\bigoplus_{n\in\omega} \mathbb{Z}/4\mathbb{Z}$ from Example 6.9. Another is the theory of the ring of dual numbers $\mathbb{C}[x]/(x^2)$ over the complex numbers.

For the general case, we recall that since T is totally transcendental, M contains a prime model P over D_M. Identifying P with its image under an elementary embedding $P \to M$, we have $D_M \subseteq P \subseteq M$. By primeness, the partial elementary bijection $f\colon D_M \to D_N$ extends to an elementary embedding $e\colon P \to N$. So we have $D_N \subseteq e(P) \subseteq N$.

Note that $D_M = \varphi(P; c) = \varphi(M; c)$, so if M is a proper elementary extension of P, then it is a proper elementary extension which does not add any new elements satisfying the formula $\varphi(x; c)$. Such a situation, $P \prec M$ with $\varphi(P; c) = \varphi(M; c)$, is called a **Vaughtian pair**. We will show that κ-categorical theories have no Vaughtian pairs. It follows that $M = P$ and by the same argument $N = e(P)$, so e is the desired isomorphism $M \cong N$.

To sum up: For a κ-categorical theory T, we will show that there is a strongly minimal formula $\varphi(x)$ with parameters in the prime model, and every model $M \models T$ is prime and minimal over $\varphi(M)$ (i.e., has no proper elementary substructure containing $\varphi(M)$, by the prohibition on Vaughtian pairs). It follows that if the strongly minimal set defined by $\varphi(x)$ has the same dimension in M and N, then $M \cong N$.

7.1 Vaughtian pairs and (κ, λ)-models

What is the relationship between the cardinality of a model and the cardinalities of definable sets in that model? The Löwenheim–Skolem theorem gives us essentially the maximum flexibility in the cardinalities of infinite models. Perhaps surprisingly, it is much harder in general to control the cardinalities of infinite definable sets.

Definition 7.1. A model $M \models T$ is **crowded** if for every non-algebraic L_M-formula $\varphi(x)$, $|\varphi(M)| = |M|$.

Note that every countable model is crowded, since every definable set must be countable. Also, every saturated model is crowded: if we have $\aleph_0 \leq |\varphi(M)| < |M|$, then the type $\{\varphi(x)\} \cup \{x \neq m \mid m \in \varphi(M)\}$ is a consistent partial type over $\varphi(M)$ which is not realized in M, so M is not saturated. But unlike saturated models, we can easily build crowded models of cardinality κ, with no hypotheses on T or κ.

Lemma 7.2. *For every infinite κ, T has a crowded model of cardinality κ.*

Proof. We build an elementary chain $(M_n)_{n\in\omega}$ of models of cardinality κ, such that every non-algebraic L_{M_n}-definable set has size κ in M_{n+1}. Let M_0 be any model of cardinality κ. Given M_n, let $\mathcal{F}$ be the set of non-algebraic L_{M_n}-formulas. For each $\varphi(x) \in \mathcal{F}$, introduce κ-many new tuples of constants $(c^\varphi_\alpha)_{\alpha<\kappa}$, each of length $|x|$. Consider the theory

$$T_{n+1} = \mathrm{EDiag}(M_n) \cup \{c^\varphi_\alpha \neq c^\varphi_\beta \mid \varphi \in \mathcal{F}, \alpha < \beta < \kappa\} \cup \{\varphi(c^\varphi_\alpha) \mid \varphi \in \mathcal{F}, \alpha < \kappa\}.$$

This is consistent by compactness, since each $\varphi \in \mathcal{F}$ is non-algebraic. Since $|M_i| = \kappa$, there are κ-many constants naming the elements of M_i, $|\mathcal{F}| = \kappa$, and for each formula $\varphi(x) \in \mathcal{F}$ there are κ-many new constants, so in total the language of T_{n+1} has size κ. Thus we can find a model $M_{n+1} \models T_{n+1}$ with $|M_{n+1}| = \kappa$. Let $N = \bigcup_{n\in\omega} M_n$. Then $M \preceq N$, $|N| = \kappa$, and N is crowded. For any non-algebraic L_N-formula $\varphi(x)$, the parameters in φ occur already in some $M_n \preceq N$, and $\varphi(x)$ is non-algebraic in M_n. Then $|\varphi(M_{n+1})| = \kappa$ and $\varphi(M_{n+1}) \subseteq \varphi(N) \subseteq N^x$, so $|\varphi(N)| = \kappa$. □

It follows that if T is κ-categorical, then every model of cardinality κ is crowded. The following definition quantifies failures of crowdedness.

Definition 7.3. Let $\kappa > \lambda$ be infinite cardinals. A (κ, λ)-model is a pair $(M, \varphi(x))$, where $M \models T$, $\varphi(x)$ is an L_M-formula, $|M| = \kappa$, and $|\varphi(M)| = \lambda$.

Robert Vaught showed that the existence of a non-crowded model (equivalently, a (κ, λ)-model for some $\kappa > \lambda$), is equivalent to a condition which is "set theory free" in the sense that it doesn't mention cardinals.

Definition 7.4. A **Vaughtian pair** is a triple $(M, N, \varphi(x))$, where $M \prec N$ and $\varphi(x)$ is a non-algebraic L_M-formula such that $\varphi(M) = \varphi(N)$.

Theorem 7.5 (Vaught's two-cardinal theorem)**.** *The following are equivalent:*

(1) T has a Vaughtian pair.

(2) T has a (κ, λ)-model for some infinite cardinals $\kappa > \lambda$.

(3) T has an $(\aleph_1, \aleph_0)$-model.

Moreover, the same formula can be used as a witness in all three conditions.

We aim to prove this theorem, along with the following extension due to Alistair Lachlan.

Theorem 7.6. *Let T be totally transcendental. If T has a Vaughtian pair, then for every uncountable cardinal κ, T has a $(\kappa, \aleph_0)$-model (using the same formula as a witness).*

Before proceeding with the proofs of these theorems, I'll show how they imply our result that κ-categorical theories have no Vaughtian pairs, and I'll give some further discussion.

Corollary 7.7. *Let κ be an uncountable cardinal. If T is κ-categorical, then T has no Vaughtian pairs.*

Proof. Suppose T is κ-categorical. By Corollary 4.21 (and Theorem 5.3), T is totally transcendental. Assume for contradiction that T has a Vaughtian pair. By Theorem 7.6, T has a $(\kappa, \aleph_0)$-model. In particular, M is not crowded. But by Lemma 7.2, T has a crowded model N of cardinality κ. Then $M \not\cong N$, contradicting κ-categoricity. $\square$

Theorem 7.6 is sufficient for our purposes, but Lachlan later proved a more general result, with weaker hypotheses and a stronger conclusion.

Theorem 7.8 (Lachlan's two-cardinal theorem)**.** *Suppose T is stable. The following are equivalent:*

(1) T has a Vaughtian pair.

(2) T has a (κ, λ)-model for some infinite cardinals $\kappa > \lambda$.

(3) T has a (κ, λ)-model for all infinite cardinals $\kappa > \lambda$.

Moreover, the same formula can be used as a witness in all three conditions.

The proof of Theorem 7.8 uses more technical tools from stability theory than we have developed so far, and we don't need it for the proof of Morley's theorem, so I will omit it. But I'd like to give an example showing that some hypothesis on T (e.g., stability) are necessary if we want to get (κ, λ)-models with κ and λ "far apart".

Example 7.9. Consider the language $L = \{S, P, \in\}$, where S and P are unary predicates and $\in$ is a binary relation. Let $M = \omega \sqcup \mathcal{P}(\omega)$, where $S^M = \omega$, $P^M = \mathcal{P}(\omega)$, and $\in$ is the usual elementhood relation between ω and $\mathcal{P}(\omega)$. Let $T = \mathrm{Th}(M)$, and note that T contains the sentence asserting "extensionality":

$$\forall y \forall z((P(y) \land P(z) \land \forall x\,(x \in y \leftrightarrow x \in z)) \to y = z)$$

The standard model M, together with the formula $S(x)$, is a $(2^{\aleph_0}, \aleph_0)$-model of T. Let $N \prec M$ be an elementary substructure of cardinality $\aleph_1$ containing ω. This is an $(\aleph_1, \aleph_0)$-model, as promised by Vaught's two-cardinal theorem, and $(N, M, S(x))$ is a Vaughtian pair. But T has no (κ, λ)-model witnessed by $S(x)$ for $\kappa > 2^\lambda$: if $|S(N)| = \lambda$, then by extensionality $|P(N)| \leq 2^\lambda$, and $|N| = |S(N)| + |P(N)| \leq 2^\lambda$.

Note that T is unstable, so this example does not contradict Lachlan's two-cardinal theorem. For example, in any model $N \models T$ with $|S(N)| = \kappa$, there are 2^κ-many complete types over $S(N)$ containing $P(x)$, one for every "real" subset of N.

Now we will proceed with the proof of Theorems 7.5 and 7.6. We first show that we can capture the notion of a Vaughtian pair for T by a first-order theory.

Definition 7.10. Fix a partitioned formula $\varphi(x; y)$. Let L^φ_{VP} be the language extending L by a new unary predicate P and a new tuple of constant symbols $c = (c_1, \ldots, c_{|y|})$ of length $|y|$. Define the L^φ_{VP}-theory T^φ_{VP} extending T by new axioms:

(1) P is an L-elementary substructure. By the Tarski–Vaught test, we can express this with the following axioms, one for each L-formula $\psi(w, z)$ with w a singleton:

$$\forall z \left(\bigwedge_{i=1}^{|z|} P(z_i) \to ((\exists w\, \psi(w, z)) \to \exists w\, (P(w) \land \psi(w, z))) \right).$$

(2) P is a proper elementary substructure: $\exists w\, \neg P(w)$.

(3) The parameters are in P: $\bigwedge_{i=1}^{|y|} P(c_i)$.

(4) P contains the set defined by $\varphi(x; c)$: $\forall x\, (\varphi(x; c) \to \bigwedge_{i=1}^{|x|} P(x_i))$.

(5) The set $\varphi(x; c)$ is infinite: for all $n \in \omega$, $\exists^{\geq n} x\, \varphi(x; c)$.

We have $(M; N, a) \models T^\varphi_{\mathrm{VP}}$ if and only if $N \prec M$, $a \in N^y$, and $\varphi(N, a) = \varphi(M, a)$ is infinite. That is, a model of T^φ_{VP} is exactly a Vaughtian pair for T witnessed by an instance of φ.

If T has a Vaughtian pair, the T^φ_{VP} is consistent for some $\varphi(x; y)$. Working with models of T^φ_{VP}, we will obtain a particularly nice Vaughtian pair, in which the models are countable and sufficiently homogeneous.

Definition 7.11. A model M is **$\aleph_0$-homogeneous** if for every partial elementary map $f\colon A \to M$ with $A \subseteq M$ finite, and every $b \in M$, f extends to a partial elementary map $A \cup \{b\} \to M$.

Lemma 7.12. *Suppose M and N are countable models of T which are $\aleph_0$-homogeneous and realize exactly the same types in $S_n(\varnothing)$ for all $n \in \omega$. Then any partial elementary map $f\colon A \to N$, where A is a finite subset of M, can be extended to an isomorphism $M \cong N$.*

Proof. Enumerate $M = (m_i)_{i\in\omega}$ and $N = (n_i)_{i\in\omega}$. We define (by back-and-forth) a sequence of partial elementary maps $(f_i)_{i\in\omega}$ extending f such that $\{m_j \mid j < i\} \subseteq \mathrm{dom}(f_i)$ and $\{n_j \mid j < i\} \subseteq \mathrm{ran}(f_i)$ for all i. Then $f_\omega = \bigcup_{i\in\omega} f_i$ is a partial elementary map extending f with $M \subseteq \mathrm{dom}(f_\omega)$ and $N \subseteq \mathrm{ran}(f_\omega)$, i.e., an isomorphism $M \cong N$.

Let $f_0 = f$. Given f_i, let $A = \mathrm{dom}(f_i)$, and write $A = \{a_0, \ldots, a_k\}$. Similarly, let $B = \mathrm{ran}(f_i)$, and write $B = \{b_0, \ldots, b_k\}$ with $f(a_j) = b_j$ for all j. Let $p_i = \mathrm{tp}(a_0, \ldots, a_k, m_i)$. Since M and N realize the same types, p_i is realized by some $b'_0, \ldots, b'_k, b'_{k+1} \in N$, so the map h defined by $a_j \mapsto b'_j$ and $m_i \mapsto b'_{k+1}$ is partial elementary.

Now
$$\mathrm{tp}(b'_0, \ldots, b'_k) = \mathrm{tp}(a_0, \ldots, a_k) = \mathrm{tp}(b_0, \ldots, b_k),$$
since f is partial elementary. So the map g defined by $b'_j \mapsto b_j$ for $j \le k$ is partial elementary. Since N is $\aleph_0$-homogeneous, g extends to a partial elementary map g' which includes b'_{k+1} in its domain, say $b'_{k+1} \mapsto b_{k+1}$. Define $f'_i = f_i \cup \{(m_i, b_{k+1})\}$. This is partial elementary because $f'_i = g' \circ h$ is the composition of two partial elementary maps.

By a symmetric argument, this time using $\aleph_0$-homogeneity of M, we can extend f'_i to a partial elementary map f_{i+1} which includes n_i in its range. $\square$

Lemma 7.13. *Suppose T has a Vaughtian pair. Then T has a Vaughtian pair $(N, M, \varphi(x))$ such that N and M are countable and $\aleph_0$-homogeneous and realize the same types over $\varnothing$.*

Proof. If T has a Vaughtian pair, then T^φ_{VP} is consistent. By Löwenheim–Skolem, T^φ_{VP} has a countable model, which is a Vaughtian pair $(N, M, \varphi(x; b))$ with N and M countable.

Write M_0 for this Vaughtian pair, viewed as a model of T^φ_{VP}. We will build an elementary chain $(M_n)_{n\in\omega}$ of countable models of T^φ_{VP}. We will alternate steps of the construction. At odd stages, we will ensure that $P(M_{n+1})$ realizes every L-type over $\varnothing$ realized in M_n. At even stages, we will enforce instances of $\aleph_0$-homogeneity for M_n and $P(M_n)$.

Odd stages: We are given M_n with n even. Let X be the set of all L-types in $\bigcup_{n\in\omega} S^L_n(\varnothing)$ realized in M_n. Since M_n is countable, it contains countably many finite tuples, so X is countable. For each type $p(z) \in X$, consider the L^φ_{VP}-type $p'(z) = p(z) \cup \{\bigwedge_{i=1}^{|z|} P(z_i)\}$. Since every formula in $p(z)$ is satisfiable in M_n and $P(M_n) \prec M_n$, every formula in $p(z)$ is satisfiable in $P(M_n)$, and hence $p'(z)$ is consistent by compactness. Thus we can find a countable elementary extension $M_n \prec M_{n+1}$ realizing $p'(z)$ for all $p(z) \in X$.

Even stages: We are given M_n with n odd. Let Y be the set of all triples (A, f, a), where A is a finite subset of M_n, $f\colon A \to M_n$ is a partial L-elementary map, and $a \in M_n$. Since M_n is countable, Y is countable. For all $(A, f, a) \in Y$, let $p_{A,f,a}(x)$ be $f_*\mathrm{tp}_L(a/A)$. If $A \subseteq P(M_n)$, $f(A) \subseteq P(M_n)$, and $a \in P(M_n)$, include the formula $P(x)$ in $p_{A,f,a}(x)$. Consistency of $f_*\mathrm{tp}_L(a/A)$ is clear, since the pushforward of a consistent type is consistent. In the case where we include $P(x)$, note that for every formula $\psi(x) \in \mathrm{tp}(a/A)$, $f_*\psi(x)$ is satisfiable in $P(M_n)$ since f is partial elementary, so $f_*\psi(x) \wedge P(x)$ is satisfiable in M_n, so $p_{A,f,a}(x)$ is consistent by compactness. Thus we can find a countable elementary extension $M_n \prec M_{n+1}$ realizing $p_{A,f,a}(x)$ for all $(A, f, a) \in Y$.

Now let $M_\omega = \bigcup_{n\in\omega} M_n$. Then M_ω is countable. Write N_ω for $P(M_\omega)$. Since $M_\omega \models T^{\varphi}_{\mathrm{VP}}$, $(N_\omega, M_\omega, \varphi(x; b))$ is a Vaughtian pair of countable models of T. I claim that N_ω and M_ω realize the same types over $\varnothing$. Since $N_\omega \prec M_\omega$, every type realized in N_ω is realized in M_ω. Conversely, let a be a finite tuple in M_ω. Then $a \in M_n$ for sufficiently large even n, and $\mathrm{tp}(a)$ is realized in $P(M_{n+1})$ and hence in N_ω.

It remains to show that N_ω and M_ω are $\aleph_0$-homogeneous. Let A be a finite subset of M_ω, $f\colon A \to M_\omega$ a partial elementary map and $a \in M_\omega$. Then the finite set $A \cup f(A) \cup \{a\}$ is contained in M_n for sufficiently large odd n, and $f_*\mathrm{tp}(a/A)$ is realized in M_{n+1}, and hence in M_ω, say by a'. Then $f \cup \{(a, a')\}$ is a partial elementary map $A \cup \{a\} \to M_\omega$, so M_ω is $\aleph_0$-homogeneous. For N_ω, the same argument works, noting that our construction ensures that $f_*\mathrm{tp}(a/A)$ is realized in $P(M_{n+1})$, and hence in N_ω. □

Proof of Theorem 7.5. (3)⇒(2): Trivial.

(2)⇒(1): Suppose $(M, \varphi(x))$ is a (κ, λ)-model for some infinite $\kappa > \lambda$. By Löwenheim–Skolem, let $N \prec M$ be an elementary substructure containing $\varphi(M)$ with $|N| = \lambda$. Then $(N, M, \varphi(x))$ is a Vaughtian pair.

(1)⇒(3): Suppose T has a Vaughtian pair. By Lemma 7.13, T has a Vaughtian pair $(N, M, \varphi(x))$ such that N and M are countable and $\aleph_0$-homogeneous and realize the same types over $\varnothing$. Note that by Lemma 7.12, $N \cong M$, but this isomorphism is not the inclusion (which is not surjective).

We build an elementary chain $(N_\alpha)_{\alpha<\aleph_1}$ such that for all α, N_α is countable and $\aleph_0$-homogeneous and realizes the same types as N over $\varnothing$, and such that $\varphi(N_\alpha) = \varphi(N)$.

Let $N_0 = N$. The conditions are trivially satisfied.

For the successor step, write A for the finite set of parameters from N appearing in $\varphi(x)$. Since $N \preceq N_\alpha$, the identity map $A \to N_\alpha$ extends to an isomorphism $N \cong N_\alpha$ by Lemma 7.12. Then N_α has a proper elementary extension $N_\alpha \prec N_{\alpha+1}$ (isomorphic to the extension $N \prec M$) such that $N_{\alpha+1}$ is countable and $\aleph_0$-homogeneous and realizes the same types over $\varnothing$ as N_α (and hence the same types as N), and such that $\varphi(N_{\alpha+1}) = \varphi(N_\alpha) = \varphi(N)$.

When $\gamma < \aleph_1$ is a limit ordinal, let $N_\gamma = \bigcup_{\alpha<\gamma} N_\alpha$. Since γ is countable, N_γ is countable. We have $\varphi(N_\gamma) = \bigcup_{\alpha<\gamma} \varphi(N_\alpha) = \varphi(N)$. Since $N \prec N_\gamma$, every type realized in N is realized in N_γ. Conversely, any finite tuple b from N_γ appears already in some N_α for $\alpha < \gamma$ and has the same type in N_γ as in N_α. Since N_α and N realize the same types, $\mathrm{tp}(b)$ is realized in N.

It remains to show that N_γ is $\aleph_0$-homogeneous. Suppose A is a finite subset of N_γ, $f\colon A \to N_\gamma$ is partial elementary, and $b \in N_\gamma$. Pick $\alpha < \gamma$ sufficiently large so that N_α contains the finite set $\mathrm{dom}(f) \cup \mathrm{ran}(f) \cup \{b\}$. Then since N_α is $\aleph_0$-homogeneous, $f\colon A \to N_\alpha$ extends to a partial elementary map $f\colon A \cup \{b\} \to N_\alpha$. This is sufficient, since $N_\alpha \prec N_\gamma$. □

The elementary chain argument in the proof of Vaught's two-cardinal theorem does not allow us to extend past $\aleph_1$, since the model of size $\aleph_1$ constructed in the proof is no longer isomorphic to the base model N in the Vaughtian pair. To obtain a $(\kappa, \aleph_0)$-model, we will use the same elementary chain idea, but we need a new way to find a proper elementary extension which does not add any new elements satisfying $\varphi(x)$.

Definition 7.14. Let $M \models T$. A partial type p is **countably satisfiable** in M if for all $\Sigma \subseteq p$ with $|\Sigma| \leq \aleph_0$, Σ is realized in M.

Recall that a partial type over M is consistent if and only if it is *finitely satisfiable* in M. Indeed, finitely satisfiable in M implies consistent by compactness. Conversely, if p is consistent, then for any $\varphi_1(x), \ldots, \varphi_n(x) \in p(x)$, the conjunction $\bigwedge_{i=1}^n \varphi_i(x)$ is consistent, so $\mathcal{U} \models \exists x \bigwedge_{i=1}^n \varphi_i(x)$, and thus also $M \models \exists x \bigwedge_{i=1}^n \varphi_i(x)$. Countably satisfiable is a natural strengthening of this condition.

Lemma 7.15. *Let T be totally transcendental. If $M \models T$ is uncountable, then there is a proper elementary extension $M \prec N$ such that every complete type over M realized in N is countably satisfiable in M.*

Proof. Let's say an L_M-formula $\varphi(x)$ is large if $\varphi(M)$ is uncountable. Otherwise, $\varphi(x)$ is small. I claim that there is an L_M-formula $\theta(x)$ in a single variable x which is minimally large in the sense that θ is large and for any other L_M-formula $\psi(x)$, either $\theta \wedge \psi$ or $\theta \wedge \neg\psi$ is small.

Suppose for contradiction that there is no minimally large L_M-formula. We build a binary tree of large formulas $(\varphi_g(x))_{g \in 2^{<\omega}}$, contradicting T totally transcendental. Let $\varphi_\varnothing$ be $\top$ in context x. This is large, since M is uncountable. Given $\varphi_g(x)$, since $\varphi_g(x)$ is not minimally large, there is some formula $\psi(x)$ such that $\varphi_g \wedge \psi$ and $\varphi_g \wedge \neg\psi$ are both large. Let $\varphi_{g0} = \varphi_g \wedge \psi$ and let $\varphi_{g1} = \varphi_g \wedge \neg\psi$.

Now fix some minimally large L_M-formula $\theta(x)$, and define:

$$p = \{\psi(x) \in F_x(M) \mid \theta \wedge \psi \text{ is large}\}.$$

I claim that this type is countably satisfiable in M. Indeed, consider a countable $\Sigma \subseteq p$, so $\theta \wedge \sigma$ is large for all $\sigma \in \Sigma$. Since θ is minimally large, $\theta \wedge \neg\sigma$ is small, i.e., $(\theta \wedge \neg\sigma)(M)$ is countable for all $\sigma \in \Sigma$. So $B = \bigcup_{\sigma \in \Sigma} (\theta \wedge \neg\sigma)(M)$ is countable. Since $\theta(M)$ is uncountable, $\theta(M) \setminus B$ is non-empty. Any element of this set satisfies Σ. It follows that p is consistent, and p is complete because $\theta(x)$ is minimally large.

Now let $a \in \mathcal{U}$ satisfy p. Note that $a \notin M$: for all $m \in M$, $\theta \wedge (x \neq m)$ is large, so $(x \neq m) \in p$. By Corollary 5.12, since T is totally transcendental, T has a model N which is prime and atomic over $M \cup \{a\}$. Then $M \prec N$, and it remains to show that every complete type over M realized in N is countably satisfiable in M.

Let $b \in N^y$ and let Δ be a countable subset of $\mathrm{tp}(b/M)$. Since N is atomic over $M \cup \{a\}$, $\mathrm{tp}(b/M \cup \{a\})$ is isolated, say by $\chi(y, a)$ (where $\chi(y, x)$ is an L_M-formula). Now $\exists y\, \chi(y, x) \in \mathrm{tp}(a/M) = p(x)$, and for all $\delta(y) \in \Delta$, since $\delta(y) \in \mathrm{tp}(b/M \cup \{a\})$, we have

$$\forall y(\chi(y, x) \rightarrow \delta(y)) \in \mathrm{tp}(a/M) = p(x).$$

Since p is countably satisfiable in M, there is some $a' \in M$ such that

$$M \models \exists y\, \chi(y, a') \quad \text{and for all } \delta(y) \in \Delta, \quad M \models \forall y(\chi(y, a') \rightarrow \delta(y)).$$

Picking $b' \in M^y$ such that $M \models \chi(b', a')$, we have $M \models \delta(b')$ for all $\delta(y) \in \Delta$. So Δ is realized in M. □

For us, the key property of the extension $M \prec N$ constructed in Lemma 7.15 is that N does not add any new elements to countable definable subsets of M. Indeed, if $\varphi(M)$ is countable, then the countable partial type

$$\varphi(x) \cup \{x \neq m \mid m \in \varphi(M)\}$$

is not realized in M and hence is also not realized in N. So $\varphi(N) = \varphi(M)$.

Proof of Theorem 7.6. Let κ be an uncountable cardinal. Suppose T is totally transcendental and has a Vaughtian pair. By Theorem 7.5, T has an $(\aleph_1, \aleph_0)$-model $(M, \varphi(x))$. That is, $M \models T$, $\varphi(x)$ is an L_M-formula, $|M| = \aleph_1$, and $|\varphi(M)| = \aleph_0$. We will build an elementary chain $(M_\alpha)_{\alpha \leq \kappa}$ such that for all $\alpha \leq \kappa$, $|M_\alpha| \leq \kappa$ and $\varphi(M_\alpha) = \varphi(M)$.

Let $M_0 = M$. Then $|M| = \aleph_1 \leq \kappa$ and $\varphi(M_0) = \varphi(M)$.

For the successor step, by Lemma 7.15, there is a proper elementary extension $M_\alpha \prec M_{\alpha+1}$ such that every type over M_α realized in $M_{\alpha+1}$ is countably satisfiable in M_α. By Löwenheim–Skolem, we may assume $|M_{\alpha+1}| = |M_\alpha| \leq \kappa$. As noted above, since $|\varphi(M_\alpha)| = \aleph_0$, $\varphi(M_{\alpha+1}) = \varphi(M_\alpha) = \varphi(M)$.

When $\gamma \leq \kappa$ is a limit ordinal, let $M_\gamma = \bigcup_{\alpha<\gamma} M_\alpha$. Then $|M_\gamma| \leq \kappa \cdot |\gamma| \leq \kappa$ and $\varphi(M_\gamma) = \bigcup_{\alpha<\gamma} \varphi(M_\alpha) = \varphi(M)$. In the case $\gamma = \kappa$, we have $|M_\kappa| = \kappa$ (since each extension in the κ-length chain is proper) and $|\varphi(M_\kappa)| = |\varphi(M)| = \aleph_0$. So $(M_\kappa, \varphi(x))$ is a $(\kappa, \aleph_0)$-model. □

7.2 A strongly minimal set in the prime model

The last component to our proof of Morley's theorem is to find a strongly minimal set with parameters in the prime model. First we note that minimal sets are ubiquitous in totally transcendental theories.

Lemma 7.16. *Suppose T is totally transcendental, and let $M \models T$. Then every infinite definable subset of M contains a minimal definable subset.*

Proof. Let $\varphi(x)$ be a non-algebraic L_M-formula, and suppose for contradiction that $\varphi(M)$ contains no minimal definable set. We build a binary tree of non-algebraic formulas $(\varphi_g(x))_{g\in 2^{<\omega}}$, contradicting T totally transcendental. Let $\varphi_\varnothing$ be $\varphi(x)$. Given $\varphi_g(x)$, since $\varphi_g(x)$ implies $\varphi(x)$, it is not minimal, so there is some formula $\psi(x)$ such that $\varphi_g \wedge \psi$ and $\varphi_g \wedge \neg\psi$ are both non-algebraic. Let $\varphi_{g0} = \varphi_g \wedge \psi$ and let $\varphi_{g1} = \varphi_g \wedge \neg\psi$. $\square$

If we apply Lemma 7.16 to the monster model $\mathcal{U}$, then we obtain strongly minimal sets. But a general model, even of a totally transcendental theory, may fail to contain any strongly minimal sets.

Example 7.17. Consider the theory from Example 6.12. This theory is totally transcendental, and its prime model M contains one equivalence class of size n for all $n \geq 1$, with no infinite classes. Every definable subset of M is finite or cofinite, so every infinite definable subset of M is minimal, but no such set is strongly minimal. Indeed, if we extend an infinite definable subset of M to $\mathcal{U}$, it will contain infinitely many equivalence classes, which are infinite and coinfinite definable subsets. The only strongly minimal sets in $\mathcal{U}$ are infinite equivalence classes (plus or minus finitely many points), and no such set is definable with parameters in the prime model.

In Proposition 6.13, we gave a criterion for minimal sets to be strongly minimal. We will now show that this criterion is always satisfied in theories with no Vaughtian pairs.

Definition 7.18. We say that T **eliminates** $\exists^\infty$ if for every partitioned formula $\varphi(x;y)$, there is a uniform bound $n_\varphi \in \omega$ such that for all $b \in \mathcal{U}^y$, if $\varphi(\mathcal{U};b)$ is finite, then $|\varphi(\mathcal{U};b)| \leq n_\varphi$.

Let me explain the terminology. Imagine we were to extend first-order logic with a new quantifier $\exists^\infty$ and the following semantics: $M \models \exists^\infty x\, \varphi(x;b)$ if and only if there are infinitely many $a \in M^x$ such that $M \models \varphi(a;b)$. If T eliminates $\exists^\infty$, then every formula in this extended logic is T-equivalent to an ordinary first-order formula, since $\exists^\infty x\, \varphi(x;y)$ is equivalent in models of T to $\exists^{>n_\varphi} x\, \varphi(x;y)$.

Theorem 7.19. *If T has no Vaughtian pairs, then T eliminates $\exists^\infty$.*

Proof. We prove the contrapositive. Suppose T fails to eliminate $\exists^\infty$. Then there is a formula $\varphi(x;y)$ such that for all $n \in \omega$, there exists $b \in \mathcal{U}^y$ such that $n < |\varphi(\mathcal{U};b)| < \aleph_0$.

We aim to show by compactness that T_{VP}^φ is consistent, and hence T has a Vaughtian pair. In fact, we show that for all $n \in \omega$, the L_{VP}^φ-theory T_n is consistent, where T_n contains (in the notation of Definition 7.10) axiom schema (1), axioms (2) through (4), and the finite subset of schema (5) given by $\exists^{\geq k} x\, \varphi(x;c)$ for $k \leq n$. Any finite subset of T_{VP}^φ is contained in T_n for some n, so this suffices.

Fixing $n \in \omega$, pick $b \in \mathcal{U}^y$ such that $n < |\varphi(\mathcal{U}; b)| < \aleph_0$. Let N be any small model containing b, and let M be any proper elementary extension of N. Interpreting the constants c as b and the predicate P as N. Since $N \prec M$, (1) and (2) are satisfied. Since $b \in N$, (3) is satisfied. Since $\varphi(x; b)$ is algebraic, $\varphi(N; b) = \varphi(M; b)$. And since $n < |\varphi(M; b)|$, $M \models \exists^{\geq k} x\, \varphi(x; c)$ for $k \leq n$. □

Theorem 7.20. *Suppose T is totally transcendental and has no Vaughtian pairs. Then there is a strongly minimal formula in a single variable with parameters in the prime model of T.*

Proof. First, recall that since T is totally transcendental, it has a prime model M (by Corollary 5.12, or, by a more elementary argument, by Fact 5.8 and Theorem 5.5). By Lemma 7.16, M contains a minimal definable set in a single variable x, say defined by the L_M-formula $\varphi(x)$. By Theorem 7.19, T eliminates $\exists^\infty$. For any formula $\psi(x; y)$, let $k = \max(n_{\varphi\wedge\psi}, n_{\varphi\wedge\neg\psi})$, so that for every $b \in M^y$, if $(\varphi(x) \wedge \psi(x; b))(M)$ is finite, it has size at most k, and if $(\varphi(x) \wedge \neg\psi(x; b))(M)$ is finite, it has size at most k. Since $\varphi(M)$ is minimal, one of these sets is finite for every $b \in M^y$, so by Proposition 6.13, $\varphi(x)$ is strongly minimal. □

7.3 Classifying models by dimensions

Suppose $M \models T$ and $\varphi(x; b)$ is a strongly minimal formula with parameters b from M. We define $\dim_{\varphi(x;b)}(M)$ to be the dimension of the strongly minimal set $D = \varphi(M; b)$, with respect to the closure operator $\mathrm{cl}(A) = \mathrm{acl}(Ab) \cap D$. As usual, we write Ab for the set A together with the elements of the tuple b.

Theorem 7.21. *The following are equivalent:*

(1) T is κ-categorical for some uncountable cardinal κ.

(2) T is κ-categorical for all uncountable cardinals κ.

(3) T is totally transcendental and has no Vaughtian pairs.

Proof. (2)⇒(1): Trivial.

(1)⇒(3): This is Corollary 4.21 (with Theorem 5.3) and Corollary 7.7.

(3)⇒(2): Suppose T is totally transcendental and has no Vaughtian pairs. By Theorem 7.20, there is a strongly minimal formula $\varphi(x; c)$ with parameters c in the prime model $P \models T$. Since P is atomic, $p = \mathrm{tp}(c/\varnothing)$ is isolated, and hence realized in every model.

Let $M \models T$ be any model, pick b realizing p in M, and let $D = \varphi(M; b)$, which is a strongly minimal set. By Corollary 5.12, T has a prime model Q over Db, which embeds elementarily into M over Db, so $D \subseteq Q \prec M$. Since T has no Vaughtian pairs, $M = Q$, so M is prime over Db.

Claim 1: If $|M|$ is uncountable, then $\dim_{\varphi(x;b)}(M) = |M|$.

Let B be a basis for D, so $\dim_{\varphi(x;b)}(M) = |B|$. Since M is prime over Db and $Db \subseteq \mathrm{acl}(Bb)$:

$$|M| = |Db| \leq |\mathrm{acl}(Bb)|.$$

Since b is finite, $|\mathrm{acl}(Bb)| = \max(\aleph_0, |B|)$. Since $|M|$ is uncountable, this implies $|M| \leq |B|$. But $B \subseteq M$, so $\dim_{\varphi(x;b)}(M) = |B| = |M|$.

Claim 2: Suppose $M' \models T$ is another model and b' is a realization of p in M'. If $\dim_{\varphi(x;b)}(M) = \dim_{\varphi(x;b')}(M')$, then $M \cong M'$.

First note that since b and b' realize the same type over $\varnothing$, there is a partial elementary bijection f mapping b to b'. Let B be a basis for D, and let B' be a basis for $D' = \varphi(M'; b')$. Since $|B| = |B'|$, we can pick a bijection between B and B'. Any finite tuple from B of length n realizes the type q_n of an independent n-tuple from D over b, and any finite tuple of length n from B' realizes the type f_*q_n of an independent n-tuple from D' over $b' = f(b)$, so f extends to a partial elementary bijection $g\colon Bb \to B'b'$. Since $D \subseteq \mathrm{acl}(Bb)$ and $D' \subseteq \mathrm{acl}(B'b')$, by Proposition 6.20, g extends to a partial elementary bijection $h\colon Db \to D'b'$.

Now since M is prime over Db, h extends to an elementary embedding $i\colon M \to N$. Then we have $D' \subseteq i(M) \prec N$. Since T has no Vaughtian pairs, $i(M) = N$, so i is an isomorphism $M \cong N$.

The result follows from the two claims. If κ is uncountable and $M, M' \models T$ with $|M| = |M'| = \kappa$, then picking b realizing p in M and b' realizing p in M', we have $\dim_{\varphi(x;b)}(M) = \kappa = \dim_{\varphi(x;b')}(M')$, so $M \cong M'$. Thus T is κ-categorical. $\square$

Definition 7.22. We say T is **uncountably categorical** if it satisfies the equivalent conditions in Theorem 7.21.

The notion of dimension for models of uncountably categorical theories introduced in the proof of Theorem 7.21 depends on a number of choices. First, we must pick a formula $\varphi(x; y)$ and an isolated type $p \in S_y(\varnothing)$ so that for any realization b of p, $\varphi(x; b)$ is a strongly minimal formula. This is a choice that can be fixed in advance and used uniformly for all models of T. More troublingly, to assign a dimension to $M \models T$, we must pick a realization b of p in M.

For uncountable models, we showed that $\dim_{\varphi(x;b)}(M) = |M|$, so the dimension is independent of all the choices above: every strongly minimal set in M has dimension $|M|$. But for countable models, it is less clear that $\dim_{\varphi(x;b)}(M)$ is an invariant of M, rather than of the pair (M, b). To classify the countable models of T, we would like to show $\dim_{\varphi(x;b)}(M)$ does not depend on the choice of b, from which it follows that $\dim_{\varphi(x;b)}(M) \neq \dim_{\varphi(x;b')}(M')$ implies $M \not\cong M'$.

For the rest of this section, we assume T is uncountably categorical, and we fix a formula $\varphi(x; y)$, where x is a single variable, and an isolated type $p \in S_y(\varnothing)$ so that for any realization b of p, $\varphi(x; b)$ is strongly minimal. We follow the proof in Tent and Ziegler, Section 6.3.

Theorem 7.23. *Let b be a realization of p in a countable model $M \models T$. Then $\dim_{\varphi(x;b)}(M) = \aleph_0$ if and only if M is saturated.*

Proof. Assume M is saturated. Let $a = (a_1, \ldots, a_n)$ be a finite tuple from the strongly minimal set $D = \varphi(M; b)$. Let $q(x) \in S_x(ab)$ be the unique non-algebraic type over ab containing $\varphi(x; b)$. Since M is $\aleph_0$-saturated, q is realized in M, say by c, and $c \in D \setminus \mathrm{cl}(a)$, so a is not a basis for D. Since D has no finite basis, $\dim_{\varphi(x;b)}(M) = \aleph_0$.

Conversely, suppose $\dim_{\varphi(x;b)}(M) = \aleph_0$. Since T is $\aleph_0$-stable, T has a countable saturated model M' by Theorem 2.18. Letting b' be a realization of p in M', the argument above shows that $\dim_{\varphi(x;b')}(M') = \aleph_0$. Then $M \cong M'$, so M is saturated. $\square$

It follows from the theorem that if a model M contains a strongly minimal set of dimension $\aleph_0$, then M is countable and saturated, so every strongly minimal set in M has dimension $\aleph_0$. Thus in this case too, the dimension is independent of all choices. So we can restrict attention to models in which every strongly minimal set has finite dimension.

Definition 7.24. A countable model $M \models T$ is **finite dimensional** if it is not saturated. Equivalently, by Theorem 7.23, if $\dim_{\varphi(x;b)}(M)$ is finite for all b realizing p in M.

Note that if T is $\aleph_0$-categorical, then every countable model of T is saturated, so T has no finite dimensional models.

Lemma 7.25. *Assume T is not $\aleph_0$-categorical. Then any model which is prime over a finite set is finite-dimensional.*

Proof. Suppose $M \models T$ is prime over a finite set A. Then M is countable and by Corollary 5.12, M is atomic over A. Since T is not $\aleph_0$-categorical, by Ryll-Nardzewski there is some n such that $S_n(\varnothing)$ is infinite. Then $S_n(A)$ is infinite, and hence, by compactness, there is a non-isolated type $q \in S_n(A)$, which is not realized in M, so M is not $\aleph_0$-saturated, and hence is finite dimensional. $\square$

For the analysis of finite dimensional models, it will be useful to consider a notion of "relative dimension". Suppose b is a realization of p in $M \models T$. For any set C containing b, we define $\dim_{\varphi(x;b)}(M/C)$ to be the dimension of the strongly minimal set $D = \varphi(M; b)$ viewed as a set defined over C, i.e., with respect to the closure operator

$$\mathrm{cl}(A) = \mathrm{acl}(A \cup C) \cap D.$$

Note that if $C \subseteq C' \subseteq M$, then $\dim_{\varphi(x;b)}(M/C') \leq \dim_{\varphi(x;b)}(M/C)$. For example, $\dim_{\varphi(x;b)}(M/M) = 0$.

When the base set C is a model, the relative dimension has a particularly simple description, which is also independent of all choices.

Lemma 7.26. *Let b be a realization of p in $M \models T$, and let N be a finite-dimensional elementary extension of M. Then $\dim_{\varphi(x;b)}(N/M)$ is the maximal length ℓ of an elementary chain:*

$$M = M_0 \prec M_1 \prec \cdots \prec M_\ell = N.$$

Proof. Write d for $\dim_{\varphi(x;b)}(N/M)$, and write ℓ for the maximal length of an elementary chain as in the statement of the lemma.

To show $d \leq \ell$, let $a_1, \ldots, a_d$ be a basis for $\varphi(N;b)$ over M. Let $M_0 = M$, let $M_d = N$, and for all $0 < i < d$, let M_i be a prime model over $Ma_1 \ldots a_i$. Since $Ma_1 \ldots a_i \subseteq M_{i+1}$, we may assume $M_i \preceq M_{i+1}$ for all i. By Corollary 5.12, M_i is atomic over $Ma_1 \ldots a_i$. The clopen set $[\varphi(x;b)]$ in the type space $S_x(Ma_1 \ldots a_i)$ contains infinitely many algebraic types over $Ma_1 \ldots a_i$, each of which is isolated, together with a unique non-algebraic type, which (by compactness) is not isolated, and hence is not realized in M_i. Since $a_1, \ldots, a_d$ is independent over M, $\mathrm{tp}(a_{i+1}/Ma_1 \ldots a_i)$ is not algebraic, so $a_{i+1} \notin M_i$. It follows that we have a chain of proper elementary extensions

$$M = M_0 \prec M_1 \prec \cdots \prec M_d = N,$$

so $d \leq \ell$.

To show $\ell \leq d$, suppose we have an elementary chain

$$M = M_0 \prec M_1 \prec \cdots \prec M_\ell = N.$$

For each $0 < i \leq \ell$, since T has no Vaughtian pairs, $\varphi(M_{i-1};b) \subsetneq \varphi(M_i;b)$. Pick some $a_i \in \varphi(M_i;b) \setminus \varphi(M_{i-1};b)$. By induction, since $Ma_1 \ldots a_{i-1} \subseteq M_{i-1}$, $\mathrm{acl}(Ma_1 \ldots a_{i-1}) \subseteq M_{i-1}$, so $a_i \notin \mathrm{acl}(Ma_1 \ldots a_{i-1})$. It follows by Lemma 6.27 that $a_1, \ldots, a_\ell$ is independent over M. Since any independent set can be extended to a basis, $d \geq \ell$. $\square$

Lemma 7.27. *Let b be a realization of p in $M \models T$, and let $D = \varphi(M;b)$. If $a_1, \ldots, a_n \in \varphi(\mathcal{U};b)$ are independent over Db, then they are independent over M.*

Proof. Suppose not. Then without loss of generality $\mathrm{tp}(a_1/Ma_2 \ldots a_n)$ is algebraic, so there is an algebraic formula $\psi(x_1, a_2, \ldots, a_n, c)$ in $\mathrm{tp}(a_1/Ma_2 \ldots a_n)$, where $c \in M^z$. We use stable embeddedness to move the parameters c inside D. More precisely, consider the partitioned formula $\psi(x_1, \ldots, x_n; z)$. Since T is totally transcendental, hence stable, by Proposition 3.16, $\mathrm{tp}_z^{\psi^{\mathrm{opp}}}(c/D)$ is definable, say by $\chi(x_1, \ldots, x_n; d)$ with d a tuple from D. Then

$$M \models \forall x_1, \ldots, x_n \left(\left(\bigwedge_{i=1}^{n} \varphi(x_i;b) \right) \to (\chi(x_1, \ldots, x_n; d) \leftrightarrow \psi(x_1, \ldots, x_n; c)) \right),$$

so the same is true in $\mathcal{U}$. It follows that $\varphi(x_1;b) \wedge \chi(x_1, a_2, \ldots, a_n; d)$ is an algebraic formula satisfied by a_1 with parameters in Db, so $a_1, \ldots, a_n$ is not independent over Db. $\square$

Lemma 7.28. *Let $M \preceq N$ be finite-dimensional models of T and b a realization of p in M. Then $\dim_{\varphi(x;b)}(N) = \dim_{\varphi(x;b)}(N/M) + \dim_{\varphi(x;b)}(M)$.*

Proof. Let B_M be a basis for $D = \varphi(M;b)$ over b, so $\dim_{\varphi(x;b)}(M) = |B_M|$. By Theorem 6.28, we can extend B_M to a basis B_N for $\varphi(N;b)$ over b, so

$\dim_{\varphi(x;b)}(N) = |B_N|$. It remains to show that $B = B_N \setminus B_M$ is a basis for $\varphi(N; b)$ over M, since then

$$\dim_{\varphi(x;b)}(N/M) = |B| = |B_N| - |B_M| = \dim_{\varphi(x;b)}(N) - \dim_{\varphi(x;b)}(M).$$

Generating: Since $\varphi(N; b) \subseteq \mathrm{acl}(B_N b)$ and $B_N b \subseteq BM$, $\varphi(N; b) \subseteq \mathrm{acl}(BM)$.

Independence: Since B_N is independent, B is independent over $B_M b$, and hence over $Db \subseteq \mathrm{acl}(B_M b)$. By Lemma 7.27, B is independent over M. □

Let M be a finite dimensional model of T, and let b_1 and b_2 be realizations of p in M. We define

$$\mathrm{diff}_M(b_1, b_2) = \dim_{\varphi(x;b_1)}(M) - \dim_{\varphi(x;b_2)}(M).$$

Our goal is to show that $\mathrm{diff}_M(b_1, b_2) = 0$ always.

Lemma 7.29. *The value* $\mathrm{diff}_M(b_1, b_2)$ *depends only on* $\mathrm{tp}(b_1, b_2/\varnothing)$. *That is, if* b_1', b_2' *realizes* $\mathrm{tp}(b_1, b_2/\varnothing)$ *in a finite dimensional model* $M' \models T$, *then* $\mathrm{diff}_M(b_1, b_2) = \mathrm{diff}_{M'}(b_1', b_2')$.

Proof. Let N be a prime model over $b_1 b_2$. We may assume $N \preceq M$. Now by Lemma 7.28:

$$\begin{aligned}
\mathrm{diff}_M(b_1, b_2) &= \dim_{\varphi(x;b_1)}(M) - \dim_{\varphi(x;b_2)}(M) \\
&= (\dim_{\varphi(x;b_1)}(M/N) + \dim_{\varphi(x;b_1)}(N)) \\
&\quad - (\dim_{\varphi(x;b_2)}(M/N) + \dim_{\varphi(x;b_2)}(N)) \\
&= \mathrm{diff}_N(b_1, b_2)
\end{aligned}$$

since $\dim_{\varphi(x;b_1)}(M/N) = \dim_{\varphi(x;b_2)}(M/N)$ by Lemma 7.26.

Now the partial elementary map $f(b_i) = b_i'$ extends to an elementary embedding $N \to M'$ with image $N' \preceq M'$. Since $N \cong N'$ by an isomorphism mapping b_i to b_i', $\mathrm{diff}_N(b_1, b_2) = \mathrm{diff}_{N'}(b_1', b_2')$. The same argument as above shows that $\mathrm{diff}_{N'}(b_1', b_2') = \mathrm{diff}_{M'}(b_1', b_2')$. So $\mathrm{diff}_M(b_1, b_2) = \mathrm{diff}_{M'}(b_1', b_2')$. □

We are ready to prove that the dimension is independent of the choice of parameter.

Theorem 7.30. *Let* b_1 *and* b_2 *be realizations of* p *in* $M \models T$. *Then*

$$\dim_{\varphi(x;b_1)}(M) = \dim_{\varphi(x;b_2)}(M).$$

Proof. If M is uncountable, then we showed in the proof of Theorem 7.21 that $\dim_{\varphi(x;b_1)}(M) = \dim_{\varphi(x;b_2)}(M) = |M|$. If M is countable and saturated, then $\dim_{\varphi(x;b_1)}(M) = \dim_{\varphi(x;b_2)}(M) = \aleph_0$ by Theorem 7.23. If M is countable and not saturated, then M is finite dimensional. It remains to show in this case that $\mathrm{diff}_M(b_1, b_2) = 0$.

Since we have a finite dimensional model, T is not $\aleph_0$-categorical. For any realizations b_1' and b_2' of p in $\mathcal{U}$, by Lemma 7.25 b_1' and b_2' are contained in some finite dimensional model, and we are justified in writing $\mathrm{diff}(b'_1, b'_2)$, without mentioning the model, since this value depends only on $\mathrm{tp}(b_1, b_2/\varnothing)$ by Lemma 7.29.

Let $q(y,z) = \mathrm{tp}(b_1, b_2)$. In $\mathcal{U}$, extend b_1 and b_2 to a sequence $(b_n)_{n\geq 1}$ of realizations of p. For each $n \geq 2$, since $\mathrm{tp}(b_n) = \mathrm{tp}(b_1) = p(y)$, the type $q(b_n, z)$ is consistent, and we can pick b_{n+1} realizing it. By Lemma 7.29, $\mathrm{diff}(b_n, b_{n+1}) = \mathrm{diff}(b_1, b_2)$ for all n.

Let $\alpha = \mathrm{MR}(p)$, and let $d = \mathrm{MD}(p)$. By Corollary 5.25, for any b realizing p, p has at most d generic extensions to $S_y(b)$. That is, there are at most d types $r(y,z)$ such that $r(y,b)$ extends p and has Morley rank α.

Let $p' \in S_y((b_n)_{n\geq 1})$ be a generic extension of p, so $\mathrm{MR}(p') = \alpha$, and let c realize p' in $\mathcal{U}$. For each n, $p \subseteq \mathrm{tp}(c/b_n) \subseteq p'$, so $\mathrm{tp}(c/b_n)$ has Morley rank α. Thus $\mathrm{tp}(c, b_n)$ is one of d possible types. It follows that there are some $i < j$ such that $\mathrm{tp}(c, b_i) = \mathrm{tp}(c, b_j)$, and hence $\mathrm{diff}(c, b_i) = \mathrm{diff}(c, b_j)$ by Lemma 7.29. We compute:

$$\begin{aligned}\mathrm{diff}(b_i, b_j) &= \sum_{k=i}^{j-1} \mathrm{diff}(b_i, b_{i+1})\\ &= \sum_{k=i}^{j-1} \mathrm{diff}(b_1, b_2)\\ &= (j-i)\mathrm{diff}(b_1, b_2).\end{aligned}$$

But also:

$$\begin{aligned}\mathrm{diff}(b_i, b_j) &= \mathrm{diff}(b_i, c) + \mathrm{diff}(c, b_j)\\ &= -\mathrm{diff}(c, b_i) + \mathrm{diff}(c, b_j)\\ &= 0.\end{aligned}$$

Thus $0 = (j-1)\mathrm{diff}(b_1, b_2)$, so $\mathrm{diff}(b_1, b_2) = 0$, as desired. □

By Theorem 7.30, we are justified in writing $\dim_\varphi(M)$ for the common value of $\dim_{\varphi(x;b)}(M)$ for any b realizing p in M. Using this invariant, we can give a complete description of the models of T.

Corollary 7.31. *Let M and N be models of T. Then $M \cong N$ if and only if $\dim_\varphi(M) = \dim_\varphi(N)$. Let P be the prime model of T.*

(1) If $\dim_\varphi(P) = \aleph_0$, then T is $\aleph_0$-categorical.

(2) If $\dim_\varphi(P) = n$ is finite, then for all cardinals k with $n \leq k \leq \aleph_0$, there is a unique (up to isomorphism) model $M_k \models T$ with $\dim_\varphi(M_k) = k$. In particular, T has countably many countable models, which can be arranged in an elementary chain:

$$P = M_n \prec M_{n+1} \prec M_{n+2} \prec \cdots \prec M_{\aleph_0}.$$

Proof. The first statement follows from Theorem 7.21 and the well-definedness of $\dim_\varphi$. If $\dim_\varphi(P) = \aleph_0$, then by Lemma 7.25, T is $\aleph_0$-categorical. Suppose $\dim_\varphi(P) = n$, and let b be a realization of p in P. The countable saturated model of T has dimension $\aleph_0$. For finite $k \geq n$, let $d = k - n$ and let M_k be the prime model over d elements of $\varphi(\mathcal{U}; b)$ which are independent over P. Then by Lemma 7.28, $\dim_\varphi(M) = \dim_\varphi(M/P) + \dim_\varphi(P) = d + n = k$. These models can be arranged in an elementary chain just as in the proof of Lemma 7.28. $\square$