(Yet) a(nother) course in model theory

Rosario Mennuni

Contents

0 Structures and formulas **1**

0.1 First-order structures 1

0.2 Formulas and theories 3

0.3 Multi-sorted structures 13

0.4 Building models: basic techniques 14

1 Five things **19**

1.1 Normal forms 19

1.2 Automorphisms 19

1.3 Some consequences of Löwenheim–Skolem 20

1.4 Taking unions 21

1.5 Finite structures 21

2 Quantifier elimination **25**

2.1 The first back-and-forth proof 25

2.2 Quantifier-free types 27

2.3 The Ra(n)do(m) graph 28

2.4 Syntax: eliminating quantifiers by hand 29

2.5 Types: packaging formulas together 31

2.6 Semantics: eliminating quantifiers by back-and-forth 34

2.7 Consequences: eliminating quantifiers for a purpose 36

2.8 Cheating: eliminating quantifiers by definitional expansions . . . 37

3 Examples and applications **39**

3.1 Algebraically closed fields 39

3.2 Some combinatorial structures 42

4 Realising many types **45**

4.1 Types over parameters 45

4.2 Type spaces 46

4.3 Examples 47

4.4 Saturation 53

4.5 Properties of saturated models 56

4.6 Monster models 58

4.7 Working in a monster model 61

5 Realising few types **65**
5.1 Isolated types 65
5.2 Omitting types 66
5.3 Prime models 70
5.4 The number of countable models 74
5.5 Ehrenfeucht–Mostowski models 79

6 Having few types **89**
6.1 Counting types 89
6.2 The order property 91
6.3 Local ranks 93
6.4 Definable types 94

7 Having very few models **97**
7.1 Morley rank 97
7.2 Totally transcendental theories and prime models 100
7.3 Countable ω-stable theories 102
7.4 Morley's Theorem 104

8 A taste of definable groups **107**
8.1 Consequences of the descending chain condition 107
8.2 Interpretability 108
8.3 Some forking calculus in disguise 109
8.4 The connected component 111
8.5 Generic types 113
8.6 Totally transcendental fields 114
8.7 An alternate ending 115

Readme

What is this? This document contains notes for a graduate course in model theory, held in the spring of 2022 at the Università di Pisa.

What is this not?

Going to change much anytime soon. Unless someone points out mistakes in them, I do not plan on editing these notes. If you want to let me know of any needed corrections, please write me at the email address below.

A book. There are already several beautiful books on model theory. Some are listed in the bibliography.

The only thing you should read. It is a good idea to look in the bibliography for examples, exercises, interesting topics, etc. Look, these are lecture notes, not a monograph, so, please, also consult other sources. Also, some of those books are *really* well written.

Original work. Of course nothing in these notes is an original result of the author. The exposition also owes various debts to multiple sources, notably Chapter 6 borrows heavily from [Che19] and Chapter 8 from [Mar02].

Why couldn't you choose a normal title? It's a tribute to two quite influential texts. Also, I like making references too much.

Why do pages move left and right? Because you are looking at a digital copy, but this version is made to be printed. If you really don't like this, just recompile with `twoside` replaced by `oneside` (see below for the source files).

Notation From some point on, lowercase letters like a, x, etc. will denote *tuples* of parameters or variables. To stress that a tuple has length 1, I write $|x| = 1$. On the other hand, if you read something like $x_0, \dots, x_n$, it probably (but not necessarily) means that they are single variables, and not tuples.

Acknowledgements Many thanks to Roberto Di Virgilio, Mariaclara Ragosta, and Alessandro Sferlazza for pointing out several inaccuracies in a previous version.

Info You can contact me at `R.Mennuni@posteo.net`. This version has been compiled on the 11th December 2023. To get the source code click on the leftmost paper clip. The bibliography source file is in the rightmost one.

Rosario Mennuni

Chapter 0

Structures and formulas

0.1 First-order structures

0.1.1 A "definition by example"

Example 0.1.1. The *language of ordered abelian groups* is $L_{\mathrm{oag}} \coloneqq \{+, 0, -, <\}$, where $+$ is a *binary function symbol*, 0 is a *constant symbol*, $-$ is a *unary function symbol* and $<$ is a *binary relation symbol.*

The "natural" way to define an L_{oag}-*structure* $\mathcal{R} \coloneqq (\mathbb{R}, +, 0, -, <)$ on the real numbers is as follows. The *domain* of $\mathcal{R}$ is $\mathbb{R}$, and we *interpret* $+$ as the function sending x, y to their usual sum, 0 as the usual number zero we all know and love, $-$ as the function sending x to its additive inverse, and $<$ as the usual order relation, that is, as the set of those $(x, y) \in \mathbb{R}^2$ such that x is strictly smaller than y.

The previous paragraph has a lot of words in order to avoid writing things like

$$\text{“we interpret } < \text{ as } <\coloneqq \{(x, y) \in \mathbb{R}^2 \mid x < y\}\text{”} \tag{1}$$

Here there is a lot of abuse of notation going on: the first instance of $<$ is a symbol; the second a subset of $\mathbb{R}^2$; and the third one means what you expect. If we want to distinguish between, say the symbol $<$, its interpretation in $\mathcal{R}$, and maybe we also want to be able to write $<$ to refer to the usual order of the reals (as in the third instance), we could for example use $\sqsubset$ as a symbol and write

$$\text{“we interpret } \sqsubset \text{ as } \sqsubset^{\mathcal{R}}\coloneqq \{(x, y) \in \mathbb{R}^2 \mid x < y\}\text{”}$$

In practice, once these distinctions are understood, writing things like $<^{\mathcal{R}}$ every time becomes very boring very quickly, so abuses of notation as in (1) are commonplace.

To stress the point further: it is completely legitimate to define an L_{oag}-structure $\mathcal{S}$ with domain $\mathbb{R}$ by setting the interpretation $+^{\mathcal{S}}$ to be multiplication, $-^{\mathcal{S}}$ to be the function sending x to e^x/x^3, $0^{\mathcal{S}} \coloneqq 23579$, and $<^{\mathcal{S}}$ to be the open unit disc intersected with $\mathbb{Q}^2$. But of course, while $\mathcal{R}$ is actually an ordered abelian group,[1] $\mathcal{S}$ is not, so if for some reason we want to study $\mathcal{S}$ it would be better to use a language with different symbols, instead of L_{oag}.

[1] In which sense? See section 0.2.

0.1.2 An actual definition (or two)

Let's give a couple of precise definitions.

Definition 0.1.2. A *(single-sorted, first-order) language* is a quadruple $L = (L_{\mathrm{c}}, L_{\mathrm{f}}, L_{\mathrm{r}}, \mathrm{ar}_L)$, where

1. $L_{\mathrm{c}}, L_{\mathrm{f}}, L_{\mathrm{r}}$ are pairwise disjoint sets, respectively the sets of *constant symbols*, *function symbols*, and *relation symbols*[2] of L; and
2. ar_L is a function $L_{\mathrm{f}} \cup L_{\mathrm{r}} \to \mathbb{N} \setminus \{0\}$.

- If $s \in L_{\mathrm{f}}$, we call $\mathrm{ar}_L(s)$ the *arity* of s, and say that s is a $\mathrm{ar}_L(s)$*-ary* function symbol.
- If $s \in L_{\mathrm{r}}$, we call $\mathrm{ar}_L(s)$ the *arity* of s, and say that s is a $\mathrm{ar}_L(s)$*-ary* relation symbol.

For instance, in Example 0.1.1, we have $\mathrm{ar}_{L_{\mathrm{oag}}}(+) = 2$, and we call $+$ a 2-ary function symbol. Synonyms such as "binary" instead of "2-ary", are also used.

In practice, one abuses the notation and just lists the symbols of a language in a single set and specifies in some way which symbols are constants, which are functions, and which are relations, as in Example 0.1.1. Another shorthand is to write arities as superscripts, as in "$L_{\mathrm{oag}} \coloneqq \{+^{(2)}, 0, -^{(1)}, <^{(2)}\}$".

Definition 0.1.3. Let L be a language. An *L-structure* $\mathcal{M}$ is given by the following.

1. A set M, called the *domain* (or *universe*) of $\mathcal{M}$.
2. For each constant symbol $c \in L_{\mathrm{c}}$, an element $c^{\mathcal{M}} \in M$.
3. For each function symbol $f \in L_{\mathrm{f}}$, a function $f^{\mathcal{M}} \colon M^{\mathrm{ar}(f)} \to M$.
4. For each relation symbol $R \in L_{\mathrm{r}}$, a subset $R^{\mathcal{M}} \subseteq M^{\mathrm{ar}(R)}$.

If s is a symbol, we call $s^{\mathcal{M}}$ its *interpretation* in $\mathcal{M}$.

Remark 0.1.4. Some authors only allow structures with nonempty domain. Sometimes this is convenient, sometimes it is not, see e.g. [Poi00, page 22].

Example 0.1.5. The *language of graphs* is $L_{\mathrm{graph}} \coloneqq \{E^{(2)}\}$. Your favourite graph G can be made into an L_{graph}-structure $\mathcal{G} \coloneqq (G, E^{\mathcal{G}})$, where $E^{\mathcal{G}}$ is the set of $(x, y) \in G^2$ such that there is an edge between x and y.

Example 0.1.6. Again, formally, *any* set G with *any* subset of G^2 is a perfectly legit L_{graph}-structure.

It is also commonplace to use the same notation for a structure and its domain, as in "the L_{oag}-structure $\mathbb{R}$", with the understanding that the interpretation of each symbol is clear from context. For the time being we will keep the notation distinct (but not for very long).

Remark 0.1.7. Slightly different approaches exist. For instance one may replace constant symbols by 0-ary function symbols; some authors also allow 0-ary relation symbols, which are interpreted as "always true" or "always false".

[2] Relation(symbol)s are sometimes also called *predicates*.

0.1.3 Expansions and reducts

Definition 0.1.8. Let L, L' be languages. We say that L' is a *sublanguage* of L, and write $L' \subseteq L$, iff $L'_{\mathrm{c}} \subseteq L_{\mathrm{c}}$, $L'_{\mathrm{f}} \subseteq L_{\mathrm{f}}$, $L'_{\mathrm{r}} \subseteq L_{\mathrm{r}}$, and $\mathrm{ar}_{L'} = \mathrm{ar}_L \restriction L'_{\mathrm{f}} \cup L'_{\mathrm{r}}$.

So a sublanguage of L is just a language with fewer symbols.

Definition 0.1.9. Let $\mathcal{M}$ be an L-structure and $L' \subseteq L$. The *reduct* $\mathcal{M} \restriction L'$ of $\mathcal{M}$ to L' is the L'-structure $\mathcal{M}'$ with the same domain as $\mathcal{M}$, and where every symbol $s \in L'$ is interpreted as $s^{\mathcal{M}'} = s^{\mathcal{M}}$. We call $\mathcal{M}$ an *expansion* of $\mathcal{M}'$ to L.

In other words, the reduct of an L-structure to $L' \subseteq L$ is obtained by forgetting the interpretations of symbols in $L \setminus L'$.

By the way, $s \in L'$, $L \setminus L'$, etc, are more abuses of notation: formally we should say, for example, "the constant symbols in $L_{\mathrm{c}} \setminus L'_{\mathrm{c}}$, the function symbols in…". Hopefully, by now it should be clear what kinds of pedantries are happening in the background, so we will stop commenting on them.

0.2 Formulas and theories

We still have a bunch of definitions to give but, as it is probably clear from the previous section, spelling out everything formally tends to be more lengthy than enlightening. So I am going to be brief and compensate with some examples. To see the details spelled out more precisely, see the literature, or the notes of a course in logic. Some references are [CK90, HL19, Hod93, Kir19, MT03, Mar02, Poi00, TZ12].

0.2.1 Formulas

Fix a language L, and fix a countably infinite[3] set V of *variables*, e.g. $V = \{x_0, x_1, x_2, \dots\}$.

Definition 0.2.1. Let L be a language. The set of *terms* of L is the closure of $L_{\mathrm{c}} \cup V$ under the functions of L_{f}. We write $t(x_0, \dots, x_n)$ to denote a term in which the set of variables appearing is included in $\{x_0, \dots, x_n\}$.

Example 0.2.2.
1. In L_{oag}, examples of terms are x_0, 0. Another example is $+(x_0, 0)$, but we also denote it by $x_0 + 0$. Yet another example is $(x_0 + 0) + (-x_1)$.
2. In L_{graph} the only terms are the variables. The same is true in every *relational* language, that is, a language with only relation symbols.

Remark 0.2.3. A term $t(x_0, \dots, x_n)$ need not necessarily mention all the variables $x_0, \dots, x_n$. For example it is perfectly legit to write $t(x_0, \dots, x_7) := x_0 + x_4$. It is also perfectly legit to write $t(x_0, x_4) := x_0 + x_4$. Or simply $x_0 + x_4$, but this is yet another (useful!) abuse of notation, and we may need to specify whether we regard this as a term in 2 or 8 variables (cf. Remark 0.2.5).

[3] For most purposes, countably many variables suffice, so we will assume this, but in some applications one needs larger sets of variables. Everything generalises fairly easily.

Non-Example 0.2.4. $x_0 < 0$ is *not* a term: it contains a relation symbol. $x_0 + x_1 + x_2$ is, strictly speaking, not a term: we need parentheses somewhere (no one guarantees that $+$ will be interpreted as an associative operation).

Remark 0.2.5. If $\mathcal{M}$ is an L-structure, every term $t(x_0, \dots, x_{n-1})$ of L induces a function $M^n \to M$, obtained in the obvious way.

Example 0.2.6. In the structure $\mathcal{R}$ that we encountered in Section 0.1.1, the L_{oag}-term $t(x_0, x_1, x_2) := x_0 + x_2$ induces the function $\mathbb{R}^3 \to \mathbb{R}$ summing the first and last coordinate.

Definition 0.2.7. An *atomic formula* of L is either:

1. $t_0(x_0, \dots, x_{n-1}) = t_1(x_0, \dots, x_{n-1})$, where t_0, t_1 are terms, or
2. $R(t_0(x_0, \dots, x_{n-1}), \dots, t_{m-1}(x_0, \dots, x_{n-1}))$, where the t_i are L-terms and R is an m-ary relation symbol of L.

So, in a sense, every structure is automatically equipped with the binary relation $=$. While other symbols in L can be interpreted in any way (consistent with their arity), $=$ *must* be interpreted as the diagonal.

Example 0.2.8. In L_{oag}, examples of atomic formulas are

- $x_0 + x_1 = x_2$
- $x_0 + x_1 < 0$

Definition 0.2.9. The set of *first-order L-formulas* is the closure of the set of L-atomic formulas under:

1. Boolean connectives $\wedge, \vee, \neg$.
2. First-order quantifiers $\exists x$, $\forall x$, where x is a variable.

Usual conventions about dropping parentheses apply. We also use the abbreviations $\varphi \to \psi$ for $(\neg\varphi) \vee \psi$ and $\varphi \leftrightarrow \psi$ for $(\varphi \to \psi) \wedge (\psi \to \varphi)$.

"First-order" means that quantifiers (well, variables, to begin with) range over M; that is, variables stand for elements of the domain. So, for example we cannot quantify over subset of M, topologies on M, etc. Unless otherwise specified, every formula we consider will be first-order, so we just say "formula" instead of "first-order formula".

One also defines the set of *free variables* of a formula: those which, at least once, occur not in the scope of any quantifier. This is one of the things were I will avoid giving a precise definition, refer to the literature, and supply examples instead.

Example 0.2.10. Examples of L_{oag}-formulas:

1. Those in Example 0.2.8.
2. $(x_0 < 0) \wedge (x_0 > 0)$.
3. $\exists x_0 \, \big(((x_1 + x_0 > 0) \vee (x_1 + x_0 = 0)) \wedge (\forall x_2 \, (x_2 < x_1))\big)$
4. $\exists x_3 \, \big(((x_1 + x_0 > 0) \vee (x_1 + x_0 = 0)) \wedge (\forall x_2 \, (x_2 < x_1))\big)$

5. $(\exists x_0\ (x_0 = 0)) \wedge (x_0 > 0)$.

Note that x_0 is not free in Item 3, but it is free in Item 4. The fact that x_3 is never mentioned after the quantifier in Item 4 is not a problem. In Item 5 x_0 *is* free. Formally, a variable may be used both free and bound in the same formula, but of course this has a tendency to make the reader angry, and it is good practice to use a fresh variable whenever possible.[4]

Non-Example 0.2.11. These are *not* first-order L_{oag} formulas:

1. $\exists n \in \mathbb{N}\ x_0 = n$. Formulas are allowed to talk about elements of the domain of the structure in which they will be interpreted; they do not know about natural numbers.

2. The usual formula saying that $\mathbb{R}$ is complete: we are not allowed to quantify over subsets of the domain.[5]

Remark 0.2.12. If we write $\varphi(x_0, \ldots, x_n)$ we mean that φ is a formula with free variables *included in* $\{x_0, \ldots, x_n\}$. The same abuse of notation as in Remark 0.2.3 applies, so we may for example write $\varphi(x_0) \coloneqq (x_0 < 0) \wedge (x_0 > 1)$ but also $\varphi(x_0, x_1) \coloneqq (x_0 < 0) \wedge (x_0 > 1)$. This becomes relevant when using formulas to *define* sets, see Definition 0.2.14.

One more tedious thing to define is what it means to *substitute* a term for a free variable in a formula. An example is: let $\varphi(y) \coloneqq y < 0$, let $t(x_0, x_1) \coloneqq x_0 + x_1$; then $\varphi(t(x_0, x_1))$ is $x_0 + x_1 < 0$. One just needs to be careful for variables not to be *captured*, that is, a substitution should not bound variables to a quantifier, as in substituting x_0 for y inside $\varphi(y) \coloneqq \exists x_0\ (\neg x_0 = y)$. These problems disappear if one uses fresh variables whenever possible.

Another thing that works as you expect is what it means for a point to *satisfy* a formula in a structure. To say that $(a_0, \ldots, a_{n-1}) \in M^n$ satisfies $\varphi(x_0, \ldots, x_{n-1})$ in $\mathcal{M}$, we write $\mathcal{M} \vDash \varphi(a_0, \ldots, a_{n-1})$.

Example 0.2.13. In $\mathcal{R}$, let $(a_0, a_1) = (-5, 3)$ and $\varphi(x_0, x_1) \coloneqq x_0 < x_1$. Then $\mathcal{R} \vDash \varphi(a_0, a_1)$. If $\psi(x_0, x_1) \coloneqq \exists y\ ((x_1 < y) \wedge (y < x_0))$, then $\mathcal{R} \nvDash \psi(a_0, a_1)$. Also, $\mathcal{R} \vDash (\exists z\ (\varphi \wedge \neg\psi))(a_0, a_1)$ (yes, that $\exists z$ is entirely superfluous).

The formal definition is by induction on the complexity of the formula: $\mathcal{M} \vDash (\varphi \wedge \psi)(a_0, \ldots, a_n)$ iff $\mathcal{M} \vDash \varphi(a_0, \ldots, a_n)$ and $\mathcal{M} \vDash \psi(a_0, \ldots, a_n)$, while $\mathcal{M} \vDash \exists x\ \varphi(a_0, \ldots, a_n, x)$ iff there is $b \in \mathcal{M}$ such that $\mathcal{M} \vDash \varphi(a_0, \ldots, a_n, b)$, etc.

While we are here, let us say that a formula with no free variables is called a *sentence*. If φ is a sentence, then either $\mathcal{M} \vDash \varphi$ or $\mathcal{M} \vDash \neg\varphi$ (with no need to assign a point to free variables, since there are none). For example, $\mathcal{R} \vDash \forall x\ x + 0 = x$.

So, what sentences *do* in a structure is either holding or not holding. What formulas with free variables do is *defining sets*.

[4] Well, in fact there *are* instances where minimising the number of variables used in a formula is useful, and one tries to recycle them as much as possible, but we will not talk about it in this course.

[5] Allowing that results in *second-order* logic. Allowing also quantification over families of subsets yields *third-order* logic, etc.

0.2.2 Definable sets

Definition 0.2.14. The set *defined* by an L-formula $\varphi(x_0, \ldots, x_{n-1})$ in $\mathcal{M}$ is

$$\varphi(\mathcal{M}) \coloneqq \{(a_0, \ldots, a_{n-1}) \in M^n \mid \mathcal{M} \vDash \varphi(a_0, \ldots, a_{n-1})\}$$

A subset of M^n is *definable* (in $\mathcal{M}$) iff it is defined by some L-formula.

Example 0.2.15. In $\mathcal{R}$, the formula $\varphi(x_0, x_1, x_2) \coloneqq x_0 + x_1 = x_2$ defines the graph of addition.

Non-Example 0.2.16. Using techniques that we will develop later in the course, it is possible to prove that the set $\mathbb{Z}$ is *not* definable in $\mathcal{R}$.

Quite often it is necessary to look at formulas with *parameters* from some subset $A \subseteq M$. It means what you expect, but formally this is what one does.

Definition 0.2.17. Let $\mathcal{M}$ be an L-structure and $A \subseteq M$. Define a language $L(A) \supseteq L$ by adding to L a new constant symbol c_a for every $a \in A$. Expand $\mathcal{M}$ to an $L(A)$ structure $\mathcal{M}_A$ by interpreting each c_a with a. A subset of M^n is A-*definable*, or *definable over* A iff it is definable in $\mathcal{M}_A$.

Example 0.2.18. In $\mathcal{R}$, the set $\{(x_0, x_1) \in \mathbb{R}^2 \mid x_0 < x_1 + 5\}$ is definable over $\mathbb{Z}$ (or even just over $\{5\}$).

See the literature for more lists of examples, e.g. [Mar02, Section 1.3] has some nice, more convoluted (and more interesting!) ones.

Sometimes we say that a set is $\emptyset$-definable to emphasise that it is definable without using parameters. Depending on the context, people use the work *definable* to mean "definable over $\emptyset$" or "definable over M". For now, we stick to the first meaning.

Remark 0.2.19. The set $\mathbb{Z}$ is *not* $\mathbb{Z}$-definable in $\mathcal{R}$. We cannot prove this yet, but for now observe that the natural attempt to a $\mathbb{Z}$-formula defining it would use an *infinite* disjunction $\bigvee_{i \in \mathbb{Z}} x = i$. This is *not* a first-order formula.

Remark 0.2.20. Definable sets in a given dimension form a boolean algebra, with the operations induced by the connectives $\wedge, \vee, \neg$, which of course correspond to intersection, disjunction, and complement of definable sets.

Some spoilers

Boolean algebras of definable sets, and their Stone duals,[6] are central objects of study in model theory, to the point that some people would go as far as saying that contemporary model theory *is* the study of definable sets in "tractable" structures. Of course one needs to make "tractable" precise —actually, the word "tame" is usually more popular in this context— and in fact there are several different notions of "tameness", that apply to different structures and have different consequences. For example, the fact that, in every dense linear order with no endpoints $\mathcal{M}$, for every n, there are only finitely many $\emptyset$-definable subsets of M^n, is intimately connected to the fact that all countable dense linear

[6] If you do not know what this means, then I recommend reading about Stone duality *after* you get familiar with types, later on in the course. But of course, if you are going to read about it straight away I will not try to stop you.

orders without endpoints are isomorphic to $(\mathbb{Q}, <)$ (we will see this later on in the course). Hence, looking at definable sets can allow us to say something about certain classes of structures.

Things also work in the other direction, and definable sets are at the centre of several applications of model theory: one is interested in a certain family of sets, and uses their definability in some tame structure to say something about them. For example, we will prove later that the sets which are $\mathbb{R}$-definable in $\mathcal{R}$ are precisely the *semilinear* sets: that is, boolean combinations of sets defined by inequalities between affine functions, e.g. polyhedra. One specific flavour of tameness enjoyed by the structure $\mathcal{R}$ implies, among other things, that:

- semilinear sets have certain decompositions in finitely many semilinear pieces of a nice form; this implies, for example, that semilinear sets have finitely many connected components;
- for every n, functions definable in $\mathcal{R}$ are piecewise $\mathcal{C}^n$;
- definable families of semilinear sets are learnable by certain kinds of algorithms.[7]

Now, one of the reasons $\mathcal{R}$ is our recurring example is that it is quite understandable, so, for semilinear sets, the things above can probably be proven directly and quite painlessly. The point is that the same notion of tameness applies to more complicated structures, and then the statements above become quite nontrivial. For instance, the same notion of tameness[8], hence the consequences above, hold for sets definable in the expansion of $\mathcal{R}$ by the field structure, the exponential functions, and the restrictions to bounded boxes of all analytic functions (simultaneously!).

At this point, the above will probably make little sense. That's normal. The point is that definable sets are important, and the subsection where they are introduced should definitely be at least one page long, but we still need to set up a bunch of things, so for now I could only give a couple of definitions and examples. I guess that this is enough rambling to make this longer than one page, so maybe it's time to stop.

If you do not like formulas, ~~then why are you even reading th~~ there is an alternative presentation of definable sets: see [Mar02, Proposition 1.3.4].

0.2.3 Theories

As remarked above, if we are given an L_{oag}-structure $\mathcal{M}$, there is no guarantee that, for example, the symbol $+$ will be interpreted as an associative operation. But associativity of $+$ can be expressed by a sentence, namely[9]

$$\forall x_0, x_1, x_2 \; \big((x_0 + x_1) + x_2 = x_0 + (x_1 + x_2)\big)$$

If we want to study ordered abelian groups, we may then write a set of L_{oag}-sentences such that, if a structure $\mathcal{M}$ satisfies them, then it actually is an ordered abelian group. A set of sentences which is satisfied in at least one structure is called a *theory*.

[7] The precise statement is that they have finite Vapnik–Chervonenkis dimension.

[8] Which, by the way, is called *o-minimality*. A standard reference is [vdD98a].

[9] Of course $\forall x_0, x_1, x_2$ is an abbreviation for $\forall x_0(\forall x_1(\forall x_2 \ldots))$

Definition 0.2.21. Let Φ be a set of L-sentences.

1. An L-structure $\mathcal{M}$ *satisfies* Φ, or is a *model* of Φ, written $\mathcal{M} \vDash \Phi$, iff for all $\varphi \in \Phi$ we have $\mathcal{M} \vDash \varphi$.

2. We say that Φ is *consistent* iff it has a model, i.e. iff there is $\mathcal{M}$ with $\mathcal{M} \vDash \Phi$.

3. An *L-theory* is a consistent set of L-sentences.

4. If T is a theory, its elements are called its *axioms*.

Some authors use "theory" to mean just "set of sentences", without requiring consistency.

Example 0.2.22. The *theory of ordered abelian groups* T_{oag} is the L_{oag}-theory containing the axioms:

1. $\forall x_0, x_1, x_2 \ \big((x_0 + x_1) + x_2 = x_0 + (x_1 + x_2)\big)$

2. $\forall x_0 \ \big(x_0 + 0 = x_0\big)$

3. $\forall x_0, x_1 \ \big(x_0 + x_1 = x_1 + x_0\big)$

4. $\forall x_0 \ \big(x_0 + (-x_0) = 0\big)$

5. $\forall x_0, x_1, x_2 \ \big((x_0 < x_1) \to (x_0 + x_2 < x_1 + x_2)\big)$

For some other standard examples of theories, see e.g. [Mar02, Section 1.2] or [Hod93, Section 2.2].

Definition 0.2.23. If T is an L-theory and φ is an L-sentence, we write $T \vdash \varphi$,[10] and say that φ is a *consequence* of T, iff for all $\mathcal{M} \vDash T$ we have $\mathcal{M} \vDash \varphi$. The *deductive closure* of T is the set of its consequences.

One says that Φ is an *axiomatisation* of T to mean T and Φ have the same consequences, i.e., the same deductive closure. For many purposes, it is convenient to identify a theory T with its deductive closure;[11] from now on, we will adopt this convention.

Definition 0.2.24. An L-theory T is *complete* iff, for every L-sentence φ, either $T \vdash \varphi$ or $T \vdash \neg\varphi$.

Non-Example 0.2.25. T_{oag} is not complete: if φ is the sentence

$$\varphi \coloneqq \forall x \ \exists y \ y + y = x$$

then $\mathcal{R} \vDash \varphi$, but $\mathbb{Z} \vDash \neg\varphi$ (where $\mathbb{Z}$ is made into an L_{oag}-structure in the natural way).

[10] We may also write $T \vDash \varphi$. There is a subtle difference which you should know if you have taken a course in logic. Otherwise, you may take them as synonyms.

[11] Sometimes, the distinction is important, e.g. sometimes it is important to know whether certain theories can be *finitely axiomatised*.

Example 0.2.26. Let T be the *theory of infinite sets*, that is, the theory in the empty language L (so, the only atomic formulas are equalities between variables) axiomatised by $\{\varphi_n \mid n \in \mathbb{N} \setminus \{0\}\}$, where[12]

$$\varphi_n \coloneqq \exists x_0, \dots, x_{n-1} \bigwedge_{i \neq j < n} x_i \neq x_j$$

With tools to be developed soon, it is possible to prove that T is complete.

Example 0.2.27. If $\mathcal{M}$ is any L-structure, the theory $\mathrm{Th}(\mathcal{M})$, defined to be the set of L-sentences that hold in $\mathcal{M}$, is (trivially) complete.

The example above is trivial, but typical:

Exercise 0.2.28. Prove the following statements.[13]

1. An L-theory is complete if and only if its deductive closure is maximal under inclusion (among L-theories, that is, *consistent* sets of L-formulas).

2. For every complete L-theory T, there is an L-structure $\mathcal{M}$ such that $T = \mathrm{Th}(\mathcal{M})$.

Of course, having the same complete theory, that is, satisfying the same sentences, deserves a name. And so does having the same models.

Definition 0.2.29. Two L-structures $\mathcal{M}$ and $\mathcal{N}$ are *elementarily equivalent*, written $\mathcal{M} \equiv \mathcal{N}$, iff $\mathrm{Th}(\mathcal{M}) = \mathrm{Th}(\mathcal{N})$.

Definition 0.2.30. Let T be an L-theory. Two L-sentences φ, ψ are *equivalent modulo* T iff, for every $\mathcal{M} \vDash T$, we have $\mathcal{M} \vDash \varphi \iff \mathcal{M} \vDash \psi$. Two formulas $\varphi(x_0, \dots, x_n)$ and $\psi(x_0, \dots, x_n)$ are *equivalent modulo* T iff, for all $\mathcal{M} \vDash T$, we have $\varphi(\mathcal{M}) = \psi(\mathcal{M})$.

Remark 0.2.31. Two formulas $\varphi(x_0, \dots, x_n)$ and $\psi(x_0, \dots, x_n)$ are equivalent modulo T if and only if $T \vdash \forall x_0, \dots, x_n\ (\varphi(x_0, \dots, x_n) \leftrightarrow \psi(x_0, \dots, x_n))$.

If we say that φ, ψ are *equivalent*, or *logically equivalent*, without specifying T (and without having a fixed T which is clear from context), we mean that they are equivalent modulo $T = \emptyset$, or modulo $T = \{\exists x\ x = x\}$ if we want to exclude empty structures, cf. Remark 0.1.4. Usually the second convention is used; note that, for example, $\forall x\ x = x$ and $\exists x\ x = x$ are equivalent under the second convention but not under the first one.

It is harmless, and also quite convenient, to introduce a logical symbol $\perp$ for "false". That is, $\perp$ is an atomic formula and is false in every structure.

Remark 0.2.32. Up to equivalence, we may write every formula using only $\wedge, \neg, \perp, \exists$, and recover $\forall$ and $\vee$ from them in the usual way.

This is useful when proving things by induction on (complexity of) formulas, since it allows to consider fewer cases (see e.g. the proof of Theorem 0.2.51).

[12] Below there are a bunch of abuses of notation going on, such as $x_i \neq x_j$ for $\neg(x_i = x_j)$; hopefully the meaning is clear.

[13] Hint: the second one follows from the first.

0.2.4 Substructures

Definition 0.2.33. We say that the L-structure $\mathcal{M}$ is a *substructure* of the L-structure $\mathcal{N}$ (and that $\mathcal{N}$ is an *extension*[14] of $\mathcal{M}$), and write $\mathcal{M} \subseteq \mathcal{N}$, iff:

1. $M \subseteq N$;
2. for every constant symbol $c \in L$, we have $c^{\mathcal{M}} = c^{\mathcal{N}}$;
3. for every n-ary function symbol $f \in L$, we have $f^{\mathcal{M}} = f^{\mathcal{N}} \upharpoonright M^n$; and
4. for every n-ary relation symbol $R \in L$, we have $R^{\mathcal{M}} = R^{\mathcal{N}} \cap M^n$; in other words, for every $a_0, \ldots, a_{n-1} \in M$, we have $\mathcal{M} \vDash R(a_0, \ldots, a_{n-1}) \iff \mathcal{N} \vDash R(a_0, \ldots, a_{n-1})$.

Example 0.2.34. Seen as L_{oag}-structures with the usual interpretations, we have $\mathbb{Z} \subseteq \mathbb{Q}$ and $\mathbb{Q} \subseteq \mathbb{R}$.

Example 0.2.35. If $\mathcal{G}$ is a graph, viewed as an L_{graph} structure in the natural way[15], then a substructure of $\mathcal{G}$ is the same as an *induced* subgraph of $\mathcal{G}$.

Non-Example 0.2.36. Let $\mathcal{G}$ be the complete graph on $\mathbb{N}$, and let $\mathcal{H}$ be a graph on $\mathbb{N}$ with no edge between 3 and 64. Then $\mathcal{H}$ is *not* a substructure of $\mathcal{G}$.

Non-Example 0.2.37. Let $\mathcal{P}$ be a poset in the language $\{\leq\}$, and suppose $a, b \in P$ are not comparable, that is, $\mathcal{P} \vDash (\neg(a \leq b)) \wedge (\neg(b \leq a))$. Let $\mathcal{P}'$ be some linear order with domain P extending the order of $\mathcal{P}$. Then $\mathcal{P}$ is *not* a substructure of $\mathcal{P}'$ (nor the other way around).

Example 0.2.38. If L is a relational language, and $\mathcal{M}$ an L-structure, then every $A \subseteq M$ can be made into an L-substructure of $\mathcal{M}$ in a unique way.

Example 0.2.39. More generally, if $A \subseteq M$, and B is the closure of A under the functions and constants of L, then B can be (uniquely) made into a substructure of M. Of course, this substructure is called the *substructure of $\mathcal{M}$ generated by A*.

Definition 0.2.40. An injective map $M \to N$ is an *embedding* of $\mathcal{M}$ into $\mathcal{N}$ iff its image is a substructure of $\mathcal{N}$.

Almost by definition, embeddings are injective maps preserving atomic formulas, that is, an injective map $\iota \colon M \to N$ is an embedding $\mathcal{M} \to \mathcal{N}$ if and only if, for every atomic formula $\varphi(x_0, \ldots, x_n)$ and $a_0, \ldots, a_n \in M$, we have

$$\mathcal{M} \vDash \varphi(a_0, \ldots, a_n) \iff \mathcal{N} \vDash \varphi(\iota(a_0), \ldots, \iota(a_n)) \tag{2}$$

We will not use them, but it is worth mentioning that *morphisms* of L-structures are defined similarly, by dropping the requirement of injectivity and weakening (2) by replacing $\iff$ with $\implies$. For instance, in Non-Example 0.2.37, the identity map $P \to P$ is a morphism (but not an embedding) $\mathcal{P} \to \mathcal{P}'$. And I do not want to risk offending any of you by telling you what an *isomorphism* is, or what *automorphisms* are.

Anyway, we were saying, embeddings preserve atomic formulas. A bit more is true.

[14] Not to be confused with *expansion*, cf. Definition 0.1.9.

[15] I'll stop writing this kind of stuff. If an interpretation is not specified, it's supposed to be the natural one.

Definition 0.2.41. A formula $\varphi(x_0, \ldots, x_n)$ is *quantifier-free* if no quantifier appears in φ.

Exercise 0.2.42. If $\mathcal{M} \subseteq \mathcal{N}$ and $\varphi(x_0, \ldots, x_n)$ is quantifier-free, then for every $a_0, \ldots, a_n \in M$ we have $\mathcal{M} \vDash \varphi(a_0, \ldots, a_n) \iff \mathcal{N} \vDash \varphi(a_0, \ldots, a_n)$.

The assumption that φ is quantifier-free is important:

Example 0.2.43. Let $L = \{<\}$, $\mathcal{M} = (\mathbb{Z}, <)$ and $\mathcal{N} = (\mathbb{Q}, <)$. Then $\mathcal{M} \subseteq \mathcal{N}$. Let $\varphi(x_0, x_1)$ be the formula $\exists y\ \big((x_0 < y) \land (y < x_1)\big)$. Then $\varphi(0, 1)$ holds in $\mathcal{N}$, but not in $\mathcal{M}$.

In particular, $\varphi(\mathcal{M}) \neq \varphi(\mathcal{N}) \cap M^2$. Even if $0, 1 \in M$, whether $\varphi(0, 1)$ holds or not depends on whether we check in $\mathcal{M}$ or in $\mathcal{N}$.

Substructures where this never happens are called *elementary*.

Definition 0.2.44. We say that $\mathcal{M}$ is an *elementary substructure* of $\mathcal{N}$ (and $\mathcal{N}$ an *elementary extension* of $\mathcal{M}$), written $\mathcal{M} \preceq \mathcal{N}$, iff $\mathcal{M} \subseteq \mathcal{N}$ and, for every formula $\varphi(x_0, \ldots, x_n)$ and $a_0, \ldots, a_n \in M$, we have

$$\mathcal{M} \vDash \varphi(a_0, \ldots, a_n) \iff \mathcal{N} \vDash \varphi(a_0, \ldots, a_n)$$

An embedding $\mathcal{M} \to \mathcal{N}$ is an *elementary embedding* iff its image is an elementary substructure of $\mathcal{N}$.

These easy observations are essentially exercises in spelling out definitions, but are quite important:

Remark 0.2.45. Let $\mathcal{M} \subseteq \mathcal{N}$.

- $\mathcal{M} \preceq \mathcal{N}$ if and only if, for every formula $\varphi(x_0, \ldots, x_{n-1})$, we have $\varphi(\mathcal{M}) = \varphi(\mathcal{N}) \cap M^n$.[16]
- $\mathcal{M} \preceq \mathcal{N}$ if and only if they have the same $L(M)$-theory.
- In particular, if $\mathcal{M} \preceq \mathcal{N}$, then $M \equiv N$.

At the risk of offending someone, let me point out that isomorphisms are elementary embeddings. Nevertheless, elementarity is really a condition on the embedding, and not just on the isomorphism type:

Example 0.2.46. Let $\mathcal{N} \coloneqq (\mathbb{Z}, <)$ and $\mathcal{M} = (2\mathbb{Z}, <)$. Then $\mathcal{M} \subseteq \mathcal{N}$, $\mathcal{M} \cong \mathcal{N}$, but $\mathcal{M} \not\preceq \mathcal{N}$, as can be checked by looking at the formula $\exists x\ 0 < x < 2$.[17]

How does one check that a substructure is elementary? See the Tarski–Vaught test below, Theorem 0.2.51.

0.2.5 Diagrams

Recall the natural expansions by constants defined in Definition 0.2.17.

Definition 0.2.47. Let $\mathcal{M}$ be an L-structure.

1. Its *elementary diagram* $\mathrm{ED}(\mathcal{M})$ is the complete $L(M)$-theory of $\mathcal{M}_M$.
2. Its *diagram*[18] $\operatorname{diag}(\mathcal{M})$ is the subset of $\mathrm{ED}(\mathcal{M})$ given by atomic formulas

[16] If you are particularly categorically-minded, you may like to think of definable sets as (a particular kind of) functors from the category of L-structures with elementary embeddings to the category of sets with injective maps.

[17] Which of course is an abbreviation for $\exists x\ ((0 < x) \land (x < 2))$, but I guess it's time to start being a bit less pedantic.

[18] Also known as *atomic diagram*.

and negations of atomic formulas.

Note that $\text{ED}(\mathcal{M})$ is, by definition, always a complete $L(M)$-theory. On the other hand, $\operatorname{diag}(\mathcal{M})$ need not be (Exercise 0.2.50).

Exercise 0.2.48. If φ is a quantifier-free $L(M)$-sentence and $\mathcal{M}_M \vDash \varphi$, then $\operatorname{diag}(\mathcal{M}) \vdash \varphi$.

The point of these definitions is that models of the (elementary) diagram of $\mathcal{M}$ correspond to (elementary) extensions of $\mathcal{M}$:

Proposition 0.2.49. Let $\mathcal{M}$ be an L-structure, and $\mathcal{N}$ be an $L(M)$-structure. Let $\iota\colon M \to N$ be the map $m \mapsto c_m^{\mathcal{N}}$. Then:

1. ι is an embedding if and only if $\mathcal{N} \vDash \operatorname{diag}(\mathcal{M})$.

2. ι is an elementary embedding if and only if $\mathcal{N} \vDash \text{ED}(\mathcal{M})$.

Proof. Exercise (easy). □

This is useful, because it allows us to build elementary extensions of $\mathcal{M}$ with certain properties by writing down suitable theories containing[19] $\text{ED}(\mathcal{M})$.

Exercise 0.2.50. Find an $\mathcal{M}$ such that $\operatorname{diag}(\mathcal{M})$ is not complete.

Theorem 0.2.51 (Tarski–Vaught test)**.** Let $\mathcal{N}$ be an L-structure, and suppose that $M \subseteq N$. The following are equivalent.

1. M is the domain of an elementary substructure $\mathcal{M} \preceq \mathcal{N}$.

2. For all $\varphi(x, y_0, \dots, y_n)$ and all $b_0, \dots, b_n \in M$, if there is $a \in N$ such that $\mathcal{N} \vDash \varphi(a, b_0, \dots, b_n)$, then there is $a' \in M$ such that $\mathcal{N} \vDash \varphi(a', b_0, \dots, b_n)$.

The statement of the Tarski–Vaught test (or criterion) looks very similar to the definition of $\preceq$. The difference is that, in order to check the condition in the criterion, we only need to look at which formulas are satisfied in $\mathcal{N}$: there is no "$\mathcal{M} \vDash$" in the statement; in fact, in the assumptions M is just a subset of N, and has not been given an L-structure (yet). This is subtle but important, as it allows arguments like the proof of Theorem 0.4.14 to go through.[20]

Proof. (1) $\Rightarrow$ (2) follows easily from the definition of $\preceq$.

Towards proving (2) $\Rightarrow$ (1), observe that, if $f(y_0, \dots, y_m)$ is an m-ary function symbol of L and $b_0, \dots, b_m \in M$, by using (2) with the formula $\varphi(x, y_0, \dots, y_m) \coloneqq x = f(y_0, \dots, y_m)$, we find that M is closed under the function symbols of L. A similar argument shows that M contains the interpretation of every constant, therefore M is the domain of a substructure $\mathcal{M}$ of $\mathcal{N}$.

To show elementarity, we now need to show that, whenever $\varphi \in L(M)$ is a sentence, then

$$\mathcal{M} \vDash \varphi \iff \mathcal{N} \vDash \varphi \tag{3}$$

We argue by induction on formulas. If (3) holds for φ and ψ, then it is immediate to observe that it also holds for $\neg\varphi$ and for $\varphi \wedge \psi$. Let us consider the case

[19] In languages larger than $L(M)$, since $\text{ED}(\mathcal{M})$ is already complete.

[20] By the way, that proof only uses things which we have already introduced, so you can read it right now if you wish. Just make sure to read Notation 0.4.13 first.

$\exists x\ \varphi(x)$. If $\mathcal{M} \vDash \exists x\ \varphi(x)$, then there is $a \in M$ such that $\mathcal{M} \vDash \varphi(a)$. But $\varphi(a)$ has lower complexity, so by induction $\mathcal{N} \vDash \varphi(a)$, and in particular $\mathcal{N} \vDash \exists x\ \varphi(x)$, proving $\implies$. For the converse, suppose $\mathcal{N} \vDash \exists x\ \varphi(x)$; then there is $a \in N$ such that $\mathcal{N} \vDash \varphi(a)$. By assumption, there is $a' \in M$ such that $\mathcal{N} \vDash \varphi(a')$, and again by inductive hypothesis $\mathcal{M} \vDash \varphi(a')$, hence $\mathcal{M} \vDash \exists x\ \varphi(x)$. □

0.3 Multi-sorted structures

As you may have expected from the "single-sorted" in Definition 0.1.2, there are things called "multi-sorted" (or "many-sorted") languages. This is one of those things where an example may be clearer than a definition:

Example 0.3.1. The *language of vector spaces* has two sorts, denoted by K and V, together with:

1. a constant symbol 0_K of arity K,
2. a function symbol $+_K$ of arity $K^2 \to K$,
3. a function symbol $-_K$ of arity $K \to K$,
4. a constant symbol 1 of arity K,
5. a function symbol $\cdot_K$ of arity $K^2 \to K$,
6. a constant symbol 0_V of arity V,
7. a function symbol $+_V$ of arity $V^2 \to V$,
8. a function symbol $-_V$ of arity $V \to V$, and
9. a function symbol $\cdot$ of arity $K \times V \to V$.

Instead of having a single set as a domain, a structure $\mathcal{M}$ for this language will have set $K(\mathcal{M})$ interpreting the sort K, and a set $V(\mathcal{M})$ interpreting the sort V. The constant symbol 0_K will be interpreted as an element of $K(\mathcal{M})$, the function symbol $\cdot$ as a function $K(\mathcal{M}) \times V(\mathcal{M}) \to V(\mathcal{M})$, etc.

An example of formula in this language is $\varphi(x, v) \coloneqq x \cdot v = 0$, where x is a variable *of sort* K and v is a variable *of sort* V. It defines a subset of $K(\mathcal{M}) \times V(\mathcal{M})$.

Below is a quick list of what changes from single-sorted languages to multi-sorted ones. See [End01, Section 4.3] for a formal definition.

1. Each sort has its own variables; in other words, each variable has a sort and ranges over that sort.
2. Each constant symbol has an arity, which is a sort,
3. Each function symbol has an arity, which is of the form $A \to B$, where A is a cartesian product of sorts and B is a sort; when building terms, we are only allowed to plug variables/constants/parameters in a function symbol if they come from the correct sorts.

4. Each relation symbol has an arity, which is a cartesian product of sorts; inside (atomic) formulas, variables/constants/parameters are only allowed to be plugged in a relation symbol if they come from the correct sort.

5. Inside (atomic) formulas, equality is only allowed between variables, constants, and parameters coming from the same sort.

Especially if you are already familiar with first-order logic (or if you looked into [End01, Section 4.3]), you may have observed that it is possible to "code" a multi-sorted structure inside a single-sorted one by using 1-ary relation symbols instead of sorts and writing down a suitable theory to avoid pathologies (for instance, to guarantee that the predicates are interpreted as disjoint sets).

This yields something very similar to multi-sorted logic, but with one important difference. Assume we perform such a "translation" on a language with infinitely many sorts, say $(S_i)_{i\in I}$, say "translated" as predicates $(P_i)_{i\in I}$. In some models of the translation there may be points which are not in any P_i, while in the multi-sorted version every point of a model must belong to a (unique) sort. The problem is that, if I is infinite, and T is a theory saying that the P_i are pairwise disjoint, then there will always be models of T containing points which are not in any P_i. This can be proven easily by using *compactness* (which, by the way, is about time to introduce), see Exercise 0.4.5.

Many-sorted languages are useful, but stating results for them tends to complicate notation and terminology. For this reason, we will mostly state results in the single-sorted case. The generalisation to the multi-sorted case is usually done with essentially the same proofs.

Assumption 0.3.2. Unless otherwise stated, everything below will be single-sorted.

0.4 Building models: basic techniques

0.4.1 Using magic

Theorem 0.4.1 (Compactness Theorem)**.** Let Φ be a set of sentences. Then Φ is consistent if and only if every finite $\Phi_0 \subseteq \Phi$ is consistent.

In model theory, this theorem is used all over the place. Its name comes from the fact that, after some rephrasing, it is equivalent to saying that certain topological spaces we will encounter later are compact.

We will not see the proof *of* compactness here (see the literature, or a course in logic), but here is a typical proof *by* compactness.

Example 0.4.2. There exists an elementary extension $\mathcal{M} \succeq \mathcal{R}$ containing an element m with $m > \mathbb{R}$.

Proof. Let $L = L_{\mathrm{oag}}(\mathbb{R}) \cup \{c\}$, where c is a new constant symbol. Consider the set of L-sentences

$$\Phi \coloneqq \mathrm{ED}(\mathcal{R}) \cup \{c > r \mid r \in \mathbb{R}\}$$

If $\Phi_0 \subseteq \Phi$ is finite, then it contains only finitely formulas of the form $c > r$. Let $r_0 \in \mathbb{R}$ be larger than all these finitely many r. Expand $\mathcal{R}_{\mathbb{R}}$ to an L-structure $\mathcal{S}$ by interpreting $c^{\mathcal{S}} \coloneqq r_0$. Then, by construction, $\mathcal{S} \vDash \Phi_0$, so, by definition, Φ_0 is consistent.

By compactness, Φ is consistent, hence there exists $\mathcal{M}' \vDash \Phi$. Let $\mathcal{M} := \mathcal{M}' \restriction L_{\mathrm{oag}}(\mathbb{R})$. By construction, $\mathcal{M}' \vDash \mathrm{ED}(\mathcal{R})$, so by Proposition 0.2.49 there is an elementary embedding $\mathcal{R} \to \mathcal{M}$, which we may assume, for notational convenience, to be the inclusion. Now let $m := c^{\mathcal{M}'}$. By construction, for all $r \in \mathbb{R}$ we have $\Phi \vDash c > r$, hence $\mathcal{M}' \vDash c > r$, that is, $\mathcal{M}' \vDash m > r$, hence $\mathcal{M} \vDash m > r$, and we are done. □

In this proof, I have freely confused $r \in \mathbb{R}$ with the corresponding constant symbol in $L_{\mathrm{oag}}(\mathbb{R})$ standing for it, but besides that I have been quite pedantic, and spelled out explicitly all the naming new constants and taking reducts. These steps of compactness proofs are usually very easy, and we will from now on omit them.

As is typical with compactness arguments, the proof above tells us very little about the structure we have proven to exist, but at least it allows us to conjure one basically out of thin air (hence the title of this subsection). Here is another standard compactness argument which allows us to conjure quite large things.

Corollary 0.4.3 (Upward Löwenheim–Skolem Theorem)**.** *Let T be a theory such that, for every $n \in \mathbb{N}$, there is $\mathcal{M} \vDash T$ of cardinality at least n. Then, for every cardinal κ, there is $\mathcal{M} \vDash T$ of cardinality at least κ.*

Furthermore, if T has an infinite model $\mathcal{M}_0$, we may also require that $\mathcal{M} \succeq \mathcal{M}_0$.

By the way, the cardinality of $\mathcal{M}$ is, by definition, the cardinality of M.[21]

Proof. Expand the language L of T to L' by adding new constant symbols $\{c_\alpha \mid \alpha < \kappa\}$, and let Φ be the set of L'-sentences

$$\Phi := T \cup \{c_\alpha \neq c_\beta \mid \alpha < \beta < \kappa\}$$

Every finite subset of Φ can only mention finitely many c_α, hence *by compactness* and our assumptions on T, the set Φ has a model. Its reduct to L is the required $\mathcal{M}$.

For the "furthermore" part, argue as above, but replacing T with $\mathrm{ED}(\mathcal{M}_0)$. □

Here is another standard fact which can be proven by using compactness.

Exercise 0.4.4. Let T be an L-theory. Prove that the class of models of T, together with elementary embeddings, has the *amalgamation property*: whenever A, B_0, B_1 are models of T, and $f_i \colon A \to B_i$ are elementary embeddings, there are $C \vDash T$ and elementary embeddings $g_i \colon B_i \to C$ such that $g_0 \circ f_0 = g_1 \circ f_1$.

$$\begin{tikzcd}
& B_0 \arrow[dr, dashed, "g_0"] & \\
A \arrow[ur, "f_0"] \arrow[dr, "f_1"'] & & C \\
& B_1 \arrow[ur, dashed, "g_1"'] &
\end{tikzcd}$$

[21] In the multi-sorted case, it's the sum of the cardinalities of the domains of each sort.

Exercise 0.4.5. Suppose that L contains infinitely many unary relation symbols $(P_i)_{i\in I}$, and that T is an L-theory where the P_i are nonempty and disjoint, that is, such that,

1. for every $i \in I$, we have $T \vdash \exists x\ P_i(x)$, and

2. for every $i \neq j \in I$, we have $T \vdash \neg\exists x\ (P_i(x) \land P_j(x))$.

Prove that there are some $\mathcal{M} \vDash T$ and $m \in M$ such that, for all $i \in I$, we have $\mathcal{M} \vDash \neg P_i(m)$.

Here is a small list of some more (fairly standard) things you can prove with compactness and some thinking. If this is your first encounter with the Compactness Theorem, it is a good idea to try doing these exercises, and maybe to look for more in the literature.

Exercise 0.4.6. Let $L_{\mathrm{grp}} = \{\cdot, e, ^{-1}\}$. There is no $L_{\mathrm{grp}} = \{\cdot, e, ^{-1}\}$-theory whose models are precisely the groups where every element is torsion (i.e. has finite order).

Exercise 0.4.7. There is no L_{graph}-theory whose models are precisely the graphs of finite diameter.

Exercise 0.4.8. There is an elementary extension of $(\mathbb{Z}, <)$ in which we may embed $(\mathbb{R}, <)$. Can such an embedding be elementary?

Exercise 0.4.9. Fix a theory T, a formula $\varphi(x)$, and suppose that, for every $M \vDash T$, the set $\varphi(\mathcal{M})$ is finite. Then there is $n \in \omega$ such that, for every $M \vDash T$, the set $\varphi(\mathcal{M})$ has size at most n.

Exercise 0.4.10. Let G be a graph and $n \in \omega$. Then G is colourable with n colours if and only if each of its finite induced subgraphs is.

Exercise 0.4.11. Every partial order extends to a linear order.

0.4.2 Using bookkeeping

The "magic" from the previous subsection (that is, compactness) is very useful when we want to build "large enough" objects. Sometimes, we want things not to be *too* large, and in that case different tools are needed. One of these, is the *Downward Löwenheim–Skolem Theorem.* The technique used below to prove it is one of those ideas that will come handy again from time to time.

Definition 0.4.12. The *cardinality* $|L|$ of a language L is the cardinality of the set of L-formulas. The *cardinality* $|T|$ of an L-theory T is the same as $|L|$.

So, for example, the cardinality of $L_{\mathrm{oag}} = \{+, 0, -, <\}$ equals $\aleph_0$, and so does the cardinality of the empty language (where the only atomic formulas are equalities between variables). If T is the empty theory in the empty language, we still have $|T| = \aleph_0$.

Sometimes we will want to consider languages which have finitely many symbols. In that case, we will simply say that "L is finite".

By the way, I hope that by now the difference between an L-structure $\mathcal{M}$ and its domain M is clear enough to introduce some abuse of notation.

Notation 0.4.13. From now on, we will freely use the same symbol (typically "M") to denote both a structure and its domain.[22] We will also start writing tuples as single letters, as in $y = (y_0, \dots, y_{|y|-1})$, where $|y|$ is the length of y.

Theorem 0.4.14 (Downward Löwenheim–Skolem)**.** Let M be an L-structure and $A \subseteq M$. There is an elementary substructure $M_0 \preceq M$ with $A \subseteq M_0$ and $|M_0| \leq |A| + |L|$.

Proof. We do an inductive construction, starting with $B_0 \coloneqq A$. For every L-formula $\varphi(x, y)$ with $|x| = 1$ and tuple $b \in B_n^{|y|}$, if there is $m \in M$ such that $M \vDash \varphi(m, b)$, put one such m in B_{n+1}. Note that $|B_{n+1}| \leq |L(B_n)| = |B_n| + |L|$, hence inductively $|B_{n+1}| \leq |A| + |L|$. Therefore, $M_0 \coloneqq \bigcup_{n\in\omega} B_n$ has the required cardinality.

By the Tarski–Vaught test (Theorem 0.2.51) we need to check that, whenever b is a tuple from M_0 and $M \vDash \exists x\ \varphi(x, b)$, then there is $a \in M_0$ such that $M \vDash \varphi(a, b)$. Since b is a finite tuple from $M_0 = \bigcup_{n\in\omega} B_n$, it must be contained in some B_n, and by construction we can find the required a inside B_{n+1}. $\square$

The following exercises can be solved with a combination of magic and bookkeeping.

Exercise 0.4.15. Let M be an infinite L-structure. If $A \subseteq M$ and κ is a cardinal with $|A| + |L| \leq \kappa \leq |M|$, then there is $M_0 \preceq M$ with $A \subseteq M_0$ and $|M_0| = \kappa$.

Exercise 0.4.16. Let M be an infinite L-structure. If κ is a cardinal with $\kappa \geq |M| + |L|$, then there is $N \succeq M$ with $|N| = \kappa$.

Exercise 0.4.17 (Vaught's test)**.** Let T be an L-theory with no finite models. If there is a cardinal $\kappa \geq |L|$ such that T has a unique model of cardinality κ up to isomorphism, then T is complete.

Exercise 0.4.18. Prove that the theory of infinite sets (defined in Example 0.2.26) is complete.

[22] If you are particularly allergic to abuses of notation, you may have noticed that, with this convention, if now we write, for example, $A \subseteq M$, it's not clear anymore if we mean that A is contained in the domain of M, or that A is a substructure of M. Fortunately, this is very unlikely to create problems, cf. Example 0.2.39.

Chapter 1

Five things everyone should see at least once in a model theory course (Corollary 1.3.2 will SHOCK you!!!)

1.1 Normal forms

Here are some standard (and quite useful) logic facts.

Fact 1.1.1. Every formula can be put in *prenex normal form*. That is, every L-formula $\varphi(x_0, \dots, x_n)$ is equivalent to one of the form

$$Q_0 y_0\, Q_1 y_1 \dots Q_m y_m\, \theta(x_0, \dots, x_n, y_0, \dots, y_m,)$$

where the Q_i are either $\exists$ or $\forall$ and θ is quantifier-free.

Notation 1.1.2. If φ is a formula, denote $\varphi^0 := \neg\varphi$ and $\varphi^1 := \varphi$

Fact 1.1.3. Every boolean combination of the formulas $\varphi_0, \dots, \varphi_n$ is equivalent to one in *disjunctive normal form*, that is, of the form $\bigvee_{i<m} \bigwedge_{j<k_i} \varphi_{\alpha_{i,j}}^{\beta_{i,j}}$, and to one in *conjunctive normal form*, that is, of the form $\bigwedge_{i<\ell} \bigvee_{j<h_i} \varphi_{\alpha_{i,j}}^{\beta_{i,j}}$.

Definition 1.1.4. A formula is *basic* iff it is atomic or the negation of an atomic formula. A formula is *negation normal* iff it is in the closure of basic formulas under $\exists, \forall, \wedge, \vee$.

In other words, the negation normal formulas are those where the only occurrences of $\neg$ are immediately before atomic formulas. The normal forms above in particular imply:

Corollary 1.1.5. Every formula is equivalent to a negation normal one.

1.2 Automorphisms

We still cannot prove Remark 0.2.19, but we can already prove something weaker.

Notation 1.2.1. If $\mathcal{M}$ is an L-structure and $A \subseteq M$, we denote by $\operatorname{Aut}(\mathcal{M}/A)$ the pointwise stabiliser of A, that is, the group of automorphisms f of $\mathcal{M}$ such that, for every $a \in A$, we have $f(a) = a$.

Exercise 1.2.2. Let $X \subseteq M^n$ be A-definable in $\mathcal{M}$. Then every automorphism fixing A pointwise fixes X setwise. That is, if $f \in \operatorname{Aut}(\mathcal{M}/A)$, then $f(X) = X$.[1]

Corollary 1.2.3. The set $\mathbb{Z}$ is not $\emptyset$-definable in $\mathcal{R}$.

Proof. For any positive $\lambda \in \mathbb{R}$, the map $x \mapsto \lambda \cdot x$ is an automorphism of $\mathcal{R}$. If $\lambda \neq 1$, this automorphism does not fix $\mathbb{Z}$ setwise. $\square$

1.3 Some consequences of Löwenheim–Skolem

Corollary 1.3.1. There is $M \preceq \mathcal{R}$ which is not complete.

Proof. By Löwenheim–Skolem, there is a countable $M \preceq \mathcal{R}$ with $\mathbb{Q} \subseteq M$. Since M is an ordered abelian group embedded in the reals, it must be Archimedean, and it follows easily by considering the group generated by any $m \in M \setminus \{0\}$ that M is unbounded in $\mathbb{R}$. Therefore, if $r \in \mathbb{R} \setminus M$, the set $\{m \in M \mid m < r\}$ is bounded in M. But, since M includes $\mathbb{Q}$, it is dense in $\mathbb{R}$, hence $\{m \in M \mid m < r\}$ has no supremum in M. $\square$

You may be wondering why we didn't take $M = \mathbb{Q}$ directly, instead of invoking Löwenheim–Skolem. The answer is that, at the moment, we do not known whether $\mathbb{Q} \preceq \mathcal{R}$. This is in fact true, and we will soon develop tools to prove it.

As usual with clickbaits, the fact below is probably something you have already heard.

Corollary 1.3.2 (Skolem paradox)**.** If ZFC has a model, then it has a countable one.

The reason this is called a paradox, is that ZFC proves the existence of uncountable sets. The catch here is that, if $M \vDash \mathsf{ZFC}$ is countable, and $a \in M$ is such that $M \vDash$ "a is uncountable", the only thing we can conclude is that M has no bijection between a and its set of natural numbers. But, of course, this does not prevent a bijection between ω and $\{b \in M \mid M \vDash b \in a\}$ to exist outside of M.

By the way, using compactness and Löwenheim–Skolem, we can make the example above even more pathological, and find a countable $M \vDash \mathsf{ZFC}$ which contains nonstandard natural numbers, that is, elements a such that $M \vDash$ "a is a natural number" but for every $n \in \omega$ we have $M \vDash a > n$. And even containing an infinite descending membership chain $a_0 \ni a_1 \ni a_2 \ni \ldots$; the reason this does not contradict the fact that M satisfies the Axiom of Foundation is that there will be no $b \in M$ such that $\{a_i \mid i < \omega\} = \{a \mid M \vDash a \in b\}$.

[1] Of course here the notation is being abused yet again: we should replace f with the map $(f, f, \ldots, f) \colon M^n \to M^n$.

1.4 Taking unions

Definition 1.4.1. A poset $(I, <)$ is *upward directed* iff for every $i, j \in I$ there is $k \in I$ such that $k \geq i$ and $k \geq j$.

Example 1.4.2. All linear orders are upward directed.

Proposition 1.4.3. Let $(I, <)$ be an upward directed poset, and $(M_i \mid i \in I)$ an I-sequence of L-structures such that, if $i < j$, then $M_i \subseteq M_j$. The union of the domains of the M_i can be uniquely made into an L-structure $M \coloneqq \bigcup_{i \in I} M_i$ such that every M_i is a substructure of M.

Proof. If R is a relation symbol and $a_0, \ldots, a_n \in M$, then each a_j belongs to some M_{i_j}. Since I is upward directed, an easy induction shows that there is $i \in I$ such that $i_0, \ldots, i_n \leq i$. By assumption, $a_0, \ldots, a_n \in M_i$. We set $M \vDash R(a_0, \ldots, a_n)$ iff $M_i \vDash R(a_0, \ldots, a_n)$. Of course we need to check that this does not depend on i, but this follows again by our assumptions: if j satisfies the same assumptions as i above, then there is $k \geq i, j$. Both M_i, M_j are substructures of M_k, hence $M_i \vDash R(a_0, \ldots, a_n) \iff M_k \vDash R(a_0, \ldots, a_n) \iff M_j \vDash R(a_0, \ldots, a_n)$.

One may use an analogous argument to define the interpretations of function symbols and constant symbols in M. Or, we can use the following standard trick.

We reduce to the previous case by assuming that the language is relational. This is done by changing the language, from L to L', say, by replacing every n-ary function symbol f of L with an $n+1$-ary relation symbol of L', to be interpreted as the graph of f, and every constant symbol of L by a 1-ary predicate, to be interpreted as a singleton. You may object that this involves checking that all the notions we are interested in (substructure, elementary substructure...) are preserved by this translation, and that the conclusion may be translated back. It is a good idea to convince yourself that this is indeed true, since this kind of trick is used fairly often.[2]

Finally, the condition that M_i is a substructure of M easily implies uniqueness. $\square$

If we sprinkle a modicum of Tarski–Vaught test in the proof of the previous proposition, we obtain the analogous statement for elementary embeddings. Spelling out the details of the proof is left as an exercise.

Exercise 1.4.4. Let $(I, <)$ be an upward directed poset, and let $(M_i \mid i \in I)$ be an I-sequence of L-structures such that, if $i < j$, then $M_i \preceq M_j$. Let $M \coloneqq \bigcup_{i \in I} M_i$. Then, for every $i \in I$, we have $M_i \preceq M$.

1.5 Finite structures

A lot of things in this course are stated for theories with infinite models, e.g. Exercise 0.4.17. The reason is that on finite structures, by the proposition below, a lot of the questions we will consider have trivial answers.

[2] Careful though, since some notions *do* change after this translation. An example is the notion substructure, hence that of "substructure generated". Furthermore, while the translation *preserves* substructures, it does not necessarily *reflect* them: a substructure in the original language is automatically a substructure in the translated one, but the converse need not hold.

Proposition 1.5.1. If M, N are L-structures and M is finite, then $M \equiv N \iff M \cong N$.

Proof. The implication $\Leftarrow$ is easy, and does not need finiteness. Suppose that $|M| = n$. We first prove $\Rightarrow$ in a special case, but in a stronger form.

Claim 1.5.2. Assume that L has only finitely many symbols. Then there is an L-sentence φ_M such that, if $N \vDash \varphi_M$, then $N \cong M$.

Proof of the Claim. Common sense dictates that this is the kind of "obvious but boring" thing that is usually left to the reader, since it is usually easier (and possibly instructive) to convince oneself that such a formula can be written, than to write it explicitly. Anyway, today I happened to leave common sense at home.

If $n = 0$, then the required sentence is $\forall x\ x \neq x$. Otherwise, the idea is to exploit finiteness of L to write a sentence saying "there are exactly n elements and they satisfy the diagram of M". Define the formula

$$\psi_n(x_0, \dots, x_{n-1}) \coloneqq \left(\bigwedge_{i<j<n} x_i \neq x_j\right) \land \left(\forall y \bigvee_{i<n} y = x_i\right)$$

Enumerate the elements of M in a tuple $a = (a_0, \dots, a_{n-1})$.[3] Note that $M \vDash \psi_n(a)$ and, conversely, if $N \vDash \psi_n(b)$, then $b = (b_0, \dots, b_{n-1})$ is an enumeration of all elements of N.

Observe that whenever N satisfies the sentence[4] $\exists x\ \psi_n(x)$ then we must have $|N| = n$. Of course this is not enough to guarantee the existence of an isomorphism $M \to N$, so we need a longer formula.

Let φ_M be $\exists x_0, \dots, x_{n-1}\ \psi_M(x_0, \dots, x_{n-1})$, where ψ_M is defined below.[5]

$$\begin{aligned}
\psi_M \coloneqq{} & \psi_n(x_0, \dots, x_{n-1}) \\
& \land \bigwedge_{\substack{i<n \\ c \in L_{\mathrm{c}} \\ M \vDash c = a_i}} c = x_i \\
& \land \bigwedge_{f \in L_{\mathrm{f}}} \left(\bigwedge_{\substack{i<n \\ h\colon \operatorname{ar}_L(f) \to n \\ M \vDash f(a_{h(0)}, \dots, a_{h(\operatorname{ar}_L(R)-1)}) = a_i}} f(x_{h(0)}, \dots, x_{h(\operatorname{ar}_L(R)-1)}) = x_i \right) \\
& \land \bigwedge_{R \in L_{\mathrm{r}}} \left(\bigwedge_{\substack{h\colon \operatorname{ar}_L(R) \to n \\ M \vDash R(a_{h(0)}, \dots, a_{h(\operatorname{ar}_L(R)-1)})}} R(x_{h(0)}, \dots, x_{h(\operatorname{ar}_L(R)-1)}) \right. \\
& \left. \land \bigwedge_{\substack{h\colon \operatorname{ar}_L(R) \to n \\ M \vDash \neg R(a_{h(0)}, \dots, a_{h(\operatorname{ar}_L(R)-1)})}} \neg R(x_{h(0)}, \dots, x_{h(\operatorname{ar}_L(R)-1)}) \right)
\end{aligned}$$

[3]Enumerations are assumed to be without repetitions unless otherwise stated.

[4]Recall that lower case letters are allowed to be tuples. I will not recall this further.

[5]For ease of notation, we use the identification $m = \{0, \dots, m-1\}$.

If $N \vDash \varphi_M$, there are $b_0, \ldots, b_{n-1}$ such that $N \vDash \psi_M(b_0, \ldots, b_{n-1})$, and by construction the map $a_i \mapsto b_i$ is an isomorphism. $\square_{\text{CLAIM}}$

In order to reduce the general case to that where L is finite, we use (of course) compactness. Consider the language $L(N) \cup \{c_0, \ldots, c_{n-1}\}$, where the c_i are new constant symbols. In this language, consider the theory

$$\mathrm{ED}(N) \cup \bigcup_{\substack{L_0 \subseteq L \\ L_0 \text{ finite}}} \psi_{M \restriction L_0}(c_0, \ldots, c_{n-1})$$

By the Claim and compactness, this theory is consistent, hence has a model $\tilde{N}$. The restriction $N' := \tilde{N} \restriction L$ is an elementary extension of N, but since N satisfies $\exists x\ \psi_n(x)$, so does N'. Since N is a substructure of N' and both have the same finite cardinality, we must have $N = N'$. It follows that the map sending $a_i \mapsto c_i^{\tilde{N}}$ is the required isomorphism. $\square$

Finite structures will still play an auxiliary role every now and then, but usually we will not look at their complete theories. This does not mean that model theory has nothing to say about finite structures: *finite model theory* is related to questions in computer science, especially in the area of computational complexity. See for example [EF95].

Chapter 2

First-order quantifiers and where to eliminate them

2.1 The first back-and-forth proof

What we do in this section may *prima facie* look completely unrelated to the title of this chapter, or to model theory in general, for that matter. Except it is very much not, as we will see later. For now, observe that we are proving that a certain theory (defined below) has only one countable model. The main focus is not on the theorem itself, but on its proof.

Definition 2.1.1. Let $L = \{<\}$, where $<$ is a binary relation symbol. The theory DLO of *dense linear orders without endpoints* has the following axioms:

1. $<$ is a *strict order*: an irreflexive, transitive relation;[1]

2. $<$ is *linear*: $\forall x, y\ ((x < y) \lor (x = y) \lor (x > y))$;[2]

3. $<$ has no *endpoints*: it has no maximum and no minimum;

4. $<$ is *dense*: $\forall x, y\ \big((x < y) \to (\exists z\ (x < z < y))\big)$.

Ok, the above definition has an hidden statement: I said "the *theory* DLO", so we should check it has a model. But, clearly, $(\mathbb{Q}, <) \vDash \mathsf{DLO}$.

Legend has it that the first back-and-forth proof was by Cantor, who invented the method to prove the theorem below. Except this is false, and Cantor managed to prove it by only going "forth". Also, I have no idea whether the proof below is the first proof by back-and-forth ever written, but nowadays it is usually the first one people see. Anyway, here is the proof.

Theorem 2.1.2 (Cantor)**.** All countable dense linear orders with no endpoints are isomorphic (to $(\mathbb{Q}, <)$).

[1] I leave it as an (easy) exercise to write these as first-order L-sentences. Also, note that the conjunction of irreflexivity and transitivity implies asymmetry.

[2] Of course $\lor$ is associative, so we may use fewer parenthesis. But I guess at some point above I promised not to stress such pedantries anymore.

Proof. Let $(M,<)$ and $(N,<)$ be countable dense linear orders with no endpoints, viewed as L-structures with $L=\{<\}$. Since they are dense (or, if you prefer, since they have no endpoints), M and N must both be infinite. Fix enumerations $(a_i)_{i<\omega}$ of M and $(b_j)_{j<\omega}$ of N. We build an isomorphism $f\colon M\to N$ inductively, by extending *partial isomorphisms*.

Start with f_0 being the empty function. If you prefer, f_0 is an isomorphism between the empty substructure of M and the empty substructure of N. We inductively define f_n in such a way that, for every $n\in\omega\setminus\{0\}$,

1. $f_n\colon A_n\to B_n$, where A_n is a finite substructure of M and B_n is a finite substructure of N;
2. $A_n\subseteq A_{n+1}$, $B_n\subseteq B_{n+1}$, and $f_n\subseteq f_{n+1}$;
3. f_n is an isomorphism of L-structures;
4. if $n=2m$, then $a_m\in A_n$;
5. if $n=2m+1$, then $b_m\in B_n$.

Suppose we manage to do this for every $n\in\omega$. If you think about it for ≈ 30 seconds, you will realise that this is enough to conclude. But, to be more formal:

Because $A_n\subseteq A_{n+1}$, the union $\bigcup_{n\in\omega}\operatorname{graph}(f_n)$ is the graph of a function, call it f, with domain a subset of M and codomain N. In fact, by Item 4 its domain is the whole M, and its image is the whole of N by Item 5. If $m<m'<\omega$, then $a_m,a_{m'}\in A_{2m'}$ and by Item 3 we have

$$\begin{aligned} M\vDash a_m<a_{m'} &\iff A_{2m'}\vDash a_m<a_{m'} \iff B_{2m'}\vDash f_{2m'}(a_m)<f_{2m'}(a_{m'})\\ &\iff N\vDash f_{2m'}(a_m)<f_{2m'}(a_{m'}) \iff N\vDash f(a_m)<f(a_{m'})\end{aligned}$$

Therefore, $f\colon M\to N$ is an isomorphism of L-structures.

Let us do this inductive construction then. Suppose we have build an isomorphism $f_{n-1}\colon A_{n-1}\to B_{n-1}$ as above. Write $A_{n-1}=\{a_{i_0}<a_{i_1}<\ldots<a_{i_k}\}$ and $B_{n-1}=\{b_{j_0}<b_{j_1}<\ldots<b_{j_k}\}$, and recall that for all $i\le k$ we have $a_i\in M$ and $b_i\in N$. If n is even, say $n=2m>0$, we take care of the "forth" part, that is, we extend f_{n-1} to $A_n\coloneqq A_{n-1}\cup a_m$. We have four cases:

a) If we already have $a_m\in A_{n-1}$, do nothing. Or, more formally, set $A_n\coloneqq A_{n-1}$, $B_n\coloneqq B_{n-1}$, and $f_n\coloneqq f_{n-1}$.

b) $a_m<a_{i_0}$. In this case, since N has no endpoints, in particular it has no minimum, hence there must be some $b\in N$ with $N\vDash b<b_{i_0}$. Send a_m to b. Or, more formally, put $A_n\coloneqq A_{n-1}\cup\{a_m\}$, $B_n\coloneqq B_{n-1}\cup\{b\}$, and $f_n\coloneqq f_{n-1}\cup\{(a_m,b)\}$.

c) $a_m>a_{i_k}$. Similarly, N has no maximum, so it contains some $b>b_{i_k}$ where to send a_m. Or, more formally,... well, ok, you know what needs to be written here.

d) There is $\ell<k$ with $M\vDash a_{i_\ell}<a_m<a_{i_{\ell+1}}$. Because N is dense, there is $b\in N$ with $N\vDash b_{i_\ell}<b<b_{i_{\ell+1}}$. Send a_m to b.

This takes care of the "forth" part. The "back" part, that is, the odd stages of the construction, are handled in the same way, with the roles of M and N reversed;[3] the only subtlety is that, for $n = 1$, there are no i_0, j_0. In that case, we start by simply choosing the preimage of b_0 arbitrarily, e.g. we can take $f_1(a_0) = b_0$. □

Here is a consequence of Cantor's theorem.

Corollary 2.1.3. DLO is complete.

Proof. By combining Theorem 2.1.2 with Exercise 0.4.17. □

In fact, we can squeeze more than just completeness from the *proof* of Theorem 2.1.2, and not just for dense linear orders. The rest of the chapter performs this squeezing.

Exercise 2.1.4. Prove that every countable linear order embeds into $(\mathbb{Q}, <)$.

2.2 Quantifier-free types

Let us look at the proof of Theorem 2.1.2. The crucial step was replicating the "position" of a_m with respect to A_{n-1}. The word "position", makes perfect sense in linear orders, but in general we will need something more adequate.

Definition 2.2.1. Let M be an L-structure, $A \subseteq M$, and $a = (a_0, \ldots, a_n)$ a tuple in M. Fix variables $x_0, \ldots, x_n$. The *quantifier-free type of a over A in M* is the set of formulas

$$\mathrm{qftp}^M(a/A) \coloneqq \{\varphi(x_0, \ldots, x_n) \in L(A) \mid \varphi \text{ quantifier-free, } M \vDash \varphi(a_0, \ldots, a_n)\}$$

In other words, $\mathrm{qftp}^M(a/A)$ is obtained by taking all quantifier-free formulas $\varphi(a_0, \ldots, a_n)$ true in M, and replacing each a_i with a free variable x_i. The formula $\varphi(x_0, \ldots, x_n)$ is allowed to contain parameters from A.

Remark 2.2.2. Of course there is nothing special in the variables $x_0, \ldots, x_n$, and we may have used $y_0, \ldots, y_n$ instead; for many purposes, the quantifier-free types obtained in these two ways are identified.[4] You can also think of $\mathrm{qftp}^M(a/A)$ as the collection of A-definable subsets of M containing a. This has the advantage of not needing to fix variables, but it makes it more difficult to compare quantifier-free types over different structures. Both points of view are useful.

At any rate, the key property we exploited in the proof was the following.

Exercise 2.2.3. Let a, b be tuples of the same length from M, N respectively.

[3] It is usual, in back-and-forth proofs, to have symmetrical hypothesis, hence to only do the "forth" part and say that the "back" part is analogous. So, why is this not called "forth-and-back"? I guess that's because "back-and-forth" is an existing English sentence but, if you want, the first step where we actually send elements somewhere is step 1, which is a "back".

[4] But sometimes being careless with identifications may result in trouble; more about this at the end of Section 2.5.

1. Assume that $\mathrm{qftp}^M(a/\emptyset) = \mathrm{qftp}^N(b/\emptyset)$. Check that $a_i \mapsto b_i$ induces an isomorphism between the substructure of M generated by[5] a and the substructure of N generated by b, defined in the obvious manner: e.g. if f is a function symbol and c a constant symbol then $f(a_0, a_1, c)$ is sent to $f(b_0, b_1, c)$.

2. Check that, conversely, if this map is well-defined[6] and an isomorphism, then $\mathrm{qftp}^M(a/\emptyset) = \mathrm{qftp}^N(b/\emptyset)$.

3. Check that, if in addition A is a substructure of both M and N, then $\mathrm{qftp}^M(a/A) = \mathrm{qftp}^N(b/A)$ if and only if the map sending $a_i \mapsto b_i$ induces (in the way as above) a well-defined isomorphism between the substructure of M generated by[7] Aa and the substructure of N generated by Ab.

In other words, in the proof we did the following. Use f_{n-1} to identify A_{n-1} with B_{n-1}. Take the next point a_m to be considered, and look at $p(x) \coloneqq \mathrm{qftp}^M(a_m/A)$. Find, inside N, a *realisation* of $p(x)$.

Definition 2.2.4. Let N be an L-structure, A a subset of N, and $p(x)$ a quantifier-free type over A. A tuple b in N with $|b| = |x|$ is said to *realise* $p(x)$, written, $b \vDash p(x)$, iff $N \vDash p(b)$. That is, for every $\varphi(x) \in p$, we have $N \vDash \varphi(b)$.

Example 2.2.5. Let $M = (\mathbb{Q}, <)$ and $A = \{-1/n \mid n \in \omega \setminus \{0\}\}$. Let $p(x) \coloneqq \mathrm{qftp}^M(2/A)$. Then $3 \vDash p(x)$. In fact, all positive rationals have the same quantifier-free type over A. More generally, $b, c \in M$ have the same quantifier-free type over A if and only if for every $a \in A$ we have $M \vDash b \geq a \iff M \vDash c \geq a$. In other words, $\mathrm{qftp}^M(a/A) = \mathrm{qftp}^M(b/A)$ if and only if a, b fill the same cut of A (in the degenerate case where $a \in A$ by the cut of a in A we mean just $\{a\}$).

2.3 The Ra(n)do(m) graph

Before developing the theory further, here is a good exercise to get familiar with back-and-forth. Work in $L_{\mathrm{graph}} = \{E\}$.

Definition 2.3.1. Let T_{rg} be the set of L_{graph}-formulas:

1. E is a graph (i.e. irreflexive and symmetric).

2. For every $(n, m) \in \omega^2 \setminus \{(0, 0)\}$, the formula

$$\forall x_0, \ldots, x_{n-1}, y_0, \ldots, y_{m-1} \left(\bigwedge_{\substack{i<n \\ j<m}} x_i \neq y_j \right) \rightarrow \left(\exists z \left(\bigwedge_{i<n} E(x_i, z) \right) \wedge \left(\bigwedge_{j<m} \neg E(y_j, z) \right) \right)$$

[5] By the way, have I already said that there is another abuse of notation going on, where a tuple is sometimes treated as a set, as in "the substructure generated by a"? Formally, we should say "generated by $\{a_i \mid i < |a|\}$". Anyway, the important thing to keep in mind is that tuples *are* allowed repetitions (while sets are not).

[6] Think of what happens if a satisfies $f(x) = g(x)$ but b does not.

[7] That's right, another abuse of notation: AB stands for $A \cup B$. Composed with the previous abuse of notation, Aa means $A \cup \{a_i \mid i < |a|\}$.

In words, T_{rg} says that its models are graphs where, for every finite sets U, V, if $U \cap V = \emptyset$ then there is a point with an edge to all elements of U and no edge to any element of V.

Exercise 2.3.2.

1. Prove that T_{rg} is consistent.[8]

2. Prove that T_{rg} has a unique countable model (up to isomorphism).

The unique countable model of T_{rg} is known as the *Random Graph*, or *Rado Graph*.

3. Prove that every countable graph embeds into the Random Graph as an induced subgraph.

The name "Random Graph" is due to the following fact: fix a countable set, and put an edge between any two distinct points with fixed probability $0 < p < 1$, independently. Then, with probability 1, the resulting graph is (isomorphic to) the Random Graph.

2.4 Syntax: eliminating quantifiers by hand

Recall that, for technical convenience, we added to our logic a symbol $\bot$, which is a (quantifier-free) atomic sentence in every language, and it is always false. We also write $\top$ for $\neg\bot$.

Definition 2.4.1. The L-theory T has *quantifier elimination* iff, for every $n \in \omega$, and every L-formula $\varphi(x)$ with $|x| = n$, there is an L-formula $\psi(x)$ without quantifiers such that $T \vdash \forall x\ \varphi(x) \leftrightarrow \psi(x)$. An L-structure M has *quantifier elimination* iff $\mathrm{Th}(M)$ does.

Remark 2.4.2.

- Note that ψ is required to have the "same" (cf. Remark 0.2.12) free variables as φ.

- The semantical counterpart to this (syntactical) definition is: T has quantifier elimination if and only if every $\emptyset$-definable set is a boolean combination of sets defined by atomic formulas. See also Remark 2.4.7 for a more geometric interpretation.

- The reason we added $\top, \bot$ to the logic is that, otherwise, if L has no constant symbols, there are no quantifier-free sentences. This happens for example in the language of orders, or the language of graphs.

Below, we will see some methods to prove that a theory has quantifier elimination. But first, some examples.

Example 2.4.3. Let $L = \{+, 0, -, \cdot, 1, <\}$ and $T = \mathrm{Th}(\mathbb{R})$. Consider the formula

$$\varphi(x_0, x_1, x_2) \coloneqq \exists y\ (x_2 \cdot y^2 + x_1 \cdot y + x_0 = 0)$$

[8]Hint: do an inductive construction, or, if you like probability, see below.

where y^2 is an abbreviation for $y \cdot y$. Then $\varphi(x)$ is equivalent modulo T to the quantifier-free formula

$$\begin{aligned}&(x_2 = x_1 = x_0 = 0)\\ \vee\ &(x_2 = 0 \wedge x_1 \neq 0)\\ \vee\ &(x_2 \neq 0 \wedge x_1^2 - (1+1+1+1) \cdot x_2 \cdot x_0 > 0)\end{aligned}$$

Example 2.4.4. Let $L = \{+, 0, -, \cdot, 1\}$ and T be the L-theory of fields. It is easy (but it takes a while) to write an existential formula $\varphi(x_0, \ldots, x_{n^2-1})$ saying that the x_i are (in that order) the entries of an invertible $n \times n$ matrix. This φ is equivalent modulo T to a quantifier-free formula, saying that this matrix has nonzero determinant.

One way to eliminate quantifiers is to take them out one at a time by induction on formulas. Some steps are always the same: for example, if $\varphi(x)$ and $\psi(x)$ are quantifier-free, clearly so is $\varphi(x) \wedge \psi(x)$. The next lemma packages together all the easy steps, and tells us where we the actual work needs to go.

Definition 2.4.5. A formula $\psi(x)$ is *primitive* iff it is of the form $\exists y\ \bigwedge_{i<k} \varphi_i(x, y)$, where every φ_i is basic.

Lemma 2.4.6. Suppose that every primitive formula $\exists y\ \bigwedge_{i<k} \varphi_i(x, y)$ with $|y| = 1$ (and the $\varphi_i(x, y)$ basic) is equivalent modulo T to a quantifier-free formula. Then T has quantifier elimination.

Proof. By induction on formulas. If $\psi(x)$ is atomic, there is nothing to do. If ψ is of the form $\neg\varphi_0$, by induction there is a quantifier-free θ such that $T \vdash \forall x\ \varphi_0(x) \leftrightarrow \theta(x)$. Clearly, ψ is equivalent modulo T to $\neg\theta$, which is quantifier-free. The case where ψ is of the form $\varphi_0 \wedge \varphi_1$ is dealt with similarly.

We are left to deal with the case $\exists y\ \varphi(x, y)$, with $|y| = 1$. Inductively, $\varphi(x, y)$ is equivalent to a quantifier-free formula $\theta(x, y)$. Using disjunctive normal form, $\theta(x, y)$ is equivalent to a formula $\bigvee_i \bigwedge_j \varphi_{i,j}(x, y)$, with the $\varphi_{i,j}$ basic. Since $\exists y\ (\alpha(x, y) \vee \beta(x, y))$ is equivalent to $(\exists y\ \alpha(x, y)) \vee (\exists y\ \beta(x, y))$, we reduce to the case where $\theta(x, y)$ is a conjunction of basic formulas. But then $\exists y\ \theta(x, y)$ is primitive with $|y| = 1$, hence it is equivalent to a quantifier-free formula by assumption. □

Remark 2.4.7. Geometrically, the quantifier $\exists$ corresponds to a projection. By the previous lemma, quantifier elimination is equivalent to the following: if $X \subseteq M^{n+1}$ is an intersection of subsets of M^{n+1} sets defined by basic formulas, and we consider the projection $\pi\colon M^{n+1} \to M^n$ on the first n coordinates (say), then $\pi(X)$ can be written as a boolean combination of subsets of M^n defined by atomic formulas.

Let us look at one easy example of quantifier elimination "by hand".

Example 2.4.8. The theory of infinite sets has quantifier elimination.

Proof. By Lemma 2.4.6 and the fact that the only atomic formulas are of the form $x_i = x_j$, we just need to eliminate the quantifier from formulas $\varphi(x_0, \ldots, x_{n-1})$ of the form

$$\exists y \left(\bigwedge_{i \in I} y = x_i \wedge \bigwedge_{j \in J} y \neq x_j \wedge \bigwedge_{(k_0,k_1) \in K} x_{k_0} = x_{k_1} \wedge \bigwedge_{(h_0,h_1) \in H} x_{h_0} \neq x_{h_1}\right)$$

for some $I, J \subseteq n$ and $K, H \subseteq n \times n$. If $I \neq \emptyset$, say because $i_0 \in I$, we may discard $\exists y$, replace every occurrence of y by x_{i_0}, and obtain an equivalent formula of the same form as above, but where $I = \emptyset$. So we may assume $I = \emptyset$. Set $\psi(x) \coloneqq \bigwedge_{(k_0,k_1)\in K} x_{k_0} = x_{k_1} \wedge \bigwedge_{(h_0,h_1)\in H} x_{h_0} \neq x_{h_1}$. Because y does not appear in $\psi(x)$, we have that $\varphi(x)$ is equivalent to $\left(\exists y \bigwedge_{j\in J} y \neq x_j\right) \wedge \psi(x)$. Since $T \vdash \forall x\, \exists y \bigwedge_{j\in J} y \neq x_j$, we have that $\varphi(x)$ is equivalent to $\psi(x)$. □

Some proofs of quantifier elimination "by hand" are in [TZ12, Section 3.3].

This way of proving quantifier elimination can be very efficient, but in some cases using this technique may involve dealing with complicated formulas, several distinctions by cases, preliminary lemmas, etc.[9] So we better have more tools at our disposal.

2.5 Types: packaging formulas together

Perhaps counterintuitively, it turns out that sometimes it is easier to manage complete theories than single formulas. Complete theories are sets of sentences, while formulas $\varphi(x)$ are allowed free variables. If we are interested in formulas with free variables, and want to pass through complete theories, the standard trick is to introduce new constants c and replace $\varphi(x)$ with $\varphi(c)$.

In this order of ideas, quantifier elimination becomes: quantifier-free types are enough to determine a *complete type.*

Definition 2.5.1. Let T be an L-theory and $n \in \omega$. Let $c_0, \dots, c_{n-1}$ be new constant symbols.

1. A *partial n-type* is an $L \cup \{c_0, \dots, c_{n-1}\}$-theory containing T.

2. A *complete n-type* is a complete $L \cup \{c_0, \dots, c_{n-1}\}$-theory containing T.

Remark 2.5.2.

1. In the literature, the word "type" is used sometimes as a synonymous of "partial type" and sometimes as a synonymous of "complete type". We will go with the *second* convention.

2. Also, some authors allow partial types to be inconsistent (i.e., not a theory).

3. Soon we will concentrate on complete T, but the definition above allows to talk of types of incomplete theories, which we will need. For complete T, it also makes sense to talk of types over parameters. We will see this later.

4. A 0-type is the same as a completion of T.

[9]As a baby example, try to prove that the theory of the Random Graph eliminates quantifiers with an argument similar to that of Example 2.4.8. You will probably end up having to prove that in the Random Graph, if U, V are finite and $U \cap V = \emptyset$, then there are infinitely many x connected to all points of U and no point of V (the axioms only state the existence of *one* such x).

Lemma 2.5.3. Let $c_0, \dots, c_{n-1}$ be constant symbols not in L. Let T be an L-theory, and T' be the deductive closure of T in $L \cup \{c_0, \dots, c_{n-1}\}$. For all L-formulas $\varphi(x)$ with $|x| = n$, the following are equivalent:

1. $T \vdash \forall x\ \varphi(x)$
2. $T' \vdash \varphi(c)$.

Proof. Since $T \subseteq T'$ we immediately have (1) $\Rightarrow$ (2). We prove $\neg$(1) $\Rightarrow$ $\neg$(2). Suppose that $T \nvdash \forall x\ \varphi(x)$. This means that $T \cup \{\exists x\ \neg\varphi(x)\}$ is consistent, so it has a model M. This M is a model of T, and there is $a \in M^{|x|}$ such that $M \vDash \neg\varphi(a)$. Expand M to an L'-structure by interpreting $c_i^{M'} \coloneqq a_i$. By definition, $T \vdash T'$, hence $M' \vDash T' \cup \{\neg\varphi(c)\}$, and we have $\neg$(2). ☐

Corollary 2.5.4. Let $c_0, \dots, c_{n-1}$ be constant symbols not in L. Let T be an L-theory, and T' be the deductive closure of T in $L \cup \{c_0, \dots, c_{n-1}\}$. Let $\Phi(x)$ be a set of formulas $\varphi(x)$ with $|x| = n$. The following are equivalent.

1. $\Phi(c)$ is a partial type.
2. For every $\varphi_0(x), \dots, \varphi_m(x) \in \Phi(x)$, the set $T \cup \left\{\exists x\ \bigwedge_{i \le m} \varphi_i(x)\right\}$ is consistent.

Proof. By the previous lemma, compactness, and a pinch of logic. ☐

Definition 2.5.5. Let $M \vDash T$ and $a \in M^n$. The *type of a in M*, denoted by $\mathrm{tp}^M(a)$, is $\{\varphi(x_0, \dots, x_{n-1}) \in L \mid M \vDash \varphi(a_0, \dots, a_{n-1})\}$.

In other words, $\mathrm{tp}^M(a)$ is the collection of all L-formulas defining a set to which a belongs (in a fixed tuple of variables[10]).

Using Corollary 2.5.4, you can check that, by replacing every x_i in $\mathrm{tp}^M(a)$ with c_i, we obtain a type in the sense of Definition 2.5.1. A standard abuse of notation, to which we will immediately start conforming, is to confuse x_i with c_i, and write types with variables instead of extra constants. So, for example, we may say that $p(x) = \mathrm{tp}^M(a)$ is a type of T. The converse holds as well: all (complete!) types are types of tuples in some model:

Proposition 2.5.6. For every n-type $p(x)$ there are $M \vDash T$ and $a \in M^n$ such that $p(x) = \mathrm{tp}^M(a)$.

Proof. This is so trivial it almost hurts: by assumption a type is a complete $L \cup \{c_0, \dots, c_{n-1}\}$-theory containing T. Take a model M of this theory, set $a_i \coloneqq c_i^M$, and (obviously) take $a = (a_0, \dots, a_{n-1})$. ☐

This does *not* mean that every type is realised in every model. We will come back to this at length later on in the course.

Let us now look at an easy but important fact.

Exercise 2.5.7. If $M \preceq N$ and $a \in M^n$, then $\mathrm{tp}^M(a) = \mathrm{tp}^N(a)$.

Note that, if $|a| = n$, and $m < n$, then $\mathrm{tp}^M(a)$ decides in particular all the m-types of its subtuples of length m; for $m = 0$, this means that $\mathrm{tp}^M(a)$ implies $\mathrm{Th}(M)$, that is, it decides a completion of T.

At last, the theorem promised at the beginning of this section.

[10] Or maybe not. See Remark 2.2.2 and the last part of this section.

Theorem 2.5.8. The following are equivalent.

1. T has quantifier elimination.
2. For all models M, N of T, and all $n \in \omega$, whenever $a \in M^n$ and $b \in N^n$ are such that $\mathrm{qftp}^M(a/\emptyset) = \mathrm{qftp}^N(b/\emptyset)$, then $\mathrm{tp}^M(a) = \mathrm{tp}^N(b)$.

Proof. (1) $\Rightarrow$ (2) is an immediate consequence of the definitions, so let us focus on (2) $\Rightarrow$ (1).

Fix T and an L-formula $\psi(x)$, say with $|x| = n$, from which we want to eliminate quantifiers. If $T \vdash \neg\exists x\ \psi(x)$, then $\psi(x)$ it is equivalent to the quantifier-free formula $\perp$ and we are done. Otherwise, consider the set of quantifier-free consequences of $\psi(x)$

$$\Psi(x) := \{\theta(x) \text{ quantifier-free} \mid T \vDash \forall x\ \psi(x) \to \theta(x)\}$$

By definition, $\psi(x) \vdash \Psi(x)$, where this notation means that, for suitable constants c, we have $T \cup \{\psi(c)\} \vdash \Psi(c)$. The heart of the proof lies in the following claim.

Claim 2.5.9. $\Psi(x) \vdash \psi(x)$.

Proof of the Claim. If not, there is a model (M, a) of $T \cup \Psi(c) \cup \{\neg\psi(c)\}$, where a denotes the interpretation of c. Let us look at $\pi(x) := \mathrm{qftp}^M(a/\emptyset)$. By our hypothesis, $T \cup \pi(c)$ should imply a complete type. We will reach a contradiction by showing that this is not the case.

Subclaim 2.5.10. $T \cup \pi(x) \cup \{\psi(x)\}$ is consistent.

Proof of the Subclaim. Otherwise, by compactness, there is a finite conjunction $\bigwedge_{j<\ell} \varphi_j(x)$ of formulas in $\pi(x)$ such that $T \vdash \forall x\ (\bigwedge_{j<\ell} \varphi_j(x) \to \neg\psi(x))$. Taking the contrapositive, $T \vdash \forall x\ (\psi(x) \to \bigvee_{j<\ell} \neg\varphi_j(x))$. Since $\bigvee_{j<\ell} \neg\varphi_j(x)$ is quantifier-free, by definition it belongs to $\Psi(x)$. But now, on one hand, by choice of M and a, we have $M \vDash \Psi(a)$, and in particular $M \vDash \bigvee_{j<\ell} \neg\varphi_j(a)$. On the other hand, every $\varphi_j(x)$ belongs to $\pi(x) = \mathrm{qftp}^M(a/\emptyset)$, hence $M \vDash \bigwedge_{j<\ell} \varphi_j(a)$, a contradiction. $\square_{\text{SUBCLAIM}}$

Therefore, there is $(N, b) \vDash T \cup \pi(x) \cup \{\psi(x)\}$. As promised, this is a contradiction: $N \vDash \pi(b)$, that is, b satisfies the same quantifier-free formulas as a; by our hypothesis, this guarantees the same formulas, even with quantifiers, are satisfied by a (in M) and by b (in N); but $M \vDash \neg\psi(a)$ and $N \vDash \psi(b)$. $\square_{\text{CLAIM}}$

By the Claim and compactness, there is a finite conjunction $\bigwedge_{i<k} \psi_i(x)$ of formulas in $\Psi(x)$ such that $T \vdash \forall x\ \big(\bigwedge_{i<k} \psi_i(x) \to \psi(x)\big)$. By definition of Ψ, all the ψ_i are quantifier-free and $T \vdash \forall x\ \big(\psi(x) \to \bigwedge_{i<k} \psi_i(x)\big)$. We conclude that $\psi(x)$ is equivalent modulo T to the quantifier-free formula $\bigwedge_{i<k} \psi_i(x)$. $\square$

Types are one of the most used tools in model theory, and we will deal with them at great length later on in the course. Before we go back to quantifier elimination, we finish this section with some final remarks about types.

Another way to think about types is: a partial n-type in M is a filter on the boolean algebra of formulas $\varphi(x)$, with $|x| = n$, modulo being equivalent modulo T. In this identification, complete n-types correspond to ultrafilters on this algebra. We will not go into details, but if you want to read about it, this algebra is called the *Lindenbaum algebra*, or *Lindenbaum–Tarski algebra*.[11] For complete T, one may equivalently fix some $M \vDash T$ and look at the boolean algebra of definable subsets of M^n.

Finally, let me clarify an abuse of notation which may seem (and usually is) harmless, but may give you some headaches down the road. When we defined types, we might as well have used different constants, say d_i instead of c_i, and when replacing constants with variables we may have used y_i instead of x_i, and we would have ended up with essentially the same notions, (see also Remark 2.2.2). Therefore, if we strive for complete, pedantic precision, it would have probably been more correct to define types in some other, constant-free and variable-free way; for example, as equivalence classes of the relation "satisfying the same formulas" (coded via some suitable set-theoretic trick—those equivalence classes are proper class-sized).

Still, formulas and variables are very convenient to handle, but, in some situations, care is needed. For example: are $p(x_0, x_1)$ and $p(y_0, y_1)$ the same type or not? Usually these two types are identified, unless they are used jointly, e.g. to define a type $q(x_0, x_1, y_0, y_1)$ as $p(x_0, x_1) \cup p(y_0, y_1) \cup \{(x_0 = y_0) \wedge (x_1 = y_1)\}$. So one may say that types are really to be considered up to change of variables/constants, but we should be careful not to take quotients too early. This phenomenon is already present at the level of formulas: is $\varphi(x_0, x_1)$ the same as $\varphi(x_1, x_0)$? If, for instance, we want to write $T \vdash \forall x\ \varphi(x_0, x_1) \leftrightarrow \varphi(x_1, x_0)$ to say that the set defined by φ is symmetric with respect to the diagonal, we better not identify $\varphi(x_0, x_1)$ with $\varphi(x_1, x_0)$ too early. One may use "variable-free" presentations of types like the one in the previous paragraph (and even of formulas and partial types), but at a price: for example, defining the partial type q above becomes more cumbersome.

2.6 Semantics: eliminating quantifiers by back-and-forth

Sometimes, dealing with substructures is easier than dealing with formulas; for example, because we are doing model theory of some algebraic structures, and we want to exploit facts that the algebraists have already proven about them. In those cases, the main theorem of this section allows us to prove quantifier elimination by using the back-and-forth method.

Definition 2.6.1. Let M, N be L-structures.

1. A *partial isomorphism* between M and N is an isomorphism between a substructure $A \subseteq M$ and a substructure $B \subseteq N$.

2. A family F of partial isomorphisms between M and N has the *back-and-forth property* iff for every $f \in F$

(forth) for every $a \in M$ there is $g \in F$ with $a \in \operatorname{dom} g$ and $g \supseteq f$, and

(back) for every $b \in N$ there is $g \in F$ with $b \in \operatorname{im} g$ and $g \supseteq f$.

[11] If you are curious about Stone duality, now could be a good moment to read about it. Or you may wait until we talk about type *spaces*.

Theorem 2.6.2. Let T be an L-theory. Suppose that, for every $M_0 \vDash T$ and $N_0 \vDash T$, there are $M \succeq M_0$ and $N \succeq N_0$ such that the family of all partial isomorphisms between finitely generated substructures of M and N has the back-and-forth property. Then T eliminates quantifiers.

Proof. Towards a contradiction, assume this is not the case. By Theorem 2.5.8, there are finite tuples[12] $a \in M_0$ and $b \in N_0$ with

$$\mathrm{qftp}^{M_0}(a/\emptyset) = \mathrm{qftp}^{N_0}(b/\emptyset) \tag{2.1}$$

but

$$\mathrm{tp}^{M_0}(a) \neq \mathrm{tp}^{N_0}(b) \tag{2.2}$$

The last inequality must be witnessed by some L-formula; by Lemma 2.4.6 the offending formula may be taken of the form $\exists y\ \varphi(x, y)$, with $\varphi(x, y)$ quantifier-free and $|y| = 1$. We use the "forth" in "back and forth" to deal with the case when

$$M_0 \vDash \exists y\ \varphi(a, y) \qquad N_0 \vDash \neg\exists y\ \varphi(b, y) \tag{2.3}$$

The case where $M_0 \vDash \neg\exists y\ \varphi(a, y)$ but $N_0 \vDash \exists y\ \varphi(b, y)$ is dealt with in the same way, by using the "back" instead.

Since $M \succeq M_0$ and $N \succeq N_0$, by definition of $\preceq$ (and Exercise 2.5.7, if you want) (2.1), (2.2), and (2.3) still hold after replacing M_0 by M and N_0 by N.

Because $\mathrm{qftp}^{M}(a/\emptyset) = \mathrm{qftp}^{N}(b/\emptyset)$, by Exercise 2.2.3 the map sending $a_i \mapsto b_i$ extends to an isomorphism $f\colon A \to B$, where $A \subseteq M$ and $B \subseteq N$ are (finitely) generated by a, b respectively. Since $M \vDash \exists y\ \varphi(a, y)$, there is $d \in M$ such that $M \vDash \varphi(a, d)$. Let $\hat{A}$ be the substructure of M generated by ad. Because $\varphi(x, y)$ is quantifier-free, by Exercise 0.2.42 $\hat{A} \vDash \varphi(a, d)$. Clearly, $\hat{A}$ is finitely generated and contains A. By the "forth" property there is an isomorphism $g \supseteq f$ with domain $\hat{A}$. Let $\hat{B} \coloneqq \mathrm{im}(g)$; note that it is a substructure of N containing B. Since g is an isomorphism and $g(a) = b$, we have $\hat{B} \vDash \varphi(b, g(d))$. Again by Exercise 0.2.42, this yields $N \vDash \varphi(b, g(d))$, and in particular $N \vDash \exists y\ \varphi(b, y)$. This contradicts the fact that, by (2.3) and elementarity, $N \vDash \neg\exists y\ \varphi(b, y)$. □

Remark 2.6.3. Some comments and a spoiler:

1. It may (and will) happen that for some M_0 and N_0, for all $M \succeq M_0$ and $N \succeq N_0$, the family of partial isomorphisms between finitely generated substructures of M and N is empty. Vacuously, the empty family does have the back-and-forth property. Note that Theorem 2.6.2 does not need such families to be nonempty (if you don't believe me, check the proof).

2. It may (and will) happen that, even if T has quantifier elimination, the family of partial isomorphisms between finitely generated substructures of some M_0 and N_0 does not have the back-and-forth property. In other words, passing to an elementary extension is in general necessary.

3. If L is relational, one may avoid passing to an elementary extension by weakening the back-and-forth property; I won't elaborate here, but if you are interested search for *Ehrenfeucht–Fraïssé games.*

4. The converse of the previous theorem is also true; to prove it, one takes M, N to be *ω-saturated*, a notion we will introduce later.

[12] From now, I will start writing e.g. $a \in M_0$ instead of $a \in M_0^{|a|}$ whenever convenient.

2.7 Consequences: eliminating quantifiers for a purpose

In the next chapter, we will see some applications of quantifier elimination in concrete structures. Here we look at some more general consequences.

Theorem 2.7.1. Suppose that T is an L-theory such that

1. T eliminates quantifiers, and
2. for all models M, N of T, there is an L-structure A which embeds in both M and N.

Then T is complete.

Proof. We need to show that, for all models M, N of T, we have $M \equiv N$, so we fix an L-sentence φ, we assume that $M \vDash \varphi$, and we aim to show that $N \vDash \varphi$. By assumption, there is a quantifier-free sentence ψ such that $T \vdash \varphi \leftrightarrow \psi$. In particular, $M \vDash \psi$. Take A as in the assumptions of the theorem, and assume for notational convenience that the embeddings of A in M and N are inclusions. Because ψ is quantifier-free, $M \vDash \psi$ implies $A \vDash \psi$, which in turn implies $N \vDash \psi$. But N is a model of T, hence $N \vDash \varphi \leftrightarrow \psi$, so $N \vDash \varphi$. □

Remark 2.7.2.

1. In the proof above, we may have $A \vDash \neg\varphi$. This is due to the fact that A is not required to be a model of T, hence, in A the sentence $\varphi \leftrightarrow \psi$ need not hold.
2. We will encounter examples of incomplete theories with quantifier elimination; by the previous theorem, this can only happen if some pair of models of T share no common substructure, even up to embeddings (compare also with point 1 of Remark 2.6.3).
3. If L has no constant symbol[13], the empty structure is an L-structure, and a perfectly good A to use in this theorem.

If for some reason you only need to prove completeness of a theory, and don't care about quantifier elimination, the following exercise may come out handy.

Exercise 2.7.3.
1. Suppose that F is some[14] family of partial isomorphisms between M and N with the back and forth property. Prove that every $f \in F$ is an *elementary map*, that is, for every L-formula $\varphi(x)$ and $a \in (\operatorname{dom} f)^{|x|}$, we have $M \vDash \varphi(a) \iff N \vDash \varphi(f(a))$.[15]
2. Deduce that, if for all models M, N of T there is some *nonempty* F as above, then T is complete.

Remark 2.7.4. If T eliminates quantifiers, then it is *model complete*: namely, all embeddings between models of T are elementary.

[13] Those who have already read Section 2.8 may also want to assume that L has no 0-ary relation symbol.

[14] Not necessarily that of all partial isomorphisms between finitely generated substructures.

[15] This does not mean that f is an elementary embedding: we may have $\operatorname{dom} f \neq M$.

We finish the section with a characterisation.

Definition 2.7.5. A theory T is *substructure complete* iff for every $M \vDash T$ and every substructure $A \subseteq M$, the theory $T \cup \operatorname{diag}(A)$ is complete.

Theorem 2.7.6. The following are equivalent for an L-theory T.

1. T is substructure complete.
2. For every $M \vDash T$ and every finitely generated substructure $A \subseteq M$, the theory $T \cup \operatorname{diag}(A)$ is complete.
3. If A is a substructure of two models M, N of T, and a is a finite tuple from A, then $\operatorname{tp}^M(a) = \operatorname{tp}^N(a)$.
4. T has quantifier elimination.

Proof. This theorem has been, in a sense, already proven. In fact, (1) $\Rightarrow$ (2) is trivial, (2) $\Rightarrow$ (3) is an easy consequence of the definitions, and (3) $\Rightarrow$ (4) follows from Theorem 2.5.8 up to replacing embeddings with inclusions. As for (4) $\Rightarrow$ (1) observe that, if T has quantifier elimination, it follows easily from Lemma 2.5.3 that so does $T \cup \operatorname{diag}(A)$;[16] the embedding provided by Proposition 0.2.49 then allows us to invoke Theorem 2.7.1. $\square$

2.8 Cheating: eliminating quantifiers by definitional expansions

There is a reason if the title of this chapter has a "where" in it: namely, whether quantifier elimination holds or not, heavily depends on the language L in which we are working.

In order for everything below to go through smoothly, we need to allow 0-ary relation symbols. If R is 0-ary, then R is an atomic formula. In a fixed structure M, it can be interpreted as $\top$ or as $\bot$.

Definition 2.8.1. Let T be an L-theory. A *definitional expansion* of T is a theory T_Φ obtained as follows.

1. Fix a set of L-formulas Φ.
2. Let L_Φ be obtained by adding to L, for every $\varphi(x) \in \Phi$, an $|x|$-ary relation symbol $R_{\varphi(x)}$.
3. Let $T_\Phi \coloneqq T \cup \{\forall x\ (\varphi(x) \leftrightarrow R_{\varphi(x)}(x)) \mid \varphi(x) \in \Phi\}$.

The *Morleyisation* of T is the definitional expansion obtained by setting Φ to be the set of all L-formulas.

In other words, the definitional expansion given by Φ makes all formulas in Φ equivalent to an atomic formula.

Remark 2.8.2.

[16] More generally, if T has quantifier elimination and $T' \supseteq T$ is a theory in a language $L' \supseteq L$ where the only new symbols are constant symbols, it is easy to show that T' still eliminates quantifiers.

1. Φ is allowed to contain formulas with different free variables, e.g. a sentence φ and some $\psi(x)$ with $|x| = 6$.

2. One may similarly define a definitional expansion M_Φ of M, and, for fixed Φ, if $M \vDash T$ then $M_\Phi \vDash T_\Phi$.

3. The reason we allowed 0-ary relation symbols is to cover the case where $\varphi \in \Phi$ is a sentence; this wouldn't have been necessary if we only considered complete T, since in that case every sentence is, modulo T, equivalent to $\top$ or equivalent to $\bot$. Note that, in a language L without constant symbols, but with 0-ary relation symbols, there is more than one way to make $\emptyset$ into an L-structure: we need to decide which 0-ary relation symbols are true and which ones are false. Compare with point 3 of Remark 2.7.2, and think about what happens when you start with the empty theory in the empty language (it is not complete!), and then Morleyise.

The next remark is as trivial as it is important.

Remark 2.8.3. Let T_Φ be the Morleyisation of T.

1. T_Φ has quantifier elimination. In particular, all embeddings between models of T_Φ are elementary.

2. Every model of T expands uniquely to a model of T_Φ.

Hence, the models of T and those of its Morleyised, even if they are structures in different languages, are essentially the same structure, in the sense that they have the same definable sets; what changes is the embeddings between them.

Morleyisation is very useful when proving abstract model-theoretic facts, since it allows us to assume quantifier elimination for arbitrary structures, let me stress this again, *without changing the definable sets.*[17]

But then you may ask: why going through all the proofs in this chapter if we could just establish quantifier elimination by brute force? Because Morleyisation is just as useful in the abstract as it is useless in the concrete. Or, in other words, having quantifier elimination in a simple language allows us to understand the definable sets, while forcing quantifier elimination tells us nothing in this regard.

For instance, you may prove as an exercise that DLO eliminates quantifiers in $L = \{<\}$ (ok, this is a lousy exercise: basically, read the proof of Theorem 2.1.2 again, then invoke Theorem 2.6.2). By inspecting the quantifier-free formulas in $L = \{<\}$, we find that all subsets of $\mathbb{Q}^1$ are finite unions of intervals (possibly unbounded) and points (if you don't see it yet, use disjunctive normal form). By contrast, take $(\mathbb{N}, +, \cdot)$. Surely, if we Morleyise this structure, every definable set becomes quantifier-free definable. But understanding what is $R_{\varphi(x)}(\mathbb{N})$ is just as difficult as understanding what is $\varphi(\mathbb{N})$.

Even if Morleyising does not help us understanding the definable sets of a given structure, other definitional expansion may do. For example, an L-theory may not have quantifier elimination in L, but maybe we can eliminate quantifiers in a "reasonable" definitional expansion. E.g., if in T every formula is equivalent to a boolean combination of existential ones (that is, of the form $\exists x\ \varphi(x)$ with x quantifier-free), then we can take Φ to be the set of existential formulas and prove quantifier elimination "down to Φ", that is, for the definitional expansion induced by Φ.

[17]As opposed to, for those in the know, adding Skolem functions, for example.

Chapter 3

Some examples and a few applications

3.1 Algebraically closed fields

Definition 3.1.1. Let $L_{\text{ring}} \coloneqq \{+, 0, -, \cdot, 1\}$. The theory of *algebraically closed fields* ACF has axioms

1. axioms of fields
2. for every $n > 0$, the axiom

$$\forall y\ \exists x\ (x^n + y_{n-1}x^{n-1} + \ldots + y_1 x + y_0 = 0)$$

If p is a prime number, we define

$$\mathsf{ACF}_p \coloneqq \mathsf{ACF} \cup \{\underbrace{1 + 1 + \ldots + 1}_{p \text{ times}} = 0\}$$

Finally, we define

$$\mathsf{ACF}_0 \coloneqq \mathsf{ACF} \cup \{\underbrace{1 + 1 + \ldots + 1}_{n \text{ times}} \neq 0 \mid n > 0\}$$

As usual, one should check that these have a model. But, as you know from algebra,

Fact 3.1.2. Every field K has an *algebraic closure* K^{alg}: an algebraic extension which is algebraically closed.

By the way, you can prove this in a very model-theoretic fashion: first show (using algebra[1]) that there is an algebraic extension K_1 of $K_0 \coloneqq K$ where every polynomial over K_0 of positive degree has a root. Then iterate this and take the union of the chain you built.[2]

Clearly, ACF is not complete, since it does not decide whether $1 + 1 = 0$. If you know a little bit of field theory, you will recall that an algebraically closed

[1] No free lunches. I mean, at some point we should use that these are fields, no?

[2] In fact, one step suffices. See [Gil68].

field is determined up to isomorphism by its characteristic and its transcendence degree over its prime field.[3] If we take this for granted, we see immediately that ACF_0 and each ACF_p are complete: for every uncountable κ, they have a unique model of size κ, and we may apply Vaught's test (Exercise 0.4.17).

But even without taking this fact for granted, we can prove something stronger, namely quantifier elimination, even for the incomplete theory ACF. The characterisation of its completions will then follow easily.

If you want to do this by using as little algebra as possible, you can: you will need Lemma 2.4.6 and some elbow grease. But since we know a lot of things about the algebra of fields, we may as well exploit it to do a neat back-and-forth proof.

Theorem 3.1.3. ACF eliminates quantifiers.

Proof. We use Theorem 2.6.2. Given $M_0, N_0 \vDash \mathsf{ACF}$, by Löwenheim–Skolem there are uncountable $M \succeq M_0$ and $N \succeq N_0$. Since our assumptions are symmetrical, up to reversing the roles of M and N we only need to take care of the "forth" part.

Let $A \subseteq M$ and $B \subseteq N$ be finitely generated substructures, which in this language means finitely generated subrings, and $f_0 \colon A \to B$ an isomorphism.[4] If K, L are the fields they generate in M, N, we can easily (and uniquely) extend f_0 to an isomorphism $f \colon K \to L$.

If $a \in M \setminus K$ is transcendental over K, denote by $K[a]$ the ring generated by Ka and by $K[X]$ the ring of polynomials over K in one variable X. Since a is transcendental, $\mathrm{id}_K \cup \{a \mapsto X\}$ extends to an isomorphism $g_0 \colon K[a] \to K[X]$. Clearly, f extends to an isomorphism $g_1 \colon K[X] \to L[X]$ mapping X to X. Since L is finitely generated, its algebraic closure is countable, hence by choice of N there is $b \in N$ transcendental over L. By transcendence, the map $g_2 \colon L[X] \to L[b]$ sending a polynomial to its value in b is an isomorphism, hence $g_2 \circ g_1 \circ g_0$ is the required extension of f.

If $a \in M \setminus K$ is algebraic over K, let $g(X)$ be its minimal polynomial. Let $h(X)$ be its image under f. Since N is algebraically closed, $h(X)$ has a root b in N, and we conclude similarly as above, by using $K[X]/(g(X))$ instead of $K[X]$. □

This is an example where it was necessary to pass to elementary extensions: if M_0 contains an element a transcendental over the prime field F, and N_0 is F^{alg}, there is no partial isomorphism $M_0 \to N_0$ with a in its domain.

Corollary 3.1.4. The completions of ACF are ACF_0 and, for each p prime, ACF_p.

Proof. If $T \supseteq \mathsf{ACF}$ is complete, for each $n > 0$ it needs to decide whether

$$\underbrace{1 + 1 + \ldots + 1}_{n \text{ times}} = 0$$

[3]The subfield generated by 1, that is, either $\mathbb{Q}$ or $\mathbb{F}_p$. Up to isomorphism, of course. But maybe it's time to stop saying "up to isomorphism" every time.

[4]If M, N have different characteristic, then there is no partial isomorphism between them, since every finitely generated substructure needs to contain the interpretation of the constant 1, which is preserved by isomorphisms. So in this case there is no such f_0 and we are already done.

holds or not. Field theory tells us that this can hold for at most one n, and that such an n must be prime. This shows that each completion contains some ACF_p or ACF_0, so we only need to show that these are complete. But this follows from Theorem 2.7.1, since $\mathbb{Z}$ embeds in every field characteristic 0 and $\mathbb{F}_p$ in every field of characteristic p. $\square$

Corollary 3.1.5 (Chevalley–Tarski)**.** If $K \vDash \mathsf{ACF}$ and $X \subseteq K^{n+1}$ is *constructible*, that is, a Boolean combination of Zariski-closed sets, then its projection on the first n coordinates is still constructible.

Proof. This is essentially a restatement of quantifier elimination, after observing that "constructible" is the same as "quantifier-free definable". $\square$

Corollary 3.1.6 (Lefschetz principle)**.** Let φ be a sentence in L_{ring}. The following are equivalent:

1. $\mathbb{C} \vDash \varphi$.
2. $\mathsf{ACF}_0 \vdash \varphi$.
3. For cofinitely many primes p we have $\mathsf{ACF}_p \vdash \varphi$.
4. For infinitely many primes p we have $\mathsf{ACF}_p \vdash \varphi$.

Proof. (1) $\Leftrightarrow$ (2) holds because ACF_0 is complete. If $\mathsf{ACF}_0 \vdash \varphi$, then by compactness a finite subset of ACF_0 suffices to entail φ. This finite subset can only say that the characteristic is different from finitely many primes, so we get (2) $\Rightarrow$ (3). Since (3) $\Rightarrow$ (4) is trivial, we conclude by proving $\neg(2) \Rightarrow \neg(4)$. Again because ACF_0 is complete, if $\mathsf{ACF}_0 \nvdash \varphi$ then $\mathsf{ACF}_0 \vdash \neg\varphi$. By the previous implications, for cofinitely many primes p we have $\mathsf{ACF}_p \vdash \neg\varphi$, and by consistency $\mathsf{ACF}_p \nvdash \varphi$. $\square$

Combining this with a standard algebraic fact yields a proof of (one of the several forms of) the Nullstellensatz.

Corollary 3.1.7. Let $K \vDash \mathsf{ACF}$. If $\mathfrak{m} \subseteq K[X_0, \dots, X_{n-1}]$ is a maximal ideal, then there is $a \in K^n$ where all $f \in \mathfrak{m}$ are 0.

Proof. Clearly, all $f \in \mathfrak{m}$ have a zero in an extension of K, namely in the field $K[X_0, \dots, X_{n-1}]/\mathfrak{m}$, and a fortiori in $L \coloneqq (K[X]/\mathfrak{m})^{\mathrm{alg}}$. By Hilbert's Basis Theorem, $\mathfrak{m}$ is finitely generated, say $\mathfrak{m} = (f_0, \dots, f_k)$, hence a annihilates all $f \in \mathfrak{m}$ if and only if a annihilates all f_i. By construction, $L \vDash \exists x_0, \dots, x_{n-1} \bigwedge_{j \le k} f(x) = 0$. By quantifier elimination, the embedding $K \hookrightarrow L$ is elementary, hence $K \vDash \exists x_0, \dots, x_{n-1} \bigwedge_{j \le k} f(x) = 0$. $\square$

We conclude this section with a beautiful model-theoretic proof (in fact, the first one to be found) of an algebraic fact. First, an easy observation.

Exercise 3.1.8. Let $(I, <)$ be upward directed, $(M_i \mid i \in I)$ be a family of L-structures such that $i < j \implies M_i \subseteq M_j$, and $M \coloneqq \bigcup_{i \in I} M_i$. Let φ be a $\forall\exists$-sentence, that is, one of the form $\forall x\, \exists y\, \psi(x, y)$, with $\psi(x, y)$ quantifier-free. If, for every $i \in I$, we have $M_i \vDash \varphi$, then $M \vDash \varphi$.[5]

[5] We will not prove it here, but you should be aware that this has a converse: if the class of models of T models is closed under unions of chains (we don't even need to check arbitrary directed sets), then T is $\forall\exists$-axiomatisable. See [Mar02, Exercise 2.5.15] for a proof sketch.

Theorem 3.1.9 (Ax)**.** Let $f\colon \mathbb{C}^n \to \mathbb{C}^n$ be a polynomial function, that is, a function $(f_0, \dots, f_{n-1})$ where every f_i is a polynomial in (the same) n variables. If f is injective, then it is surjective.

Proof. By quantifying over coefficients as in Definition 3.1.1, it is easy to see that, for fixed n and $d \coloneqq \max_{i<n} \deg f_i$, the conclusion may be expressed by an L_{ring}-sentence. If you actually write down this sentence and put in prenex normal form, you will in all likelihood end up with a $\forall\exists$ sentence,[6] call it $\varphi_{n,d}$.

A moment's thought reveals that $\varphi_{n,d}$ holds over every finite field. Since $\mathbb{F}_p^{\mathrm{alg}}$ can be written as a directed union of finite fields, by Exercise 3.1.8 $\mathbb{F}_p^{\mathrm{alg}} \vDash \varphi_{n,d}$. Since ACF_p is complete, it follows that $\mathsf{ACF}_p \vDash \varphi_{n,d}$. This is true for every p, so we conclude by Corollary 3.1.6. □

Exercise 3.1.10. Consider $\mathbb{R}$ with its natural L_{ring}-structure. Prove that $\exists y\; x = y^2$ is not equivalent to a quantifier-free formula.

Exercise 3.1.11. Fix a field K. The *language of K-vector spaces* is $L_{K-\mathsf{VS}} \coloneqq \{+, 0, -\} \cup \{\lambda \cdot - \mid \lambda \in K\}$. Each K-vector space is made into an $L_{K-\mathsf{VS}}$-structure by interpreting $+, 0, -$ as the functions giving its underlying abelian group and $\lambda \cdot -$ as the 1-ary function "scalar multiplication by λ". Denote by $K{-}\mathsf{VS}$ the common theory of all *infinite* K-vector spaces. Prove that $K{-}\mathsf{VS}$ eliminates quantifiers and is complete.

If instead of a field K we take a ring R, and look at R-modules in a similar fashion, then quantifier elimination *tout court* can fail, but there is still quantifier elimination down to *positive primitive* formulas. There are many sources for this, e.g. [Poi00, Theorem 6.26], [TZ12, Theorem 3.3.5], or the extensive monograph [Pre88].

Remark 3.1.12. If K is a field and V a vector space, and we are interested in studying the model theory of V, we have two natural choices: viewing V as an $L_{K-\mathsf{VS}}$-structure, or throwing K inside the structure and looking at (K, V), see Example 0.3.1. The two resulting structures behave *very* differently, unless K is finite: if K is infinite, there are elementary extension of (K, V) where the field sort grows!

3.2 Some combinatorial structures

It is often useful to have at the ready an array of understandable structures and theories to test conjectures, understand new definitions, etc. Usually we like these structures to have quantifier elimination in a reasonable language, so that we understand, at least to some extent, their definable sets. This section contains some theories you can use for this purpose, and to get some practice with proofs of quantifier elimination (and of consistency!).

Definition 3.2.1. Recall that, by definition, $2^{<\omega} \coloneqq \{f\colon n \to \{0,1\} \mid n \in \omega\}$. Let $L = \{P_\sigma \mid \sigma \in 2^{<\omega}\}$, where each P_σ is a unary predicate. Let $T_{2^{<\omega}}$ be the theory with axioms[7]

---

[6] If you don't, and you tried writing it in good faith, please let me know what you wrote. My email address is at page vi.

[7] I am not aware of any standard name/notation for $T_{2^{<\omega}}$, I just put down the first that came to mind; suggestions are welcome.

1. $\forall x\ P_\emptyset(x)$, where we think of $\emptyset$ as the unique function $0 \to \{0,1\}$;
2. each P_σ is infinite;
3. whenever $\sigma_0 \subseteq \sigma_1$,[8] the axiom $\forall x\ P_{\sigma_1}(x) \to P_{\sigma_0(x)}$;
4. for all $\sigma \in 2^{<\omega}$, the axiom[9]

$$(\neg\exists x\ P_{\sigma^\frown 0}(x) \wedge P_{\sigma^\frown 1}(x)) \wedge \forall x\ (P_\sigma(x) \to (P_{\sigma^\frown 0}(x) \vee P_{\sigma^\frown 1}(x)))$$

Definition 3.2.2. Let κ be a nonzero cardinal (possibly finite) and let $L \coloneqq \{E_i \mid i < \kappa\}$, where each E_i is a binary relation symbol. The theory of κ *generic equivalence relations* is axiomatised by

1. every E_i is an equivalence relation with infinitely many classes;
2. for every finite nonempty $I \subseteq \kappa$, an axiom saying that, whenever X_i is an equivalence class of E_i, the intersection $\bigcap_{i\in I} X_i$ is infinite.

Definition 3.2.3. In $L = \{0,1,\cap,\cup,\subseteq,(\cdot)^\complement\}$, the theory of *atomless boolean algebras* is the theory of the Boolean algebras B such that $B \setminus \{0\}$ has no $\subseteq$-minimal elements.

We already said that DLO is complete and eliminates quantifiers. If you have already done Exercise 2.3.2, you will have probably realised that T_{rg} does too. And of course, so does the theory of infinite sets. Here are some useful variants:

1. The theory of a DLO together with a dense and codense unary predicate P.
2. For a fixed cardinal κ, the theory of κ-many DLO's $<_i$ on the same underlying set, where the intersection of finitely many intervals, each relative to a different $<_i$, is nonempty.
3. The theory of the *densely ordered random graph*: in the language $L = \{E, <\}$, take DLO together with a strengthened version of the axioms in Definition 2.3.1, stating not only the existence of one z with the required edges, but of a dense set of such z.

Exercise 3.2.4. Choose a (preferably nonempty) subset of the set of theories introduced in this section. Prove that the theories in this subset

1. are indeed theories, that is, they have a model,
2. eliminate quantifiers, and
3. are complete.

Exercise 3.2.5. Consider $T \coloneqq \mathrm{Th}(\mathbb{Z},<)$.

1. Prove that T does not eliminate quantifiers.
2. Find an expansion of $(\mathbb{Z},<)$ by one symbol only which has the same definable sets and quantifier elimination.

---

[8] That is, σ_1 extends σ_0.

[9] If $\operatorname{dom}\sigma = n = \{0,\dots,n-1\}$, we denote by $\sigma^\frown i$ the function with domain $n+1$ which restricts to σ and maps $n \mapsto i$.

Chapter 4

Realising many types

4.1 Types over parameters

Notation 4.1.1. Unless otherwise stated, T denotes a complete L-theory with infinite models, and M, N, M_0, etc. models of T.

We may still repeat that T is complete for emphasis.

We already saw what a type is in Definition 2.5.1. A type over a set of parameters is just what you expect:

Definition 4.1.2. Let $M \vDash T$ and $A \subseteq M$. A partial (respectively, complete) *n-type over A* is a partial (respectively, complete) n-type in $\mathrm{Th}(M_A)$.

Some easy but important observations:

Remark 4.1.3.

1. Let $\Phi(x)$ be a set of $L(M)$-formulas. Then $\Phi(x)$ is a partial type over M if and only if $\{\varphi(M) \mid \varphi(x) \in \Phi(x)\}$ has the *finite intersection property*, that is, every intersection of finitely many of its elements is nonempty.
2. Every partial type over A can be extended to a complete type over A, since every theory extends to a complete theory.

If you solved Exercise 0.4.4, you already know how to solve this:

Exercise 4.1.4. The class of models of a complete T with elementary embeddings has the *joint embedding property*: given any two models M_0, M_1 of T there are $N \vDash T$ and elementary embeddings $M_0 \to N$ and $M_1 \to N$.

If $N \succeq M$, then $\mathrm{Th}(N_A) = \mathrm{Th}(M_A)$. Therefore, the types over $A \subseteq M$ do not change when passing to an elementary extension of M. For this reason, we may drop the M in "tp^M":

Definition 4.1.5. Let $A \subseteq M \vDash T$ and $b \in M^n$. The *type of b over A* is

$$\mathrm{tp}(b/A) \coloneqq \{\varphi(x_0, \dots, x_{n-1}) \in L(A) \mid M \vDash \varphi(b_0, \dots, b_{n-1})\}$$

If $p(x)$ is a type over A, we say that b *realises* $p(x)$ iff $p(x) = \mathrm{tp}(b/A)$; in this case, we write $b \vDash p$.

4.2 Type spaces

Definition 4.2.1. Let $A \subseteq M \vDash T$, and fix a tuple of variables x. The space $S_x(A)$ is the set of $|x|$-types $p(x)$ over A, equipped with the topology generated by the basis of open sets $\{[\varphi(x)] \mid \varphi(x) \in L(A)\}$, where

$$[\varphi(x)] := \{p(x) \in S_x(A) \mid p(x) \vdash \varphi(x)\}$$

Remark 4.2.2.

1. This is indeed a basis for a topology, as opposed to just a prebasis. In fact, it is even closed under finite intersections, since $[\varphi(x)] \cap [\psi(x)] = [\varphi(x) \wedge \psi(x)]$. Similarly, it is closed under finite unions, since $[\varphi(x)] \cup [\psi(x)] = [\varphi(x) \vee \psi(x)]$. By definition of basis, open sets are those of the form $\bigvee_{i \in I}[\varphi_i(x)]$.

2. $S_x(A)$ is Hausdorff, since if $p(x) \neq q(x)$ there must be $\varphi(x) \in p(x)$ such that $\varphi(x) \notin q(x)$. Because $q(x)$ is complete, then $\neg\varphi(x) \in q(x)$. Therefore $p(x) \in [\varphi(x)]$, $q(x) \in [\neg\varphi(x)]$, and since types are consistent, we clearly have $[\varphi(x)] \cap [\neg\varphi(x)] = \emptyset$.

3. Each $[\varphi(x)]$ is clopen, since it has complement $[\neg\varphi(x)]$.

4. It follows from the previous points that the $[\varphi(x)]$ also form a basis for the closed sets.

5. Nonempty closed sets correspond to partial types. More precisely, every nonempty closed set is of the form $F = \bigcap_{\varphi(x) \in \Phi(x)}[\varphi(x)]$, for $\Phi(x)$ a partial type over A. In other words, $p(x) \in F$ if and only if $p(x)$ is a completion of $\Phi(x)$. We denote F by $[\Phi(x)]$.

6. $S_x(A)$ is compact, because of... compactness. To see this, recall that an equivalent definition of (topological) compactness is "every family of closed sets with the finite intersection property has nonempty intersection". Spelling this out, if we think of types as complete $L \cup \{c\}$-theories, this means precisely that if every finite subset of $\Phi(c)$ is consistent, then $\Phi(c)$ is consistent.

7. Restricting a type $p(x, y)$ to the formulas which do not involve y yields a continuous, surjective map $S_{xy}(A) \to S_x(A)$. Similarly, if $A \subseteq B \subseteq M$, then the restriction map $p \mapsto p \restriction A \colon S_x(B) \to S_x(A)$ is surjective and continuous.

8. $S_{xy}(A)$ is *not* the product $S_x(A) \times S_y(A)$. In other words, even if $p(x)$ and $q(y)$ are complete, $p(x) \cup q(y)$ need not be. This depends on the fact that not every formula $\varphi(x, y)$ can be written as a boolean combination of formulas of the form $\psi(x)$ or $\theta(y)$. An easy example is the formula $x = y$. If for example $A = M$ is a model, and $p(x)$ is a *nonrealised* type, that is, a type extending $\{x \neq a \mid a \in M^{|x|}\}$, then $p(x) \cup p(y)$ has one completion containing[1] $x = y$ and (at least) one completion containing $x \neq y$. If you are familiar with the Zariski topology, this is a akin to the fact that the Zariski topology on $\mathbb{A}^2$ is not the product of the Zariski topology on $\mathbb{A}^1$ with itself.

[1] If $|x| > 1$ this is an abbreviation for $\bigwedge_{i<|x|} x_i = y_i$.

9. If $A = M$ is a model, then the set of *realised* types, that is, those containing a formula $x = a$ for some $a \in M$, is dense. In fact, let $[\varphi(x)]$ be a basic open set. If $[\varphi(x)]$ is nonempty, it contains some $p(x)$. This $p(x)$ is realised in some $N \succeq M$, say by b. In particular, $N \vDash \varphi(b)$, so $N \vDash \exists x\ \varphi(x)$, hence $M \vDash \exists x\ \varphi(x)$. If $a \in M$ is such that $M \vDash \varphi(a)$, then $\{x = a\}$ implies a complete type, which is clearly realised, and clearly contained in $[\varphi(x)]$.

Clearly, up to homeomorphism, $S_x(A)$ only depends on $|x|$, and not on the specific tuple of variable used. Therefore, if we do not care about the particular variables used, or if they are clear from context, we also use the following notation.

Notation 4.2.3. We write $S_n(A)$ to denote the topological space $S_x(A)$ for some x with $|x| = n$. We denote by $S_{<\omega}(A)$, or simply by $S(A)$, the disjoint union of the $S_n(A)$ for $n \in \omega$.

Exercise 4.2.4. All clopen subsets of $S_x(A)$ are of the form $[\varphi(x)]$, for some $\varphi(x) \in L(A)$.

A crucial idea behind modern model theory (if not *the* idea which kickstarted modern model theory) is that certain topological properties of the spaces $S_x(A)$ are intimately connected to the behaviour of T and of its models. We will see some of this later. As a warm up, try to answer the following question.

Question 4.2.5. What does it mean for $\{p(x)\}$ to be an isolated point of $S_x(A)$

1. for an arbitrary A?
2. in the special case where $A = M$ is a model?

4.3 Examples

Before we develop the theory further, it is time to familiarise ourselves with type spaces by looking at a bunch of examples.

4.3.1 Infinite sets

Let T be the theory of infinite sets, for which by now you should be able to prove quantifier elimination in at least two different ways. Fix $A \subseteq M \vDash T$, and let us look at $S_1(A)$.

First, let us look at the case $A = \emptyset$. A direct inspection of the possible quantifier-free formulas with no parameters and with only one free variable should convince you that $S_1(\emptyset)$ has only one element, implied by the formula $x = x$. But there is a more elegant way of proving this, which has the advantage of working just as quickly also in some cases where the language is a bit more complicated.

Exercise 4.3.1.
1. If there is $f \in \mathrm{Aut}(M/A)$ such that $f(a) = b$, then $\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$.
2. More generally, if there are $N \succeq M$ and $f \in \mathrm{Aut}(N/A)$ such that $f(a) = b$, then $\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$.

3. Find T, M, A with $A \subseteq M \vDash T$, some $f \in \mathrm{Aut}(M)$, and $a, b \in M$ such that
   (a) f fixes A *setwise*,
   (b) $b = f(a)$, and
   (c) $\mathrm{tp}(a/A) \neq \mathrm{tp}(b/A)$.

Take now your favourite $M \vDash T$, that is, your favourite infinite set. An $f \in \mathrm{Aut}(M)$ is nothing more than a permutation of M, that is, a bijection $M \to M$, and it follows from the previous exercise that there is only one type over $\emptyset$ realised in M. Since M was arbitrary, by Exercise 4.1.4 and point 2 of Exercise 4.3.1 there is only one type over $\emptyset$, full stop.

What about $S_1(A)$ for nonempty A? Again by inspection, or again by using automorphisms, we see that

1. for every $a \in A$, there is a 1-type $p_a(x)$ implied by the formula $x = a$. In particular, each $\{p_a\}$ is open, that is, isolated, in $S_1(A)$;

2. there is one, and only one, element not of the form above, the *generic type* $p_{\mathrm{g}}(x)$, axiomatised by $\{x \neq a \mid a \in A\}$; if A is infinite, it is a (well, the only) nonisolated point.

You may have thought that these very disconnected compact spaces are either trivial (e.g. $S_1(\emptyset)$ above, which has just one point), or very difficult to visualise. This is not true. For example, if $|A| = \aleph_0$, the description above may be very quickly turned into an homeomorphism between $S_1(A)$ and $\{0\} \cup \{1/n \mid n \in \omega \setminus \{0\}\}$ (with the usual subspace topology inherited from $\mathbb{R}$), sending p_{g} to 0.

While we are here, observe the following.

Remark 4.3.2. In *every* theory, if $S_1(A)$ is infinite, then it *must* have at least one nonisolated point: otherwise it would be an infinite discrete space, so it wouldn't be compact.

This kind of considerations will play a crucial role in the next chapter. But now, let's go back to examples.

What about $S_n(A)$, for $n \geq 2$? Clearly, an n-type $p(x)$ will, to begin with, determine n 1-types $p \upharpoonright x_i$, obtained by considering only the formulas with no free variable other than x_i. If all the $p \upharpoonright x_i$ are realised in A, then this already determines $p(x)$. But if, for example, $p \upharpoonright x_0 = p_{\mathrm{g}}(x_0)$ and $p \upharpoonright x_1 = p_{\mathrm{g}}(x_1)$, then there are at least two completions of $(p \upharpoonright x_0) \cup (p \upharpoonright x_1)$, since we need to decide whether $x_0 = x_1$ or $x_0 \neq x_1$ will be in our completion.

Long story short, an element $p(x) \in S_n(A)$ is determined by:

1. which x_i are equal to some point $a \in A$ (and, for these, which $a \in A$ they equal), and

2. for the other x_i, for which pairs (i, j) we have $p(x) \vdash x_i = x_j$, and for which instead $p(x) \vdash x_i \neq x_j$.

Clearly, every type needs to specify the information above. But why is that enough to entail a complete type? You can prove this in two ways:

1. use quantifier elimination and show that every $L(A)$ formula $\varphi(x)$ is decided by fixing the information above; or

2. show that if $a, b \in M^{|x|}$ agree on the conditions mentioned above, then there are $N \succeq M$ and $f \in \mathrm{Aut}(N/A)$ such that $f(a) = b$ (where $f(a) = (f(a_0), \dots, f(a_{|a|-1}))$.

Do we really need to pass to an elementary extension in order to use these automorphism arguments? Well, in general yes: for example if $A = M$ then p_{g} is not realised in M. "Ok —you may say— but you took the whole of M, what if A is *small*? For example, what if all types over A are realised in M?" Keep reading.

4.3.2 DLO

Let $T = \mathsf{DLO}$. Let us look at spaces of 1-types $S_1(A)$. By quantifier elimination, we may equivalently look at quantifier-free types. If $p(x) \in S_1(A)$, since $p(x)$ is complete, if $\varphi(x) \in \psi(x) \in L(A)$ and $p(x) \vdash \varphi(x) \lor \psi(x)$ then we must have $p(x) \vdash \varphi(x)$ or $p(x) \vdash \psi(x)$. By disjunctive normal form and the axioms of DLO, it follows that $p(x)$ is determined by which formulas of the form $x > a$, $a > x$, $x = a$, $x \neq a$ it contains.

The types containing $x = a$ are, by definition, the realised ones. Every other type determines (and is uniquely determined by) a cut in A, that is, a pair (L, R) with $A = L \sqcup R$ with $L < R$, including the degenerate cases where L or R are empty. In detail, each such cut $C = (L, R)$ determines a nonrealised 1-type by setting $p_C(x) \coloneqq \{x > a \mid a \in L\} \cup \{x < a \mid a \in R\}$. Conversely, each nonrealised type $p(x)$ determines a cut $C_p = (L_p, R_p)$, where $L_p = \{a \in A \mid p(x) \vdash x > a\}$ and $R_p = \{a \in A \mid p(x) \vdash x < a\}$. These maps are clearly inverses of each other.

Let us look at three very concrete cases.

Example 4.3.3. $A = \emptyset$. There is a unique 1-type $p(x)$, implied by the formula $x = x$.

Example 4.3.4. $A = \mathbb{Q}$. We have different kinds of types:

1. For each $a \in \mathbb{Q}$, a realised type $p_a(x)$, implied by $x = a$.

2. The type $p_{+\infty}(x) \coloneqq \{x > a \mid a \in \mathbb{Q}\}$, corresponding to the cut with $R = \emptyset$, and the type $p_{-\infty}(x)$, corresponding to $L = \emptyset$.

3. For each $a \in \mathbb{Q}$, a type $p_{a^+}(x) \coloneqq \{x > a\} \cup \{x < b \mid b > a\}$, corresponding to the cut with $R = (a, +\infty)$, and a type $p_{a^-}(x)$, corresponding to the cut with $R = [a, +\infty)$.

4. Types corresponding to *irrational cuts*, that is, cuts (L, R) where L has no maximum, R has no minimum, and both are nonempty. If you prefer, these are precisely the cuts of the form $\{x > q \mid q < r\} \cup \{x < q \mid q > r\}$, for $r \in \mathbb{R} \setminus \mathbb{Q}$.

Example 4.3.5. If $A = \mathbb{R}$, and $C = (L, R)$ is a cut with L and R both nonempty, then completeness of $\mathbb{R}$ tells us that L must have a supremum r, which will either be in L or in R. It follows that over $\mathbb{R}$ there are no irrational cuts. All other kinds of 1-types described above are clearly still possible, and there are no other kinds of 1-types.

What about the topological structure? You may imagine $S_1(A)$ as some sort of very disconnected completion of A: open sets are generated by those of the form $[x = a]$, $[x > a]$, and $[x < a]$.[2]

And what about n-types for $n > 1$? Again by quantifier elimination and inspection of the quantifier-free L-formulas, we see that an n-type $p(x)$ is determined by its 1-subtypes $p \upharpoonright x_i$, together with its restriction $p(x) \upharpoonright \emptyset$. In other words an n-type $p(x)$ over A is determined by

1. in which cut of A it places each x_i (including the degenerate case where $p(x) \vdash x_i = a$ for some $a \in A$), and
2. in which order it puts its variables, including the case where some of them are identified; in other words, which formulas of the form $x_i < x_j$ or $x_i = x_j$ it implies.

4.3.3 A digression: binarity

Careful: the trick we used for infinite sets and DLO, namely, reducing an n-type over A to n 1-types over A and one n-type over $\emptyset$ does not always work. But what is it that we used exactly?

Exercise 4.3.6. For a complete T, the following are equivalent.

1. Every formula $\varphi(x)$ is equivalent to a boolean combination of formulas with at most two free variables.[3]
2. For all tuples a, b and all sets A, we have $\mathrm{tp}(a/A) \cup \mathrm{tp}(b/A) \cup \mathrm{tp}(ab/\emptyset) \vdash \mathrm{tp}(ab/A)$.
3. For all tuples $a^0, a^1, \ldots, a^k$ and all sets A, we have $\mathrm{tp}(a^0/A) \cup \mathrm{tp}(a^1/A) \cup \ldots \cup \mathrm{tp}(a^k/A) \cup \mathrm{tp}(a^0, \ldots, a^k/\emptyset) \vdash \mathrm{tp}(a^0, \ldots, a^k/A)$.

Exercise 4.3.7. DLO is binary. More generally, any complete theory which eliminates quantifiers in a language L where

1. there are no function symbols, and
2. every relation symbol has arity at most 2

is binary.

For the sake of simplicity, most (but not all!) examples in this section will be binary.

4.3.4 The random graph

Since the random graph eliminates quantifiers in a binary relational language, it is binary.[4] By quantifier elimination, n-types over $\emptyset$ are easily described: a type $p(x)$ over $\emptyset$ needs to say which x_i coincide and, for the pairs with $p(x) \vdash x_i \neq x_j$, whether $p(x) \vdash E(x_i, x_j)$ or $p(x) \vdash \neg E(x_i, x_j)$.

So we are left to describe $S_1(A)$ for arbitrary A. Clearly, a 1-type $p(x)$ will need to decide

[2] Careful, a here is an element of A, not a type!

[3] Not necessarily always the same two variables, e.g. $\varphi(x_0, x_1) \wedge \varphi(x_1, x_2)$ is fine.

[4] If you prefer, you can use quantifier elimination directly.

1. whether $x = a$ for some $a \in A$, and
2. if this is not the case, for which $a \in A$ we have $E(x, a)$, and for which $a \in A$ we have $\neg E(x, a)$.

It follows from quantifier elimination that providing this information determines a complete 1-type. But does *every* choice give a 1-type, or are there some inconsistent ones? The Random Graph axioms and compactness tell us that any choice will do:

Exercise 4.3.8. For every $B \subseteq A$ there is $p \in S_1(A)$ such that

1. $p(x) \vdash \{x \neq a \mid a \in A\}$,
2. $p(x) \vdash \{E(x, a) \mid a \in B\}$, and
3. $p(x) \vdash \{\neg E(x, a) \mid a \in A \setminus B\}$.

Exercise 4.3.9. Consider the subspace $X \subseteq S_1(A)$ of nonrealised types, that is, the closed subspace given by the partial type $\{x \neq a \mid a \in A\}$. Prove that X is homeomorphic to $2^{|A|}$, that is, the product of $|A|$-many copies of the discrete space $\{0, 1\}$, with the product topology.

If $|A| = \aleph_0$, you may have recognised that the space X above is nothing more that Cantor space, that is, the Cantor set with the subspace topology inherited from $\mathbb{R}$. If you want a full type space homeomorphic to the Cantor space, without needing to pass to subspaces, here is an example.

Exercise 4.3.10. Prove that, if $T_{2^{<\omega}}$ is as in Definition 3.2.1, then $S_1(\emptyset)$ is homeomorphic to Cantor space.

4.3.5 Generic equivalence relation

Let T be the theory of a generic equivalence relation E. This is the case $\kappa = 1$ of Definition 3.2.2, except that here we just write write E instead of E_0.

Elements of $S_1(A)$ can be of three kinds:

1. Realised. You know the drill. Isolated, etc etc.
2. For each $a \in A$ there is a "generic type of the class of a", axiomatised by $\{x \neq a \mid a \in A\} \cup \{E(x, a)\}$, that is, the type of a new point in the class of A. If $\{b \in A \mid E(b, a)\}$ is infinite, then this point is not isolated.
3. A single "generic" type, axiomatised by $\{\neg E(x, a) \mid a \in A\}$, that is, the type of a point in a new equivalence class. Similarly, if A/E is infinite, then this point is not isolated.

Spoiler 4.3.11. You may object: "ok, but $\{x \neq a \mid a \in A\} \cup \{E(x, a)\}$ is only nonisolated because of $\{x \neq a \mid a \in A\}$; that is, in the subspace of nonrealised types, this type is isolated by $[E(x, a)]$." Congratulations, you are halfway through the road to the definition of Morley rank. Keep reading.[5]

While we are on the subject of spoilers: soon we will be interested in cardinalities of type spaces. The following exercise is recommended.

[5] Or just go straight away to Definition 7.1.1.

Exercise 4.3.12.

1. Assuming that A is infinite, compute the cardinality of $S_1(A)$.

2. Do the same for the theory of κ generic equivalence relations, for every nonzero cardinal κ.

Note that the number of equivalence relations here is fixed by the language. But, as in the case of vector spaces (cf. Remark 3.1.12), we can also make a different choice: what if we allow the equivalence relations to be part of the model?

Definition 4.3.13. Let $L = \{P, R, E\}$, where P, R are unary predicates ("Points" and "Relations") and E is a ternary relation symbol. The theory T^*_{feq} has the following axioms.

1. $\forall x, y, z\ \big(E(x, y, z) \to (P(x) \land P(y) \land R(z))\big)$.[6]

2. The predicate R is infinite.

3. Every $E(-, -, z)$ (for fixed z) is an equivalence relation on P with infinitely many classes.

4. For every $n \in \omega$, an axiom saying that pairwise distinct equivalence relations $R(-, -, z_0), \dots, R(-, -, z_n)$ interact generically (as in Definition 3.2.2).

Exercise 4.3.14. 1. Prove that T^*_{feq} is indeed a theory.

2. Prove that T^*_{feq} is complete and has quantifier elimination.

3. Prove that T^*_{feq} is not binary.

4. Count how many 1-types there are over a model M.[7]

4.3.6 Algebraically closed fields

Let p be either a prime or 0, and let $T = \mathsf{ACF}_p$. Fix $M \vDash T$, and let us look at $S_1(M)$. By quantifier elimination, a type $p(x)$ is determined by which polynomials with coefficients in M are 0 in x, and which are not. Since we are working over a model M, this means that we have two possibilities:

1. For some $f(X) \in M[X]$ of positive degree, we have $p(x) \vdash f(x) = 0$. Since M is algebraically closed, there are $a_0, \dots, a_{\deg f - 1} \in M$ such that $p(x) \vdash \bigvee_{i < \deg f} x = a_i$. Since $p(x)$ is complete, it needs to choose one of these disjuncts, hence for some $i < \deg f$ we have $p(x) \vdash x = a_i$.

2. The only remaining option is the generic type $p(x) = \{f(x) \neq 0 \mid f(X) \in M[X], \deg f > 0\}$. Again, by compactness this type cannot be isolated.

[6] If you want to use multi-sorted structures, this is a good place to view P, R as sorts instead of predicates, and E as a relation of arity $P^2 \times R$.

[7] Hint: you can do this even without obtaining a complete description of all types.

Note anything strange? $S_1(M)$ is essentially the same as the space of 1-types over an infinite set M with no structure. What about $S_1(A)$, for A not a model? In this case, it is not true that every isolated type is realised in A. For example, if $p = 0$ and $A = \mathbb{Q}$, then $x \cdot x = 2$ implies a complete type, but of course no $a \in \mathbb{Q}$ realises it. More generally, if $f(X) \in \mathbb{Q}[X]$ is irreducible, then $f(X) = 0$ will imply a complete type.

What is $S_n(A)$ in general? ACF_p is *not* binary,[8] so we cannot resort to the same trick we used over and over in this section, and we need a bit of algebra. It is easy to see that $S_n(A)$ is essentially the same as $S_n(\langle A\rangle)$, where $\langle A\rangle$ is the structure generated by A, which in this language means the ring generated by A. By quantifier elimination we only need to deal with formulas of the form $f(X) = 0$, and by clearing denominators we see that we may pass to the fraction field of $\langle A\rangle$. Long story short, we only need to look at $S_n(K)$ for K a field, not necessarily algebraically closed.

So we need a convenient way to describe a consistent, complete choice of formulas of the form $f(X) = 0$ and $f(X) \neq 0$, where $f \in K[X]$ and $X = (X_0, \dots, X_{n-1})$. Of course, we already know the answer from algebra: the types $p(x) \in S_n(K)$ are in bijection with the prime ideals of $K[X]$. The "prime" here depends on completeness of $p(x)$: by completeness, there is some $N \succ M$ and some $a \in N^n$ such that $p(x) = \operatorname{tp}(a/K)$. If $f(a) \cdot g(a) = 0$; then clearly $f(a) = 0$ or $g(a) = 0$. I will leave the details as an exercise.

Exercise 4.3.15. The map $p(x) \mapsto \{f(X) \in K[X] \mid p(x) \vdash f(x) = 0\}$ is a bijection between $S_n(K)$ and the prime ideals of $K[X]$, where $X = (X_0, \dots, X_{n-1})$.

So $S_n(K)$, as a set, is essentially the same as $\operatorname{Spec}(K[X_0, \dots, X_{n-1}])$. If you have never seen this notation before, you may safely skip to the next section.

On the other hand, if you are a bit familiar with algebraic geometry, you may ask yourself whether, if $\operatorname{Spec}(K[X_0, \dots, X_{n-1}])$ is equipped with the Zariski topology, then this bijection is a homeomorphism. The answer is no: the topology induced by the bijection above coincides with the *constructible* one: each $[f(X) = 0]$ is clopen. In fact, it *is* possible to view $\operatorname{Spec}(K[X_0, \dots, X_{n-1}])$ with the Zariski topology as a type space, but this requires the notion of "type space" to be generalised: in particular, we need to allows for non-Hausdorff spaces. See [DST19, Section 14]. In fact, by changing the logic, one may view every spectral space as a type space. See for example [Hay19, Kam22].

4.4 Saturation

Proposition 2.5.6 tells us that all types over A are realised by some element in some model of $\operatorname{Th}(M_A)$. By Exercise 4.1.4 (or, if you prefer by proving this fact directly), we find that

Remark 4.4.1. Every type over $A \subseteq M$ is realised in some elementary extension of M.

In general, passing to an elementary extension is necessary:

[8] Here is a hint to prove it: let $A = (\mathbb{Q}(c))^{\mathrm{alg}}$, with c transcendental over $\mathbb{Q}$. Take $a, b \in \mathbb{C}$ algebraically independent over $\mathbb{Q}^{\mathrm{alg}}$, but such that $a - b = c$. Use one of the equivalent forms of binarity from Exercise 4.3.6.

Example 4.4.2. The partial type[9] $\pi(x) := \{x \neq m \mid m \in M\}$ is not realised in M.

Therefore, we cannot hope for all partial types over M to be realised in M. In fact, the example above shows that this is never true, unless M is finite.

Definition 4.4.3. Let κ be an infinite cardinal. We say that M is *κ-saturated* iff, whenever $A \subseteq M$ is such that $|A| < \kappa$, and $n \in \omega$, then every n-type over A is realised in M.

Notation 4.4.4. Even if κ-saturation depends only on the cardinality of κ, and not on its order type, it is common to say ω-saturated instead of $\aleph_0$-saturated. Things like "ω_1-saturated" instead of "$\aleph_1$-saturated" also appear in the literature.

Saturation can be checked on 1-types:

Proposition 4.4.5. Suppose that, whenever $A \subseteq M$ is such that $|A| < \kappa$, then every 1-type over A is realised in M. Then M is κ-saturated.

Proof. Let $p(x) \in S_{n+1}(A)$, where $A \subseteq M$ and $|A| < \kappa$. Let $q(x_0, \dots, x_{n-1}) := \{\varphi(x_0, \dots, x_{n-1}) \mid p(x) \vdash \varphi(x_0, \dots, x_{n-1})\}$ be its restriction to the first n coordinates. Inductively, there is $a \in M^n$ such that $a \vDash q$. By substituting a_i for x_i inside $p(x)$, we find a 1-type $r(x_n)$ over Aa,[10] and by assumption there is $a_n \in M$ realising $r(x_n)$. Clearly, $(a_0, \dots, a_n) \vDash p(x)$. □

Exercise 4.4.6.

1. Prove that $(\mathbb{R}, <)$ is ω-saturated, but not $\aleph_1$-saturated.
2. Characterise the ω-saturated $M \vDash \mathsf{ACF}$.
3. Which of these theories have a countable ω-saturated model? For which of these theories *all* countable models are ω-saturated?
   (a) The theory of infinite sets.
   (b) DLO.
   (c) The theory of the Random Graph.
   (d) $K{-}\mathsf{VS}$, for K a field.[11].
   (e) The theory $T_{2^{<\omega}}$ from Definition 3.2.1.
   (f) The theory of 1 generic equivalence relation.
   (g) T^*_{feq}.

Another addition to the list of trivial but important things:

Remark 4.4.7. Since types over A are in particular sets of $L(A)$-formulas, for every $n \in \omega$ there are at most $2^{|L|+|A|}$-many n-types over A.

[9] Recall that we are assuming that T has infinite models; since it is also complete, it has no finite models, so $\pi(x)$ is consistent.
[10] You may want to check as an exercise that $r(x_n)$ is indeed a (complete) type over Aa.
[11] Semi-hint: the answer may depend on K.

Lemma 4.4.8. Let $\kappa \geq |L|$ and $|M| \leq 2^\kappa$. Then there is $N \succeq M$ with $|N| \leq 2^\kappa$ and such that, for every $A \subseteq M$ with $|A| \leq \kappa$, the model N contains realisations of all types over A.

Proof. For every $n \in \omega$ and every n-type $p(x)$ over some $A \subseteq M$ of size $|A| \leq \kappa$, add to $L(M)$ a tuple of constants $c_p = (c_{p,0}, \dots, c_{p,|x|-1})$; call the resulting language L'. Since 2^κ has cofinality larger than κ, there are at most 2^κ subsets of M of size κ. Together with the previous remark, this yields $|L'| \leq 2^\kappa$, and by compactness there is $N' \vDash \mathrm{ED}(M) \cup \{p(c_p) \mid A \subseteq M, |A| \leq \kappa, p(x) \in S(A)\}$.[12] Let C be the set of interpretations in N' of all the new constants introduced above. Then $|C \cup M| \leq 2^\kappa$, and by applying downward Löwenheim–Skolem and taking the reduct to L we obtain the desired N. □

In the last step of last proof, we can justify that $N \succeq M$ in two ways: one is observing, before taking the reduct to L, that the L'-structure we built satisfies $\mathrm{ED}(M)$. Alternatively, we can use the following:

Exercise 4.4.9. Suppose $M_0 \subseteq M_1 \subseteq M_2$.

1. Suppose that $M_1 \preceq M_2$. Then $M_0 \preceq M_2$ if and only if $M_0 \preceq M_1$.
2. Find an example where $M_0 \preceq M_2$, $M_0 \preceq M_1$, but $M_1 \not\preceq M_2$.

The N we built in the previous lemma need not be κ-saturated, for the simple reason that we introduced new parameters. We fix this by repeating the construction transfinitely many times.

Theorem 4.4.10. Let $\kappa \geq |L|$ and $|M| \leq 2^\kappa$. Then there is a κ^+-saturated $N \succeq M$ of size at most 2^κ.

Proof. We do an inductive construction of length κ^+. We start with $M_0 := M$. At successor stages, we use Lemma 4.4.8 to take as $M_{\alpha+1}$ some elementary extension of M_α of size at most 2^κ and realising all types over subsets of M_α of size at most κ. At limit stages λ, we take $M_\lambda := \bigcup_{\alpha<\lambda} M_\alpha$, and observe that this is an elementary extension of each previous M_α by Exercise 1.4.4. Since λ is an ordinal of cardinality at most κ, we have $|M_\lambda| \leq \kappa \cdot 2^\kappa = 2^\kappa$, so we may continue the construction. Keep doing this for every ordinal below κ^+, and at the end set $N := \bigcup_{\alpha<\kappa^+} M_\alpha$; observe immediately that $|N| \leq \kappa^+ \cdot 2^\kappa = 2^\kappa$.

In order to check that N is κ^+-saturated, take $A \subseteq N$ of size $|A| < \kappa^+$, and let $p(x) \in S(A)$. Since κ^+ is regular, there must be $\alpha < \kappa^+$ such that $A \subseteq M_\alpha$. By construction, $M_{\alpha+1}$ contains the required b. □

Even if it does not really simplify the proof, note that we could have just worked with 1-types and obtained the same result by Proposition 4.4.5.

Exercise 4.4.11. Suppose that $|L| = \aleph_0$ and, for every n, there are at most $\aleph_0$ types over $\emptyset$. Then T has a countable ω-saturated model.[13]

The converse is trivial: a countable model can only realise countably many types over $\emptyset$, and if it is ω-saturated then it realises all types over $\emptyset$.

Here is the promised converse to Theorem 2.6.2. Alas, you will have to supply the proof yourself.[14]

[12] As usual, check that this is consistent as an exercise.

[13] Hint: prove first that, over any finite A, there are at most $\aleph_0$ types.

[14] Hint: back-and-forth is about realising quantifier-free types, no?

Exercise 4.4.12. Let T be a possibly incomplete theory with quantifier elimination. If M, N are ω-saturated models of T, then the family of all partial isomorphisms between finitely generated substructures of M and N has the back-and-forth property.

4.5 Properties of saturated models

In this section we look at consequences of saturation.

In the proof of Theorem 4.5.2 below, we will deal with types of infinite tuples or, if you prefer, types in infinitely many variables. While essentially everything translates, keep in mind that the definition of saturation only talks of finitary types.[15] Recall the notion of *elementary map* from Exercise 2.7.3.

Remark 4.5.1. Elementary maps preserve types: if $f : M \to N$ is elementary, then $\mathrm{tp}(a) = \mathrm{tp}(f(a))$.

Theorem 4.5.2. Every κ-saturated N is κ-*universal*: for every M with[16] $|M| \le \kappa$ there is an elementary embedding $M \to N$.

Proof. We do a "only forth" proof, not just in the usual sense that we say "the 'back' is analogous", but in the sense that we only need (and only have enough hypothesis to prove) the "forth". Fix an enumeration $(a_i)_{i<\kappa}$ of M of order type κ. Inductively, we define a partial elementary map $f_\alpha : a_{<\alpha} \to N$; notationally, write $a_i \mapsto b_i$. Because T is complete, the (unique) partial map with domain $\emptyset$ is elementary. At limit stages, we take unions; since f_α is elementary if and only if each of its restrictions to a finite domain is, elementarity is preserved in unions of chains. At the end, we take $f := \bigcup_{\alpha<\kappa} f_\alpha$, which will be an elementary map with domain the whole of M, that is, an elementary embedding.

So we are left to deal with the inductive definition of $f_{\alpha+1}$. By inductive hypothesis, $a_{<\alpha}$ and $b_{<\alpha}$ have the same type. Consider $p(x) := \mathrm{tp}(a_\alpha/a_{<\alpha})$, and let $q(x)$ be obtained by $p(x)$ by replacing, for $i < \alpha$, each a_i with b_i. If $q(x)$ is consistent, by saturation of N and the fact that $q(x)$ is over fewer than κ parameters we can find $b_\alpha \vDash q(x)$ in N, and we are done. So suppose that $q(x)$ is inconsistent. Hence, for some $\varphi(x, w) \in L$ such that[17] $p(x) \vdash \varphi(x, a_{<\alpha})$, we have $\mathrm{Th}(N_{b_{<\alpha}}) \vdash \neg\exists x\ \varphi(x, b_{<\alpha})$.[18] Since $p(x)$ is a type, on the other hand $\mathrm{Th}(M_{a_{<\alpha}}) \vdash \exists x\ \varphi(x, a_{<\alpha})$. This contradicts that $f_{<\alpha}$ is elementary (or, if you prefer, that $a_{<\alpha}$ and $b_{<\alpha}$ have the same type). $\square$

Note that in order for this to go through, we really need to work with types, as opposed to quantifier free types. Otherwise, there is no guarantee that $q(x)$ will be consistent. Of course, if T has quantifier elimination (which, at this level of generality, we may assume by Morleyising), then the difference is immaterial.

How saturated can a model be? By Example 4.4.2, we cannot hope to find an $|M|^+$-saturated M. What about the next best thing?

[15] Although of course one can enumerate an infinite tuple on its cardinality to show that κ-saturated models realise all types in κ variables over a set of size $< \kappa$.

[16] This is not a typo: while for κ-saturation we require a condition for sets of size $< \kappa$, for universality the inequality is not strict.

[17] Clearly, only finitely many a_i will appear in φ.

[18] If we were checking the consistency of an arbitrary set of formulas, we should have replaced φ with a finite conjunction $\bigwedge_i \varphi_i$. But $p(x)$ is closed under conjunctions, hence so is $q(x)$.

Definition 4.5.3. A model M is *saturated* iff it is $|M|$-saturated.

We will say something about the existence of such models in the next section. For now, let us look at their properties.

Theorem 4.5.4. Suppose that M and N are saturated models of the same cardinality. Then $M \cong N$. In fact, every partial elementary map $M \to N$ with domain of size $< |M|$ extends to an isomorphism $M \to N$.

Proof. Suppose M, N have cardinality κ. Fix enumerations $(a_\alpha)_{\alpha<\kappa}$ of M and $(b_\alpha)_{\alpha<\kappa}$ of N. Take the proof of Theorem 2.1.2 and the proof of Theorem 4.5.2. Add ice, shake well, strain into a chilled martini glass. Garnish with a lemon twist (optional).

For the second statement, suppose $A \subseteq M$ has cardinality $\mu < \kappa$, and f is a partial elementary map with domain A. Make sure that the enumeration $(a_\alpha)_{\alpha<\kappa}$ starts by enumerating A, that is, $A = a_{<\mu}$,[19] and that $(b_\alpha)_{\alpha<\kappa}$ starts by enumerating $f(A)$ accordingly, that is, that for every $\alpha < \mu$ we have $b_\alpha = f(a_\alpha)$. Now argue as above, but starting the back-and-forth at stage μ. $\square$

The second part of the previous theorem is particularly interesting in the special case where $M = N$.

Corollary 4.5.5. Every saturated M of cardinality κ is *strongly κ-homogeneous*: every partial elementary map from M into itself with domain of cardinality $< \kappa$ extends to an element of $\mathrm{Aut}(M)$. In particular, if $|A| < \kappa$ then $\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$ if and only if a and b are in the same $\mathrm{Aut}(M/A)$-orbit.

Proof. Everything is immediate, except perhaps the "in particular" bit, so let's spell that out. One direction is Exercise 4.3.1, and does not even need strong homogeneity. In the other direction, note that "the map $\mathrm{id}_A \cup \{a \mapsto b\}$ is elementary" is just a fancy way of saying "$\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$", and that any extension of this map to an automorphism will, by definition, belong to $\mathrm{Aut}(M/A)$. $\square$

If you are wondering why there is a "strongly" before "homogeneous" above, you may want to know that there is also a weaker notion of κ-homogeneity, requiring only that partial elementary maps with domain smaller than κ may be extended to one extra point. An argument in the same spirit as the other ones in this section shows that κ-saturated models, of whatever cardinality, are κ-homogeneous. This notion is involved in some characterisations of saturation, see e.g. [Poi00, Chapter 9].

Saturation is sometimes phrased as a matter of being "large". This is inaccurate, or at least a bit odd, since if M is "large" and $N \succeq M$, we would expect N to be "large" as well. For saturation, this is false:

Exercise 4.5.6. Find a cardinal κ, a κ-saturated M, and an $N \succeq M$ which is not κ-saturated.

Saturation is closer to being compact, since it tells us that the intersection of certain families with the finite intersection property (cf. Remark 4.1.3) is nonempty.

[19] Usual abuses of notation about treating tuples as sets apply.

4.6 Monster models

In some contexts, we need to realise a lot of types, and deal with several models at once. A common convention is to choose a κ-saturated model $\mathfrak{U}$, for κ "larger than everything we want to consider", and embed everything in there (elementarily), which we may do by Theorem 4.5.2. That way, instead of saying, for example, "let $N \succeq M$ contain a realisation $a \vDash p(x)$", we may just convene at the start that everything we mention lives inside $\mathfrak{U}$, and simply say "let $a \vDash p(x)$". Under this convention, that is, only considering elementary substructures $M \preceq \mathfrak{U}$, applying Exercise 4.4.9 with $M_2 = \mathfrak{U}$ tells us that all inclusions between these M are elementary.

Now, the "larger than everything we want to consider" is justified by Theorem 4.4.10: if we need to consider larger sets (for example, because we want to realise a type over $\mathfrak{U}$), we may pass to an elementary extension of $\mathfrak{U}$ with a higher degree of saturation. So far, so good. But, while it may not be immediately clear why, we would like $\mathfrak{U}$ to also be strongly κ-homogeneous, since a lot of proofs may be simplified by using so-called "automorphism arguments" (we will see one of these soon). Essentially, the point is that "over small sets, types are the same as orbits" is a nice property for $\mathfrak{U}$ to have, if we want to work inside it.

By the results in the previous section, if we were able to find, for arbitrarily large κ, a saturated $\mathfrak{U}$ of cardinality κ (as opposed to, merely, a κ-saturated $\mathfrak{U}$), then we would be happy: our $\mathfrak{U}$ would be $|\mathfrak{U}|$-strongly homogeneous, and even uniquely determined by its cardinality.

If we assume additional set theoretic assumptions, then this can be done: if $\kappa \geq |L|$ and $\kappa^+ = 2^\kappa$, then by Theorem 4.4.10 there is a saturated model of size κ^+, hence if GCH holds, or at least if it holds at arbitrarily large cardinals, then we can find arbitrarily large saturated models.

But what if we want the "small subsets" of $\mathfrak{U}$ to be closed under some construction which, for example, sends A to something of size $2^{2^{|A|}}$? Clearly, taking $\mathfrak{U}$ to be saturated of size κ^+ and declaring "small" to mean "of size $\leq \kappa$" is not a good idea. Things would be better if we had a saturated model of size κ for κ a strong limit, that is, such that $\lambda < \kappa \implies 2^\lambda < \kappa$. But in the proof of Theorem 4.4.10 we used regularity of κ^+, so we would like some κ which is regular, a strong limit, and larger than $|L|$, so at the very least uncountable. In other words, we want arbitrarily large *strongly inaccessible* cardinals to exist. If we go, consistency-wise, a bit beyond ZFC, and assume a proper class of strongly inaccessible cardinals, then we are once again done. The reason is that the proof of Theorem 4.4.10 may be adapted to show the following:[20]

Exercise 4.6.1. If $\kappa \geq |L|$ is strongly inaccessible, there is a saturated model of cardinality κ.

If we want to stay within the reach of ZFC though, we cannot assume instances of GCH, let alone a proper class of inaccessibles. So we proceed as follows.[21] We show that, for every M and every $\kappa \geq |L|$, there is $\mathfrak{U} \succeq M$

[20] This will also follow from what we will do in this subsection, but if you try to solve the exercise now then you will probably stumble on the idea underlying the constructions below.

[21] Another possible approach is to work in NBG instead of ZFC, and build a class-sized, set-saturated monster model. Yet another approach (for those who know a bit more set theory) is this: several theorems have an arithmetic conclusion (code formulas in a countable language

which is κ-saturated and κ-strongly homogeneous. We may then choose a "large enough" strong limit κ, declare "small" to mean "of size $< \kappa$", and work in $\mathfrak{U}$. If we need to deal with larger things, we enlarge κ to κ' and pass to a κ'-monster $\mathfrak{U}_1 \succ \mathfrak{U}$.

Notation 4.6.2. If κ is a cardinal, we denote by C_κ the set of cardinals strictly small than κ.

Definition 4.6.3. Suppose that $|M| = \kappa$. We call M *special* iff it has a *specialising chain*, that is, iff it is the union of an elementary chain

$$M = \bigcup_{\mu \in C_\kappa} M_\mu$$

such that each M_μ is μ^+-saturated.

The idea is to fix a strong limit κ, and prove that if $\mathfrak{U}$ is special and of carefully chosen cardinality (spoiler: it will not be κ), then $\mathfrak{U}$ is κ-saturated and κ-strongly homogeneous. Since there are always arbitrarily large strong limit (not necessarily regular) cardinals, this will suffice. You may object that we need to build not just arbitrarily large $\mathfrak{U}$ but, for arbitrarily large M, some special $\mathfrak{U} \succeq M$, but we will get this for free from saturation because of Theorem 4.5.2.

Remark 4.6.4. If M is saturated, we may take as a specialising chain the one constantly M. So saturated models are special.

Theorem 4.6.5. Let $\kappa > |L|$ be a strong limit. Then there is a special $\mathfrak{U} \vDash T$ of size κ.

Proof. The strategy of proof is as follows: we build a suitable elementary chain, then take as $\mathfrak{U}$ its union. We then trim the chain we built to obtain a specialising chain for $\mathfrak{U}$.

Since κ is a strong limit, we can find an increasing $\operatorname{cof}(\kappa)$-sequence of cardinals $(\kappa_i \mid i < \operatorname{cof}(\kappa))$ such that[22]

$$\kappa = \sum_{i<\operatorname{cof}(\kappa)} \kappa_i = \sum_{i<\operatorname{cof}(\kappa)} 2^{\kappa_i}$$

where without loss of generality $\kappa_0 > |L|$. Start with M_0 any model of size κ_0.[23] At successor stages, use Theorem 4.4.10 to obtain an $|M_i|^+$-saturated $M_{i+1} \succeq M_i$ of size at most $2^{|M_i|}$. If i is a limit ordinal, set $M_i \coloneqq \bigcup_{j<i} M_j$, invoke Exercise 1.4.4 to get elementarity, and observe that

$$|M_i| = \left|\bigcup_{j<i} M_j\right| \le \sup_{j<i} 2^{\kappa_j} \le 2^{\kappa_i}$$

Set $\mathfrak{U} \coloneqq \bigcup_{i<\operatorname{cof}(\kappa)} M_i$, then trim $(M_i \mid i < \operatorname{cof}(\kappa))$ by choosing any weakly increasing function $\iota \colon C_\kappa \to \operatorname{cof}(\kappa)$ with the property that $\kappa_{\iota(\mu)} \ge \mu$. The required specialising chain is $(M_{\kappa_{\iota(\mu)}} \mid \mu \in C_\kappa)$. It is easy to check that $|\mathfrak{U}| \le \kappa$; if the inequality was strict, we would easily get a contradiction by using Theorem 4.5.2 to obtain an embedding of $\mathfrak{U}$ inside one of the pieces of the specialising chain above, say $M_{\kappa_{\iota(\mu)}}$, and observing that $M_{\kappa_{\iota(\mu^+)}}$ has larger cardinality. $\square$

inside $\mathbb{N}$). Use absoluteness results to assume GCH without loss of generality.

[22] Here i ranges on *ordinals* less than $\operatorname{cof}(\kappa)$.

[23] If you want to prove directly the existence of special $\mathfrak{U} \succeq M$, you may start with $M_0 = M$ (of course this requires taking κ large enough).

Theorem 4.6.6. If $\mathfrak{U}_0, \mathfrak{U}_1$ are special models of T of the same cardinality, then they are isomorphic.

Proof. We prove this by back and forth along carefully built enumerations. Let $\kappa := |\mathfrak{U}_0| = |\mathfrak{U}_1|$, and write our special models as unions of fixed specialising chains $\mathfrak{U}_0 = \bigcup_{\mu \in C_\kappa} M^0_\mu$ and $\mathfrak{U}_1 = \bigcup_{\mu \in C_\kappa} M^1_\mu$.

Claim 4.6.7. For $\ell < 2$, there are enumerations $(a^\ell_i)_{i<\kappa}$ of $\mathfrak{U}_\ell$, possibly with repetitions, such that, if $\mu \in C_\kappa$, then $(a^\ell_i \mid i < \mu^+) \subseteq M^\ell_\mu$.

Proof of the Claim. Clearly the assumptions are the same for $\mathfrak{U}_0$ and $\mathfrak{U}_1$, so in the proof of this claim we drop ℓ from the notation. Fix an enumeration without repetitions $(b_j \mid j < \kappa)$ of $\mathfrak{U}$. Inductively, define a_i as the first b_j in $M_{|i|}$ not yet enumerated as one of the a_- if one exists, and as a_0 otherwise. This clearly gives as the desired property, but we need to check that we have indeed enumerated all points of $\mathfrak{U}$. Towards a contradiction, let j_0 be minimal such that b_{j_0} does not equal any of the a_i. Let μ be minimum such that $b_{j_0} \in M_\mu$, and look at $(a_i \mid \mu \le i < \kappa)$. This can contain only elements b_j with $j < j_0$, and never repeating any of those twice, which is impossible because $|\kappa \setminus \mu| = \kappa > |j_0|$. □ CLAIM

We can now proceed by back-and-forth, by inductively building a partial elementary map f_i such that $a^0_i \in \operatorname{dom} f_i$ and $\operatorname{dom} f_i \subseteq M^0_{|i|}$, while $a^1_i \in \operatorname{im} f_i$ and $\operatorname{im} f_i \subseteq M^1_{|i|}$.[24] It is possible to ensure this because we are using the enumerations given by the Claim: first of all, $a^0_i \in M^0_{|i|}$. To define f_i on a^0_i, we need to realise a type over a subset of $M^1_{|i|}$ of size $|i|$; by definition of specialising chain, this can be done inside $M^1_{|i|}$, so our inductive assumption is preserved. With a symmetric argument, we then ensure $a^1_i \in \operatorname{im} f_i$, then move to $i+1$. □

While proving that certain objects are unique is always very satisfying, perhaps counterintuitively uniqueness of special models will not be especially useful *per se*, but rather because it implies strong homogeneity. This is proven via the following trick.

Exercise 4.6.8. Let $\mathfrak{U}$ be special of cardinality κ, and let $A \subseteq \mathfrak{U}$ have size $|A| < \operatorname{cof}(\kappa)$. Then $\mathfrak{U}_A$ is special.

Corollary 4.6.9. Every special $\mathfrak{U}$ of size κ is $\operatorname{cof}(\kappa)$-saturated and $\operatorname{cof}(\kappa)$-strongly homogeneous.

Proof. The $\operatorname{cof}(\kappa)$-saturation of $\mathfrak{U}$ is an easy consequence of the definition of "special", while $\operatorname{cof}(\kappa)$-strong homogeneity is a consequence of Theorem 4.6.6 and the previous exercise: if $a \mapsto b$ is a partial elementary map $M \to M$ with $|a| < \operatorname{cof}(\kappa)$, just expand L with constants c_i to be interpreted as a_i in the first copy of M and b_i in the second one. □

If you are kidnapped by an evil wizard who is about to make you magically forget everything contained in the present subsection except for one sentence of your choice, I strongly recommend you save the following corollary.

Corollary 4.6.10. For every M and every infinite cardinal κ, there is $\mathfrak{U} \succeq M$ which is κ-saturated and κ-strongly homogeneous.

[24] Note: we are not necessarily mapping $a^0_i \mapsto a^1_i$.

Proof. By Theorem 4.5.2 and Corollary 4.6.9[25] we only need to show that there are arbitrarily large strong limit cardinals of arbitrarily large cofinality. Recall that $\beth_0 := \aleph_0$, $\beth_{\alpha+1} := 2^{\beth_\alpha}$, and $\beth_\lambda := \sup_{\mu<\lambda} \beth_\mu$ for limit λ. Since $\beth$ is increasing and continuous, it has arbitrarily large fixed points, that is, cardinals μ with $\beth_\mu = \mu$. Note that fixed points of $\beth$ are strong limit cardinals. Enumerate increasingly and continuously the fixed points of $\beth$ on the ordinals, say with a class-function f. If α is a limit *ordinal*, then $f(\alpha) = \sup_{\beta<\alpha} f(\beta)$, hence $\operatorname{cof}(f(\alpha)) = \operatorname{cof}(\alpha)$. Therefore, if μ is a regular cardinal, then $\operatorname{cof}(f(\mu)) = \mu$, and we conclude by choosing any regular $\mu > \kappa$. $\square$

While saturation is clearly preserved under reducts, for κ-strong homogeneity this is *not* the case, in general. A good reason to work in special models, instead of just any κ-saturated, κ-strongly homogeneous one, is the following.

Remark 4.6.11. It follows easily from the definitions that if $L' \subseteq L$ and $\mathfrak{U}$ is special, then so is $\mathfrak{U} \upharpoonright L'$. In particular, $\mathfrak{U} \upharpoonright L'$ is $\operatorname{cof}(|\mathfrak{U}|)$-strongly homogeneous

Exercise 4.6.12. Find a cardinal κ and a structure M such that M is κ-strongly homogeneous, but some reduct of M is not.[26]

4.7 Working in a monster model

Notation 4.7.1. From now on the following conventions and notations apply.

1. We fix a strong limit cardinal κ "larger than everything we want to consider", we work inside a special model $\mathfrak{U}$ with $\operatorname{cof}(|\mathfrak{U}|) > \kappa$, and if we say that something is *small* we mean it has size $< \kappa$. We call $\mathfrak{U}$ the *monster model*.

2. We write $\vDash \varphi(a)$ for $\mathfrak{U} \vDash \varphi(a)$, etc.

3. We write $\vDash \varphi(x)$ for $\vDash \forall x\ \varphi(x)$.

4. We write $A \subset^+ \mathfrak{U}$ to mean that A is a small subset of $\mathfrak{U}$, and $M \prec^+ \mathfrak{U}$ to mean that M is a small elementary substructure of $\mathfrak{U}$. We use $A, B, \ldots$ to denote small sets.

5. We write $a \equiv_A b$ to mean that $\operatorname{tp}(a/A) = \operatorname{tp}(b/A)$.

6. All tuples, small sets, etc. are assumed to be inside $\mathfrak{U}$, and when we say "model" we mean "small elementary substructure of $\mathfrak{U}$", unless the "model" is $\mathfrak{U}$ is self, or unless we specify otherwise.[27]

7. *Definable* means "$\mathfrak{U}$-definable", and *formula* means "$L(\mathfrak{U})$-formula". If parameters are not allowed, we write "L-formula", or "$L(\emptyset)$-formula". More generally, we say "$L(A)$-definable" or "A-definable" if we only allow parameters from A.

[25]. . . , and the fact that saturation in a cardinal implies saturation in all the smaller ones, and similarly for strong homogeneity,. . .

[26]Hint: take as $\operatorname{Th}(M)$ the expansion of DLO by a predicate interpreted as an initial segment with no supremum.

[27]Well, or unless I forget to specify. Sorry.

8. If we say that two formulas are "equivalent", we mean modulo $\mathrm{ED}(\mathfrak{U})$.

If you went through the previous section too quickly,[28] recall that $A \subset^+ \mathfrak{U}$ implies that

1. every $p(x) \in S_{<\omega}(A)$ is realised in $\mathfrak{U}$, and
2. if $a \equiv_A b$, then there is $f \in \mathrm{Aut}(\mathfrak{U}/A)$ such that $f(a) = b$.

Remark 4.7.2. By compactness and saturation, every infinite definable subset of $\mathfrak{U}^n$ is not small.

Let us look at topological proof of a statement not involving topology.[29]

Proposition 4.7.3. Let $\varphi(x)$ be an $L(\mathfrak{U})$-formula. Then $\varphi(x)$ is equivalent to some $L(A)$-formula if and only if whenever $a, b \in \mathfrak{U}^{|x|}$ and $a \equiv_A b$ then $\vDash \varphi(a) \leftrightarrow \varphi(b)$.

Proof. Left to right is obvious, so let us prove right to left.

Let $\pi\colon S_x(\mathfrak{U}) \to S_x(A)$ be the restriction function. Consider the clopen subsets $[\varphi(x)]$ and $[\neg\varphi(x)]$ of $S_x(\mathfrak{U})$. Since their union is $S_x(\mathfrak{U})$, and since π is surjective, $\pi([\varphi(x)]) \cup \pi([\neg\varphi(x)]) = S_x(A)$. Now, π is a continuous function between compact Hausdorff spaces, hence it is closed. But our assumptions and $|A|^+$-saturation of $\mathfrak{U}$ imply that $\pi([\varphi(x)]) \cap \pi([\neg\varphi(x)]) = \emptyset$, hence $\pi([\varphi(x)])$ and $\pi([\neg\varphi(x)])$ are closed sets which are the complement of each other, and are therefore clopen. We conclude by Exercise 4.2.4. $\square$

Strong homogeneity tells us that types over A are the same as orbits over A. But what about formulas?

Proposition 4.7.4. Let $X \subseteq \mathfrak{U}^n$ be a definable subset of $\mathfrak{U}$. Then X is fixed setwise by every element of $\mathrm{Aut}(\mathfrak{U}/A)$ if and only if X is A-definable.

Proof. Right to left is obvious. For left to right we need to show that, if $\varphi(x)$ is a formula defining X, then $\varphi(x)$ is equivalent to some $L(A)$-formula. If not, by Proposition 4.7.3 there are $a \equiv_A b$ with $a \vDash \varphi(x)$ and $b \vDash \neg\varphi(x)$. Let $f \in \mathrm{Aut}(\mathfrak{U}/A)$ be such that $f(a) = b$. Then f does not fix X setwise, against our assumptions. $\square$

In Section 4.3 we used that two elements were conjugated under $\mathrm{Aut}(M/A)$ to show that they had the same type over A. Now that we work in a monster $\mathfrak{U}$, by strong homogeneity we also have the converse, and we may freely confuse types over small sets A with orbits of $\mathrm{Aut}(\mathfrak{U}/A)$. Let us a look at this in action by using an "automorphism argument" to prove some statement which does not involve automorphisms.

Definition 4.7.5. We say that $a \in \mathfrak{U}^{|a|}$ is

1. *definable over* A iff $\{a\}$ is A-definable; in other words, iff there is an $L(A)$-formula $\varphi(x)$ such that $\vDash \varphi(a)$ and $\varphi(x)$ has only one solution;

[28] . . . , or if you recently encountered an evil wizard,. . .

[29] The proposition above may also be proven with an argument similar to that we used for Theorem 2.5.8. Compare the lengths of the two proofs.

2. *algebraic over A* iff a belongs to a finite A-definable set; in other words, iff there is an $L(A)$-formula $\varphi(x)$ such that $\vDash \varphi(a)$ and $\varphi(x)$ has only finitely many solutions.

We denote by $\mathrm{dcl}(A)$ (respectively, $\mathrm{acl}(A)$) the set of points of $\mathfrak{U}^1$ definable (respectively, algebraic) over A.

So $\mathrm{dcl}(A)$ is the union of all A-definable singletons (in $\mathfrak{U}^1$) and $\mathrm{acl}(A)$ the union of all finite A-definable sets (again, subsets of $\mathfrak{U}^1$).

Remark 4.7.6. If $\varphi(x)$ has exactly m solutions, then there is a sentence in $\mathrm{ED}(\mathfrak{U})$ saying this. So if $b \in \mathrm{acl}(A)$ and $b \equiv_A c$, then $c \in \mathrm{acl}(A)$. In particular:

1. if $b \in \mathrm{acl}(A)$, then $\mathrm{tp}(b/A)$ has only finitely many realisations in $\mathfrak{U}$ (if $b \in \mathrm{dcl}(A)$, then "finitely many" is actually "only one"), and they need to be contained in every M with $A \subseteq M \preceq \mathfrak{U}$;

2. $\mathrm{acl}(A)$ is fixed by $\mathrm{Aut}(\mathfrak{U}/A)$ setwise (and, trivially, $\mathrm{dcl}(A)$ is fixed by $\mathrm{Aut}(\mathfrak{U}/A)$ pointwise).

Remark 4.7.7. By strong homogeneity, $a \in \mathrm{dcl}(A)$ if and only if a is fixed by $\mathrm{Aut}(\mathfrak{U}/A)$, and $a \in \mathrm{acl}(A)$ if and only if the orbit of a under $\mathrm{Aut}(\mathfrak{U}/A)$ is finite.

Here is the promised automorphism argument.

Proposition 4.7.8. The set $\mathrm{acl}(A)$ is the intersection of all models containing A. More precisely, $\mathrm{acl}(A)$ equals the intersection of all $M \prec^+ \mathfrak{U}$ with $M \supseteq A$.

Proof. The inclusion $\subseteq$ follows from Remark 4.7.6. Suppose $b \notin \mathrm{acl}(A)$. Then $\mathrm{tp}(b/A)$ has infinitely many realisations in $\mathfrak{U}$, and by compactness and saturation they cannot all be contained in a fixed small model, so there is $M' \prec^+ \mathfrak{U}$ which does not contain some $b' \equiv_A b$. Let $f \in \mathrm{Aut}(\mathfrak{U}/A)$ be such that $f(b) = b'$. Then $M \coloneqq f^{-1}(M')$ does not contain b. □

While we are here, let us also observe this.

Proposition 4.7.9. The operators dcl and acl are closure operators.

Proof. We prove this for acl, the proof for dcl is similar (and easier). Clearly, acl is *extensive*, that is, $A \subseteq \mathrm{acl}(A)$, and *monotone*, that is, if $A \subseteq B$ then $\mathrm{acl}(A) \subseteq \mathrm{acl}(B)$. We need to prove that acl is *idempotent*, that is, $\mathrm{acl}(\mathrm{acl}(A)) = \mathrm{acl}(A)$. The inclusion $\supseteq$ follows from extensivity and monotonicity. For the other inclusion, we use Remark 4.7.7. Let $a \in \mathrm{acl}(\mathrm{acl}(A))$, as witnessed by an $L(A)$-formula $\varphi(x, w)$ and parameters $b \in \mathrm{acl}(A)$. The fact that $\varphi(x, b)$ has finitely many solutions is a property of $\mathrm{tp}(b/A)$, so if $c \equiv_A b$ then $\varphi(x, c)$ still has finitely many solutions. Since $b \in \mathrm{acl}(A)$, there are only finitely many $c \equiv_A b$, and it follows that the $\mathrm{Aut}(\mathfrak{U}/A)$-orbit of a is contained in the finite definable set $\bigvee_{c \equiv_A b} \varphi(x, c)$. □

Warning: if $\varphi(x)$ is an $L(A)$-formula with finitely many solutions satisfied by b, then it is not in general true that all solutions of $\varphi(x)$ will realise $\mathrm{tp}(b/A)$. In the proof above, we only needed one inclusion. Anyway, for a careful choice of $\varphi(x)$, this is true: you can prove it as a warm up for the next chapter.

Exercise 4.7.10. Suppose $b \in \mathrm{acl}(A)$. Show that there is an $L(A)$-formula $\varphi(x)$ isolating $\mathrm{tp}(b/A)$, that is, such that in $S_x(A)$ we have $[\varphi(x)] = \{\mathrm{tp}(b/A)\}$.

Chapter 5

Realising few types

5.1 Isolated types

In the previous chapter, we built models realising many types. But what if we want to build a model where a certain type, maybe even a partial one, is not realised? Certain types must always be realised: think of the partial type $\{x = x\}$, or of the complete[1] type $\{x = 0\}$ in ACF_p. On the other hand, in ACF_0, say, the generic type over $\mathbb{Q}$ is not realised in $\mathbb{Q}^{\mathrm{alg}}$. Why can this type be *omitted*?

Definition 5.1.1. A model M *omits* a partial type $\pi(x)$ iff there is no $a \in M$ such that $M \vDash \pi(a)$.

Let us begin to clarify the matter by answering Question 4.2.5. If $p(x) \in S_x(A)$ is isolated, it means that there is $\varphi(x) \in L(A)$ such that $\{p(x)\} = [\varphi(x)]$. In other words, any realisation of $\varphi(x)$ automatically realises the whole of $p(x)$. This has the following consequence.

Proposition 5.1.2. If $p(x) \in S_x(A)$ is isolated, then every model containing A realises p.

Proof. If $\varphi(x)$ isolates $p(x)$, then in particular $\varphi(x)$ is consistent, which means that $\vDash \exists x\ \varphi(x)$. Every model containing A also contains the parameters appearing in $\varphi(x)$, so it must contain a witness a to that existential quantifier, hence $a \vDash p(x)$. □

So there is no hope to omit isolated types. What about the rest? We will deal with this shortly, but first let us finish answering Question 4.2.5.

Corollary 5.1.3. If $p(x) \in S_x(M)$ is isolated, then there is $m \in M$ with $p(x) = \{x = m\}$.

Proof. M is clearly a model containing M, and the conclusion follows easily from the previous proposition. □

We already said that realised types are always isolated, and over a model isolated types are realised. You may wonder if this characterises models. The answer is negative.

[1] As usual, up to deductive closure.

Exercise 5.1.4. Find an example where all isolated types in $S_x(A)$ are realised, but A is not a model.

Nevertheless, we *can* characterise models in a slightly different fashion.

Proposition 5.1.5. For $A \subseteq \mathfrak{U}$, the following are equivalent.

1. The set of realised types is dense in $S_1(A)$.
2. $A \preceq \mathfrak{U}$.

Proof. We already saw in point 9 of Remark 4.2.2 that one direction holds, even for n-types, with n arbitrary. For the converse, we apply the Tarski–Vaught test: saying that $\mathfrak{U} \vDash \exists x\ \varphi(x)$ means that $[\varphi(x)]$ is nonempty; by assumption, in $S_x(A)$, the set $[\varphi(x)]$ contains a realised type, that is, there is $a \in A$ such that $\mathfrak{U} \vDash \varphi(a)$. □

5.2 Omitting types

If $p(x) \in S_x(A)$ is not isolated, can we omit it? In general, the answer is no. If you insist on T being complete, there is a slightly involved counterexample which we will see later, Example 5.2.10. If you are happy to see a partial type $\pi(x)$ over $\emptyset$ in an *incomplete* theory T that cannot be omitted, even though there is no $\varphi(x) \in L(\emptyset)$ with $\varphi(x) \vdash \pi(x)$, here is the standard example.

Example 5.2.1. Let $L = \{c_i \mid i < \aleph_1\} \cup \{d_j \mid j < \omega\}$, and let T say that the c_i are pairwise distinct. Then $\pi(x) = \{x \neq d_j \mid j < \omega\}$ is realised in every model, but it is not implied by any $\varphi(x)$.

Proof. The first part is clear. For the second part, suppose $\varphi(x) \vdash \pi(x)$. Since $\varphi(x)$ can only mention finitely many d_j, there is some d_{j_0} it does not mention. It is then easy to construct $M \vDash T$ with $M \vDash \varphi(d_{j_0})$, and we are done. □

Nevertheless, if we are working over a *countable* language, then nonisolated types over $\emptyset$ can be omitted. This follows from the Omitting Types Theorem, proven below. What about countable L, but over uncountably many parameters? Again, Example 5.2.10 below shows that even $\aleph_1$ parameters may be too much.

Long story short, we need to assume that both L and A are countable. So we may as well throw A into L, that is, pass to $L(A)$, and just work over $\emptyset$. Since we are already over the empty set, we may work in a bit more generality and talk of omitting partial types in incomplete theories. Type spaces over $\emptyset$ still make sense, but now $S_0(\emptyset)$ may have more than one point, and we are not allowed to use $\mathfrak{U}$ (its theory determines a completion!).[2]

Let us give a name to "there is a consistent $\varphi(x)$ such that $\varphi(x) \vdash \pi(x)$". If $\pi(x)$ is complete, this already has a name: $\pi(x)$ is isolated. In fact, even for partial $\pi(x)$, this already has a name:

Remark 5.2.2. There is some consistent $\varphi(x)$ such that $\varphi(x) \vdash \pi(x)$ if and only if the associated closed set $[\pi(x)]$ of $S_x(\emptyset)$ has nonempty interior.

[2]Yes, I know, we just started using it. It will come back soon, I promise.

Theorem 5.2.3 (Omitting Types Theorem)**.** Let T be a possibly incomplete theory in a countable L, and $\{\Phi_n(x^n) \mid n < \omega\}$ a family of partial types over $\emptyset$, where each $|x^n|$ is finite. If every $[\Phi_n(x^n)]$ has empty interior, then there is a countable $M \vDash T$ omitting every $\Phi_n(x^n)$.

The proof of this will be slightly intricate, and will need bookkeeping rather than magic: omitting types is more difficult than realising them and, as we saw, sometimes it is so difficult it cannot even be done. According to "a not well-known model theorist" quoted in [Sac72], *"Any fool can realise a type, but it takes a model theorist to omit one."*

Proof. Let C be a countable set of fresh constants. Fix the following.

1. An enumeration $(\sigma_i \mid i < \omega)$ of all $L(C)$-sentences.
2. An enumeration with repetitions $(c^i \mid i < \omega)$ of the set $C^{<\omega}$ of all finite tuples of constants from C, with the property that every element of $C^{<\omega}$ is listed infinitely many times (build it by using your favourite bijection $\omega \to \omega^2$).

We start with $T_0 = T$, and inductively build and increasing chain of theories $(T_i \mid i < \omega)$ with the following properties.

(a) Each $T_i \setminus T$ is finite; that is, at each stage we add only finitely many sentences.[3] This will be needed to keep the construction going.

(b) $T' \coloneqq \bigcup_{i<\omega} T_i$ is complete.

(c) T' is a *Henkin theory*: for every $L(C)$-formula $\varphi(y)$ with $|y| = 1$ there is $c \in C^1$ such that $T' \vdash (\exists y\ \varphi(y)) \to \varphi(c)$.

(d) For all $n < \omega$ and all $c \in C^{|x^n|}$ there is $\varphi(x^n) \in \Phi_n(x^n)$ such that $T' \vdash \neg\varphi(c)$.

It is (lengthy but) easy to prove that a complete Henkin $L(C)$-theory T' has a model M where every $m \in M$ is the interpretation of a constant symbol in C: you take as M the quotient of C by the equivalence relation "$T' \vdash c = c'$", use completeness to decide the interpretations of the symbols of L, and then check that everything is well-defined and works.[4] Therefore, if we manage to carry out the construction, we are done: by point (d), (the reduct to L of) such an M (which is clearly countable) will omit every $\Phi_n(x^n)$.

Before the construction, for each $c \in C^{<\omega}$, write down a list $\ell(c)$ of all $\Phi_n(x^n)$ with $|x^n| = |c|$, of order type a natural number or ω. During the construction, we will cross them out one by one.[5] The $i + 1$-th stage of the construction goes as follows.

[3]Recall that the construction only has ω steps. Also, of course, here we do *not* take deductive closures of our theories.

[4]If you have never seen this construction, you may want to do this as an exercise. Otherwise, you can see the details being spelled out in [TZ12, Lemma 2.2.3], for instance.

[5]We will keep referring to $\ell(c)$ as $\ell(c)$ even after crossing out some elements, computer science-style. If you prefer an extra index to a slight abuse of notation, add an index $\ell_k(c)$, say that this "list" is an function with domain a subset of ω, and instead of saying that "d is crossed out from the list" say that $\ell_{k+1}(c) \coloneqq \ell_k(c) \upharpoonright (\mathrm{dom}(\ell_k) \setminus (\ell_k(c))^{-1}(\{d\}))$. But I think this proof already has enough indices, so I have relegated k to this footnote.

(i) Look at σ_i, from the fixed enumeration of all $L(C)$-sentences. Since inductively T_i is consistent, the union of T_i with at least one between σ_i or $\neg\sigma_i$ is still consistent; choose one between the two which is consistent with T_i, and call it σ. This will be added to T_{i+1} to ensure that T' is complete.

(ii) If σ is of the form $\exists y\ \varphi(y)$, since by inductive assumption we only added finitely many formulas to T, there is $c \in C$ which we have not used so far. Let $T'_i \coloneqq T_i \cup \{\sigma\} \cup \{(\exists y\ \varphi(y)) \to \varphi(c)\}$. If σ is not of that form, just set $T'_i \coloneqq T_i \cup \{\sigma\}$. This will ensure that T' is Henkin. Since c had not been mentioned yet, T'_i is easily seen to be consistent.

(iii) Look at c^i from our enumeration with repetitions of $C^{<\omega}$. Look at the list $\ell(c^i)$, and let $\Phi_n(x^n)$ be the first one which we have not crossed out yet. Again, inductively we only added finitely many formulas to T. Let $\psi(c')$ be their conjunction, where c' is the tuple of all constants in C we mentioned so far, (including the constants in c^i) so $T'_i = T \cup \{\psi(c')\}$. Write $c' \coloneqq (c^i, \tilde{c})$. Because $[\Phi_n(x^n)]$ has empty interior, we have $T \cup \{\exists z\ \psi(c^i, z)\} \nvdash \Phi_n(c^i)$, and by Lemma 2.5.3 $T \cup \psi(c') \nvdash \Phi_n(c^i)$. Therefore, there must be $\varphi(x^n) \in \Phi_n(x^n)$ such that $T \cup \{\psi(c')\} \cup \{\neg\varphi(c^i)\}$ is consistent. Set $T_{i+1} \coloneqq T'_i \cup \{\neg\varphi(c^i)\}$, and cross $\Phi_n(x^n)$ out of $\ell(c^i)$. Note that $T_{i+1} \setminus T_i$ has size at most 3, so inductively $T_{i+1} \setminus T$ is finite.

Fix $c \in C^{<\omega}$. Since c appears as c^i for infinitely many i, and the list $\ell(c)$ is of order type a natural number or ω, every $\Phi_n(x^n)$ with $|x^n| = |c|$ will eventually get crossed out of $\ell(c)$, hence point (d) is taken care of. Congratulations, you are now a model theorist! $\square$

By the way, the "countable" in the statement may have been added a posteriori, without knowing anything about the proof, using Löwenheim–Skolem, because of the following easy observation:

Remark 5.2.4. If N omits p and $M \preceq N$, then M omits p.

Remark 5.2.5. What about the converse of the Omitting Types Theorem? If T is complete, it holds: we have already shown in Proposition 5.1.2 that isolated types cannot be omitted, and a similar proof shows that neither can partial types with nonempty interior. On the other hand, if T is not complete, $\exists x\ \varphi(x)$ may be true in some $M \vDash T$ and false in some $N \vDash T$; even if $\varphi(x)$ isolates a type, it will be omitted in N.

But what about omitting an *arbitrary* family of, say, nonisolated complete types in a complete theory? This is too much to hope for:

Exercise 5.2.6. Find a complete theory T with infinite models such that no element of $S_1(\emptyset)$ is isolated.[6]

If $a \in M \vDash T$, with T as above, then $\mathrm{tp}(a/\emptyset)$ is clearly not omitted. But ok, the set of types we tried to omit here was clearly too fat, namely, it was the whole type space. *Meagre* sets can instead be omitted.

[6] If you want to do this exercise, do it now, since a solution is buried in the next few pages. Hint: a solution is also buried in the previous ones.

Corollary 5.2.7 (Omitting Types Theorem on steroids)**.** Let T be a possibly incomplete theory in a countable L. For every $m \in \omega$, let X_m be a meagre subset of $S_m(\emptyset)$. Then there is a countable $M \vDash T$ omitting every element of every X_m.

Recall that X is *meagre* iff it is contained in a countable union of closed sets, each with empty interior. Recall also that compact Hausdorff spaces are in particular locally compact, and that the Baire Category Theorem holds for locally compact Hausdorff spaces: meagre subsets of $S_n(\emptyset)$ have empty interior, and no nonsense is happening here.

Proof. By assumption, for each n there are partial types $\Phi_{m,n}(x^{m,n})$ with empty interior such that $X_m \subseteq \bigcup_{n<\omega}[\Phi_{m,n}(x^{m,n})]$. If a model M omits all $\Phi_{m,n}(x^{m,n})$, then *a fortiori* M also omits all the elements of each X_m. Now choose your favourite bijection $\omega^2 \to \omega$ and apply the Omitting Types Theorem. □

For a more topological proof, see [Poi00, Section 10.1]. But now it is time for counterexamples.

Example 5.2.8. The converse of the Omitting Types Theorem on steroids is false, even for complete theories. Namely, there are non-meagre sets that can be omitted. In fact, even comeagre ones (that is, with meagre complement). For example, consider $T_{2^{<\omega}}$, and let $Y \subseteq S_1(\emptyset)$ be the set of types corresponding to eventually constant elements of 2^ω. This is clearly countable, and no point of this space is isolated, so Y is meagre. Its complement $X \coloneqq Y^{\complement}$ is by definition comeagre. Since Y intersects every clopen set in infinitely many points (there are infinitely many eventually constant functions with a given finite restriction!), it is easy to see from the axioms of $T_{2^{<\omega}}$ that Y can be made into a model of $T_{2^{<\omega}}$ omitting all types in X.

Example 5.2.9. We cannot omit partial types in infinitely many variables, not even countably many. In DLO, let $x = (x_i \mid i < \omega)$ and let $\Phi(x) \coloneqq \{x_{i+1} < x_i \mid i < \omega\}$. Clearly, $\Phi(x)$ is not implied by any single formula, for example because a single formula can only mention finitely many variables. Basically by definition, M omits $\Phi(x)$ if and only if it is well-ordered. But of course no DLO is well-ordered.

The counterexample below, from [Fuh62], is slightly involved, so I was about to just cite it, but I am not aware of any source describing it in English.

Example 5.2.10. There is a complete T, in a language L with $|L| = \aleph_1$, containing a partial type over $\emptyset$ with empty interior that cannot be omitted.

Proof. Start with a language L_0 with three sorts[7] X, Y, F and a relation symbol R of arity $X \times Y \times F$. Write an L_0-theory saying the following.

(i) Each of X, Y, F is infinite.

(ii) For each $f \in F$, the formula $R(x, y, f)$ defines the graph of a bijection between X and Y.

[7] If you have skipped Section 0.3, this could be a good point to read it. Everything can be done with one sort and predicates, but then you need to say every time that certain predicates partition the universe, that relations are trivial outside of their intended domain, etc.

(iii) For every $f \in F$, and every bijection $\gamma\colon : X \to Y$ which differs from $R(x, y, f)$ only in finitely many points, there is $g \in F$ such that $R(x, y, g)$ is the graph of γ (of course you will need one axiom for every n, where n is the size of the set where these functions differ).

Now do the following:

1. Fix a countable model M_0 of the theory above.
2. Enlarge L_0 to L_1 by adding constants $\{a_i \mid i < \omega\}$ naming all elements of $X(M_0)$. Interpret these in the obvious way, that is, consider $(M_0)_{X(M_0)}$. Call it M_0'.
3. Take an elementary extension $M_1 \succeq (M_0)_{X(M_0)}$ such that $|Y(M_1)| = \aleph_1$.
4. Enlarge L_1 to L by adding constants $\{b_j \mid j < \aleph_1\}$ naming every element of $Y(M_1)$.
5. Let M be the natural expansion of M_1 to an L-structure (that is, $M = (M_1)_{Y(M_1)}$), and take as T the complete L-theory of M.

Working in M, consider $\pi(x) \coloneqq \{x \neq a_i \mid i < \omega\}$, where x is a variable of sort X. This is a partial type over $\emptyset$ only using L_1-formulas. Clearly, $\pi(x)$ cannot be omitted in any $N \equiv M$, since F contains witnesses that X and Y are in bijection and Y is uncountable. To conclude, we need to show that there is no consistent L-formula implying it. Suppose there is. Recall that L is just $L_1(B)$, where $B = \{b_j \mid j < \aleph_1\}$, and write such a formula as $\varphi(x, b')$, where b' is a suitable finite tuple of the b_j with no repetitions and $\varphi(x, y)$ is an L_1-formula. Since $\varphi(x, b')$ is consistent, it has a solution, and since $\varphi(x, b') \vdash \pi(x)$, every solution is different from each a_i. Let $c \in B^{|b'|}$ be arbitrary. By staring long enough at axiom (iii), you can convince yourself that the map sending $b'_\ell \mapsto c_\ell$ extends to an automorphism of M_1 which fixes $X(M_1)$ pointwise.[8] This implies that $M \vDash \forall x\ \varphi(x, b') \leftrightarrow \varphi(x, c)$. Since $(b_j \mid j < \aleph_1)$ is a list of *all* elements of $Y(M)$, it is easily shown that $M \vDash \forall x\ \varphi(x, b') \leftrightarrow \psi(x)$, where $\psi(x) \coloneqq \exists y\ (\varphi(x, y) \wedge \theta(y))$, for $\theta(y)$ the formula saying that the y_i are pairwise distinct (as we chose the b'_i to be). Now, $\psi(x)$ is an $L_1(\emptyset)$-formula, satisfied by some point of X but by no a_i, and implying the $L_1(\emptyset)$-partial type $\pi(x)$. All this information never mentions the b_j, so it is written in $\mathrm{Th}(M_1)$, and by elementarity also in $\mathrm{Th}((M_0)_{X(M_0)})$. This is clearly nonsense, since in M_0 every point of X is one of the a_i. □

All these counterexamples may make you think that in uncountable setting you should really give up thinking about omitting types. You shouldn't. There *is* a version of the Omitting Types Theorem for uncountable L: a partial type which cannot be implied by less than $|L|$ formulas can be omitted, see [CK90, Theorem 2.2.19]. You can show it an exercise by adapting the proof of the "vanilla" Omitting Types Theorem.

5.3 Prime models

We go back to the usual setting of complete T with infinite models. We fix a monster model $\mathfrak{U}$ and work inside it.

[8]The automorphism is of M_1, not of M! In M, the b_j are named, so automorphisms must fix them.

Definition 5.3.1. If $M \vDash T$ and $A \subseteq M$, we call M *prime over A* iff it embeds elementarily over A in every model containing A, that is, in every model of the $L(A)$-theory $\mathrm{Th}(M_A)$. We call M *prime* iff M is prime over $\emptyset$.

What can we say about prime models?

Remark 5.3.2. By Löwenheim–Skolem, if M is prime over A, then $|M| \leq |L| + |A|$.

In general, the bound may be strict.

Example 5.3.3. Let X be your favourite infinite set. Equip X its full structure, that is, add a predicate symbol for every subset of every X^n, and make X into a structure in this language in the obvious way. Clearly, X is a prime model of its theory, but the language has cardinality $2^{|X|}$.

The Omitting Types Theorem allows us to say something else very quickly.

Definition 5.3.4. We call M *atomic over A* iff the only $p(x) \in S_x(A)$ which are realised in M are isolated. We say just *atomic* instead of *atomic over $\emptyset$*.

Proposition 5.3.5. If M is prime over A, and both L, A are countable, then M is atomic over A.

Proof. Suppose that $a \in M^n$ is such that $\mathrm{tp}(a/A)$ is nonisolated. Since everything is countable, by the Omitting Types Theorem there is a countable N omitting $\mathrm{tp}(a/A)$. Good luck embedding M (elementarily) into N. $\square$

Of course, we may have given the definitions of primality and atomicity the other way around, by defining first "prime" and "atomic", and then introducing parameters by adding them to the language. In particular, if A is countable, we may pretend to be working over $\emptyset$. Nevertheless, keeping track of parameters is important, since it allows us to state things like the following.

Proposition 5.3.6 (Monotonicity and transitivity of isolation)**.** The type $\mathrm{tp}(ab/A)$ is isolated if and only if $\mathrm{tp}(b/A)$ and $\mathrm{tp}(a/Ab)$ are isolated.

Proof. Left to right, suppose $\varphi(x, y)$ isolates $\mathrm{tp}(ab/A)$. This means that whenever $\psi(x, y)$ is in the latter, then $\vDash \forall x, y\ (\varphi(x, y) \to \psi(x, y))$. This implies two things. Firstly, that $\vDash \forall x\ \varphi(x, b) \to \psi(x, b)$, and since $\psi(x, y) \in \mathrm{tp}(ab/A)$ is arbitrary, this implies that $\varphi(x, b)$ isolates $\mathrm{tp}(a/Ab)$. Secondly, in the special case where $\psi(x, y)$ is of the form $\theta(y)$, that is, it does not mention x, it implies that $\vDash \forall y\ (\exists x\ \varphi(x, y)) \to \theta(y)$. Therefore $\exists x\ \varphi(x, y)$ isolates $\mathrm{tp}(b/A)$.

Right to left, observe that if $\varphi(x, y) \in L(A)$ is such that $\varphi(x, b)$ isolates $\mathrm{tp}(a/Ab)$, then this is written in $\mathrm{tp}(b/A)$, in the form of formulas the likes of $\forall x\ (\varphi(x, y) \to \chi(x, y))$. If additionally, as we are assuming, $\psi(y)$ isolates $\mathrm{tp}(b/A)$, then it follows easly that $\psi(y) \wedge \varphi(x, y)$ isolates $\mathrm{tp}(ab/A)$. $\square$

Corollary 5.3.7. Suppose that M is atomic over A. Then, for every finite tuple $b \in M$, we also have that M is atomic over Ab.

Proof. Fix $a \in M^n$. By assumption, $\mathrm{tp}(ab/A)$ is isolated, and we conclude by Proposition 5.3.6. $\square$

Theorem 5.3.8. Let L be countable.

1. Up to isomorphism there is at most one atomic countable model of T.
2. Countable atomic models are ω-strongly homogeneous.
3. A countable model is atomic if and only if it is prime.
4. A model is prime if and only if it is atomic and countable.

Proof. Let M, N be countable atomic models. The empty function $M \to N$ is elementary because T is complete, so we may fix enumerations of M, N of order type ω and start to build an isomorphism by back-and-forth. Use the notation $f_n \colon A_n \to B_n$ for the partial elementary function built at stage $n-1$ (surjective on B_n). In the "forth" part (and as usual, the "back" is symmetrical), at stage n, say we are presented with $a \in M$. Look at $\operatorname{tp}(a/A_n)$. Since a comes from the atomic model M, this type is isolated, say by $\varphi(x, A_n)$, with $\varphi(x, y) \in L(\emptyset)$. Now remember that A_n is a finite tuple, and observe that, as observed in the previous proof, the fact that $\varphi(x, A_n)$ isolates a complete type over A_n is written in $q(y) \coloneqq \operatorname{tp}(A_n/\emptyset)$. Inductively $A_n \equiv B_n$, hence the type obtained from $\operatorname{tp}(a/A_n)$ by replacing each element of A_n with the corresponding element of B_n via f_n is isolated. It must be therefore realised in N, and the back-and-forth can continue, proving the first part.

For second part, suppose $a_i \mapsto b_i$ is partial elementary map $M \to M$ with finite domain. Add constants c_i to the language, and expand M to M_0 by interpreting c_i as a_i, and to M_1 by interpreting c_i as b_i. By Corollary 5.3.7 M_0 and M_1 are still atomic, and we conclude by applying first part.

For the third part, we already know one implication; the converse is proven observing that arguing as above but only going "forth" allows to embed a countable atomic model in an arbitrary one.

Finally, the fourth part is immediate from the third one and Remark 5.3.2. □

Hence, in the countable case, prime models are unique, but we still haven't said anything about their existence. As you probably expect, we will end this section by proving a big theorem showing that every complete countable theory has a prime model. Just kidding, this is blatantly false, and you already know a counterexample:

Remark 5.3.9. Since there are countable theories with no isolated types over $\emptyset$, atomic models, and a fortiori prime ones, need not always exist.

Which theories have prime models then, and over which sets? When are they unique? We have already seen some partial answers, and will see more below, but you should know that we are taking a peek at a rabbit hole that is deeper than you probably expect. When everything is countable, anyway, we already have the tools to prove a very satisfying topological characterisation. Again, if we have countably many parameters we may (and will) throw them in the language and work over $\emptyset$.

Theorem 5.3.10. Let L be countable. The following are equivalent.

1. T has a prime model.
2. T has an atomic model.

3. For every n, the set of isolated n-types is dense in $S_n(\emptyset)$

Proof. The first two statements are equivalent by Theorem 5.3.8, Löwenheim–Skolem, and Remark 5.2.4. Suppose now that M is atomic, and let $\varphi(x) \in L(\emptyset)$ be consistent. Since M is a model, there is $a \in M^{|x|}$ with $\vDash \varphi(a)$. But then $\operatorname{tp}(a/\emptyset)$ is isolated and belongs to $[\varphi(x)]$.

Conversely, suppose that isolated types are dense. For $n \in \omega$, consider the set of $L(\emptyset)$-formulas

$$\Phi_n(x_{<n}) \coloneqq \{\neg\varphi(x_{<n}) \in L(\emptyset) \mid \varphi(x_{<n}) \text{ isolates a complete } n\text{-type}\}$$

If we can find a model M omitting all $\Phi_n(x_{<n})$, then we are done: a tuple $a \in M^n$ cannot satisfy $\Phi_n(x_{<n})$, hence by definition we must have $\vDash \varphi(a)$ for some $\varphi(x_{<n})$ isolating a complete type. Now, some of the Φ_n will be inconsistent, so will be automatically omitted. If we show that all the other ones have empty interior, then we can invoke the Omitting Types Theorem and conclude. So suppose that for some n the closed set $[\Phi_n] \subseteq S_n(\emptyset)$ has nonempty interior, that is, there is $\psi(x_{<n}) \vdash \Phi_n(x_{<n})$. By hypothesis, $[\psi(x_{<n})]$ contains an isolated type, say isolated by $\varphi(x_{<n})$, so in particular $\varphi(x_{<n}) \vdash \psi(x_{<n})$. By definition of $\Phi_n(x_{<n}) \vdash \neg\varphi(x_{<n})$ and by combining all of the above we get to $\varphi(x_{<n}) \vdash \neg\varphi(x_{<n})$, so $[\varphi(x_{<n})]$ is empty and cannot contain a type, let alone isolate it. □

Even in the absence of countability (in particular, we cannot use the Omitting Types Theorem) density of isolated types over *every* set is an assumption strong enough to grant prime models.

Theorem 5.3.11. Suppose that, for every A, in $S_1(A)$, the isolated points are dense. Then, for every A, there is a prime model over A.

Proof. Let $\mu \coloneqq |A| + |L|$. By counting isolating formulas, we see that $S_1(A)$ contains at most μ isolated types; list them as $(p_i \mid i < \mu)$, by possibly repeating some of them if necessary. Inductively, define a chain $(A_i \mid i < \mu)$ as follows.

1. Start with $A_0 \coloneqq A$.
2. If A_i realises p_i, let $A_{i+1} \coloneqq A_i$. Otherwise, consider the projection map $\pi\colon S_1(A_i) \to S_1(A)$. Since p_i is isolated, $\pi^{-1}(\{p_i\})$ is open. By assumption, it contains an isolated point, call it q_i. Note that q_i extends p_i. Let $a_i \vDash q_i$, and set $A_{i+1} \coloneqq A_i a_i$.
3. At limit stages, take unions.

Set $B_0 \coloneqq \bigcup_{i<\mu} A_i$.

Claim 5.3.12. If $N \supseteq A$, then there is $B_0' \subseteq N$ with $B_0' \equiv_A B_0$.[9]

Proof of the Claim. This is another "only forth" argument: we build an elementary map $B_0 \to N$ by induction on i, the same i we used to build B_0. The only nontrivial case is when A_{i+1} was obtained by adding a_i to A_i. Inductively, we may assume to have already embedded A_i into N, say as A_i'. Translate q_i into $q_i' \in S_1(A_i')$ according to this embedding, note that it is still isolated, and take $a_i' \in N$ realising it to continue the embedding. □ CLAIM

[9] Yes, we are looking at a type in infinitely many variables.

Note that $|B_0| \leq \mu$ and iterate the construction, obtaining, for every $j \in \omega$, some B_{j+1} realising all isolated types in $S_1(B_j)$ and such that whenever $B_j \subseteq N$, then B_{j+1} can be embedded in N over B_j, in the same sense as above. Therefore, $M := \bigcup_{j<\omega} B_j$ can be embedded over A in any $N \supseteq A$, and we only need to show that M is a model. By Proposition 5.1.5, it is enough to show that the realised points are dense in $S_1(M)$. So take $\varphi(x) \in L(M)$. Since formulas may only mention finitely many parameters, there is $j \in \omega$ such that $\varphi(x) \in L(B_j)$. By assumption, $[\varphi(x)] \subseteq S_1(B_j)$ contains an isolated type, which is realised in B_{j+1}, say by b. Then $\{x = b\} \in [\varphi(x)] \subseteq S_1(M)$, and we are done. □

So we better have criteria to know when the isolated types are dense. Sometimes, it is just a matter of counting.

Proposition 5.3.13. Let L be countable. If $S_n(\emptyset)$ is countable (possibly finite), then the isolated n-types are dense.

This does not just have one topological proof, but two!

First proof. The union of countably many nonisolated points is meagre. By the Baire Category Theorem for (locally) compact Hausdorff spaces, its complement is dense. □

Second proof. If the isolated n-types are not dense, there is a nonempty $[\varphi(x)]$ containing none. Inductively, partition $[\varphi(x)]$ in two nonempty clopen sets, which of course will still contain no isolated points. By iterating this, we build a complete binary tree of height ω of clopen sets, ordered by (reverse?[10]) inclusion. This tree has $2^{\aleph_0}$ branches, and the intersection of each branch is nonempty by compactness, hence $S_n(\emptyset)$ is uncountable. □

Exercise 5.3.14. Prove that if L is countable, and $S_n(\emptyset)$ is uncountable, then it must have size at least continuum.[11]

Of course, density of isolated types does not imply countability of type spaces. For example, take your favourite countable T with an uncountable type space, e.g. $T_{2^{<\omega}}$, and fix a countable model M. Then $\mathrm{ED}(M)$ is a countable theory T', and types over $\emptyset$ in T' are, by definition, the same as types over M in T; since $S_1(\emptyset)$ was uncountable in T, it is a fortiori uncountable in T' (if you prefer, look at the surjective restriction map $S_1(M) \to S_1(\emptyset)$). But since we named all elements of a model, in T' the isolated types are dense by Proposition 5.1.5.

5.4 The number of countable models

Let us begin with an easy lemma.

Lemma 5.4.1. If L is countable, then every type over $\emptyset$ can be realised in some countable model.

Proof. Given $p(x) \in S_n(\emptyset)$, let $a \models p(x)$. Use Löwenheim–Skolem to take a countable M containing a. □

[10] It depends on whether you like your trees to grow upwards or downwards.

[11] Hint: adapt the second proof, by showing that if $[\varphi(x)]$ is uncountable then it can be partitioned in two uncountable clopen sets.

A remarkable consequence of the results in the previous section is that we can prove that a countable theory has a prime model *by just counting types* (isn't that wonderful?). As you may expect, when for every n we can only find *finitely* many n-types over $\emptyset$, instead of merely countably many, something special happens. This was realised in 1959 by several[12] people independently.

Theorem 5.4.2 (Ryll-Nardzewski–Svenonius–Engeler)**.** Let L be countable. The following are equivalent.

1. For every n, the space $S_n(\emptyset)$ is finite.
2. For every n, the space $S_n(\emptyset)$ is discrete.
3. For every n, every $p \in S_n(\emptyset)$ is isolated.
4. For every n, if $|x| = n$, there are only finitely many $L(\emptyset)$-formulas $\varphi(x)$ up to equivalence modulo T.
5. For every n and every $M \vDash T$, there are only finitely many $\emptyset$-definable subsets of M^n.
6. Every $M \vDash T$ is atomic.
7. Every countable $M \vDash T$ is atomic.
8. There is $M \vDash T$ which is countable, saturated, and atomic.
9. T has a prime model, and it is saturated.
10. T has only one countable model up to isomorphism.
11. There is $M \vDash T$ which, for every n, realises only finitely many n-types.
12. There is a countable $M \vDash T$ such that $\mathrm{Aut}(M)$ is *oligomorphic*, that is, for every n the diagonal action $\mathrm{Aut}(M) \curvearrowright M^n$ has finitely many orbits.
13. For every countable $M \vDash T$, the permutation group $\mathrm{Aut}(M)$ is oligomorphic.

Proof. Finite Hausdorff spaces are discrete, every point is isolated if and only if the space is discrete, and a compact discrete space is finite, which proves $1 \Rightarrow 2 \Leftrightarrow 3 \Rightarrow 1$. Formulas over $\emptyset$ up to equivalence are the same as $\emptyset$-definable sets, and if there are only finitely many, then a finite boolean combination of them is enough to imply a complete type over $\emptyset$, showing $5 \Leftrightarrow 4 \Rightarrow 1$. Now, $\emptyset$-definable sets are the same as clopen subsets of $S_n(\emptyset)$, so $1 \Rightarrow 5$, since a finite set has only finitely many subsets. If all types are isolated, all models have no choice but to only realise isolated types, and if all models do so, in particular so do the countable ones, so $3 \Rightarrow 6 \Rightarrow 7$. By Lemma 5.4.1 every type over $\emptyset$ can be realised in a countable model, so $7 \Rightarrow 3$.

Below, we use that we have already proven the equivalences above. If for every n there are finitely many n-types, Exercise 4.4.11, tells us that there is a countable saturated model, so $7 \Rightarrow 8$. Such a model realises all types over $\emptyset$, so $8 \Rightarrow 3$. If we throw in Theorem 5.3.8, we also discover immediately that $8 \Leftrightarrow 9$ and that $7 \Rightarrow 10$. If T has only one countable model M, then it must

[12] Three, not four: "Ryll-Nardzewski" is a single surname.

have at most countably many types, so by Proposition 5.3.13 the isolated types are dense, and by Theorem 5.3.10 T has a prime model, which by 5.3.2 must be M. Again by the fact that T has countably many models, and again by Exercise 4.4.11, T has a countable saturated model, which must again be M, so $10 \Rightarrow 9$.

Again, using the previously proven equivalences, if T has only one countable model, then it must be saturated, hence ω-strongly homogeneous, so types over $\emptyset$ are the same as orbits over $\mathrm{Aut}(M)$. Since there are finitely many types, this gives $10 \Rightarrow 13$. Trivially, $13 \Rightarrow 12$, and being in the same orbit implies having the same type, so in an arbitrary model there are at most as many types as orbits, hence $12 \Rightarrow 11$. Finally, assume that M realises only finitely many n-types, say $p_0(x), \dots, p_k(x)$. Since $S_n(\emptyset)$ is Hausdorff, we can find $\varphi_0(x), \dots, \varphi_k(x)$ such that $[\varphi_i(x)]$ contains $p_i(x)$ but no $p_j(x)$ for $j \neq i$. Clearly, $\varphi_i(M) = p_i(M)$. Now, take an arbitrary $\psi(x) \in L(\emptyset)$. In any model, the set of its realisations is a (possibly infinite) union of sets of realisations of complete types over $\emptyset$. In M the only n-types are the p_i, so

$$\psi(M) = \bigcup_{\substack{i \le k \\ p_i(x) \vdash \psi(x)}} p_i(M) = \bigcup_{\substack{i \le k \\ p_i(x) \vdash \psi(x)}} \varphi_i(M)$$

Since M is a model, this implies that $\psi(x)$ is equivalent to $\bigvee_{\substack{i \le k \\ p_i(x) \vdash \psi(x)}} \varphi_i(x)$. There are at most 2^{k+1} formulas of this form, and $\psi(x)$ was arbitrary, and k does not depend on ψ, so $11 \Rightarrow 4$, and I leave to you the pleasant task of checking that the directed graph on 13 vertices we built above is connected. □

Remark 5.4.3. If you want a more succinct statement to remember, a common choice is $1 \Leftrightarrow 10$. The choice requiring the least number of definitions in order to be stated is probably $4 \Leftrightarrow 10$.

Take a moment to appreciate the different nature of the statements which we have just proven to be equivalent: some of them are topological, some are dynamical, and some, of course, model-theoretic. Some conditions can be checked on an arbitrary model, and some of them are just a matter of counting. On the other hand, some tell us about the existence of special structures, and one is a uniqueness statement. Something enjoying such a diverse array of characterisations clearly deserves a name. We take the opportunity to also give a name to other things you have already seen (and will keep seeing later on).

Definition 5.4.4. Let κ be an infinite cardinal. A theory T is *κ-categorical* iff it has at most one model of cardinality κ. We also say *ω-categorical* to mean $\aleph_0$-categorical.

Here we are assuming T complete, but recall that, by Exercise 0.4.17, if $\kappa \ge |L|$ and T has no finite models, then κ-categoricity implies completeness. Usually, when people talk of ω-categorical theories, they implicitly also mean that L is countable. This is important, since the Ryll-Nardzewski theorem does not generalise to uncountable languages.

Example 5.4.5. Let M be ω viewed as a structure in the language L with a symbol $<$ for the usual order and a unary predicate P_X for every $X \subseteq \omega$,[13] and

[13] If you prefer, just take all subsets of all ω^n.

let T be its complete L-theory. Note that $|L| = 2^{\aleph_0}$ but, by construction, T has a countable model, namely M. We prove below that every other model has size at least $2^{\aleph_0}$. This tells us two things:

1. The assumption of countability of L is necessary in the Ryll-Nardzewski theorem: in this theory, $S_1(\emptyset)$ is essentially the same as the space $\beta\omega$ of ultrafilters over ω (if you prefer, the Stone–Čech compactification of ω with the discrete topology), which notoriously has size $2^{2^{\aleph_0}}$.
2. Löwenheim–Skolem does not generalise to arbitrary cardinalities: if you assume $\neg\mathsf{CH}$, this example also shows that below $|L|$ there can be gaps in the possible cardinalities of models.

Proof. Let $\{A_i \mid i < 2^{\aleph_0}\}$ be an *almost disjoint* family of infinite subsets of ω, that is, a family such that for all $i \neq j$ the intersection $A_i \cap A_j$ is finite.[14] Observe that

1. If $i \neq j$, because $A_i \cap A_j$ is finite, it must have a maximum, call it a_{ij}.
2. Since every A_i is infinite, for every x there is $y > x$ such that $y \in A_i$.

Note that these properties are written in T. Moreover, since we named every subset of ω, in particular we named singletons, hence T says that there can be no point between n and $n+1$. Therefore, every $N \neq M$ must contain some $c > \omega$. By what we said above, for every i there must be $d \in N$ with $d > c$ and $N \vDash P_{A_i}(d)$. If $j \neq i$, since $d > c > \omega$, in particular $d > a_{ij}$, hence $N \vDash \neg P_{A_j}(d)$. So $|N| \geq 2^{\aleph_0}$. □

Remark 5.4.6. Naming finitely many parameters preserves ω-categoricity. On the other hand, naming even $\aleph_0$ many does not. You can convince yourself of both statements very quickly by counting types.

Remark 5.4.7. There is a powerful method to build ω-categorical theories (and more), known as taking a *Fraïssé limit.* By now you have seen more than enough to understand this construction, but for time reasons I (sadly) have to redirect you to the literature, see for example [Hod93, Chapter 7]. Several theories we have seen in this course are the theory of some Fraïssé limit: $\mathsf{DLO}, T_{\mathrm{rg}}$, the theory of infinite sets, the theory of κ generic equivalence relations for $\kappa \leq \aleph_0$, and T^*_{feq} are all examples.

As we saw above, countable (complete) theories with only one countable model enjoy quite striking properties. Countable (complete) theories with only *two* countable models enjoy an even more striking property: they do not exist.

Theorem 5.4.8 (Vaught's never two)**.** There is no complete countable first-order theory with exactly two countable models up to isomorphism.

Proof. Suppose T is a counterexample. Inside two countable models, for every n, there is only space to realise $\aleph_0$ many n-types over $\emptyset$. By Lemma 5.4.1 every type over $\emptyset$ is realised in a countable model, so for every n we have $|S_n(\emptyset)| \leq \aleph_0$. By Exercise 4.4.11, there is a countable saturated $M_2 \vDash T$.

[14] For example, you can build such a family by putting ω in bijection with $2^{<\omega}$ and taking the family 2^ω of branches of this tree.

Moreover, by Theorem 5.3.10 and Proposition 5.3.13, T has a prime model M_0. We now build a third model M_1. Since T has more than one countable model, by Theorem 5.4.2, for some n there is a nonisolated $p(x) \in S_n(\emptyset)$. If a is a realisation, because a is a finite tuple, every $S_m(a)$ is still countable[15], so there is M_1 which is prime over a, that is, $(M_1 a)$ is a prime model of $T_a := \mathrm{Th}(\mathfrak{U}_a)$. In particular, M_1 is countable. Since M_1 realises p and M_0 does not, we clearly have $M_1 \not\cong M_0$, so we need to show $M_2 \not\cong M_1$. But if $M_1 \cong M_2$, then M_1 would be saturated. But ω-saturated models stay saturated after naming finitely many constants, so (M_1, a), is a saturated model of T_a. Because it is also a prime model of T_a, by Theorem 5.4.2 T_a has finitely many n-types over $\emptyset$, that is, $S_n(a)$ is finite. Now take two different n-types q_0, q_1 over $\emptyset$ in T. These can be seen as a partial types over $\emptyset$ in T_a, hence be completed to distinct $\hat{q}_0, \hat{q}_1 \in S_n(a)$. Since in T the space $S_n(\emptyset)$ is infinite, this is a contradiction. $\square$

Exercise 5.4.9. Find an incomplete T in a countable L with exactly two non-isomorphic countable models.

Funnily enough, this characterises 2 among the positive natural numbers.

Exercise 5.4.10. Let $L = \{<\} \cup \{c_i \mid i < \omega\}$, and let T be $\mathsf{DLO} \cup \{c_i < c_{i+1} \mid i < \omega\}$.

1. Prove that T is complete and has quantifier elimination.
2. Prove that, up to isomorphism, the countable models of T are the expansions of $(\mathbb{Q}, <)$ obtained as follows.
   (a) For every $i < \omega$, the constant c_i is interpreted as i.
   (b) For every $i < \omega$, the constant c_i is interpreted as $-1/(i+1)$.
   (c) The sequence $(c_i)_{i<\omega}$ is interpreted as an increasing sequence converging (in $\mathbb{R}$) to an irrational number.
3. Which of these is prime? Which is saturated?
4. For $n \geq 4$, let $L_n := L \cup \{P_0, \dots, P_{n-3}\}$, where every P_j is a 1-ary predicate. Let T_n be the union of T with the axioms saying that every P_j is dense, that the P_j partition the domain, and that every c_i is in P_0. Prove that T_n has exactly n countable models up to isomorphism.

Therefore, the number of countable models of a complete theory in a countable language can be any positive natural number except two. What about infinite cardinals?

Exercise 5.4.11. Let T be a complete L-theory, with L countable. Prove that

1. Every such T has at most $2^{\aleph_0}$ countable models up to isomorphism.
2. There is such a T with $2^{\aleph_0}$ pairwise nonisomorphic countable models.
3. There is such a T with exactly $\aleph_0$ pairwise nonisomorphic countable models.

[15] Uncountably many m-types over a would yield uncountably many $|a| + m$ types over $\emptyset$.

What about cardinals between $\aleph_0$ and $2^{\aleph_0}$, of course when they exist, that is, when CH fails? This has been open for more than half a century.

Conjecture 5.4.12 (Vaught's conjecture)**.** If a complete theory in a countable language has uncountably many countable models, then it has continuum many.

By a theorem of Morley (*not* the one we will see later on in the course), the only case to exclude is that of a countable complete theory with exactly $\aleph_1$ countable models (again, obviously, assuming $\neg\mathsf{CH}$). As far as I know, the conjecture is open,[16] but it has been proven to hold in for special classes of theories. There is also a more general version of the conjecture, known as the *topological Vaught conjecture*, stated in terms of Polish groups acting on Polish spaces.

5.5 Ehrenfeucht–Mostowski models

We saw in Theorem 5.3.11 that, if over *every* A the isolated types are dense, then there is a prime model over every A, regardless of the size of L. This theorem is very powerful: for example the theory DCF_0 of differentially closed fields of characteristic 0 satisfies its assumptions, and as a consequence every differential field of characteristic 0 has a *differential closure*. We will not deal with differential fields in this course, so I refer the interested reader to the literature, see e.g. [Poi00, Section 6.2].

Nevertheless, the assumptions of Theorem 5.3.11 are quite strong, and at any rate, it does not gives *arbitrarily large* models realising few types. The main theorem of this section will do exactly that, with no restriction on $|L|$. The idea is to start with a suitable sequence $(a^i)_{i\in I}$ of points which "look all the same", and then to take some kind of model M "enveloping" this sequence. Intuitively, if the tuples from $(a^i)_{i\in I}$ all look the same, then M will realise few types. On the way to the main theorem, we will also encounter a property implying that over every A the isolated types are dense. But let us begin with precise statements.

Notation 5.5.1. If $(I,<)$ is a linear order, we denote an I-sequence of tuples of the same length by $a^I := (a^i)_{i\in I}$. If I is not specified, or clear from context, we also just say "sequence" instead of "I-sequence". If for example $I=\omega$, we also write $a^{<\omega}$, to make it clear that we are not referring to the ω-th element of some a^I indexed on, say, $I=\kappa$.

We put the index as a superscript (as in: a^i) because each a^i is a tuple, not necessarily of length 1. So, for example, a^i_1 denotes the second element of the i-th tuple in a^I. In the literature it is also common to write a_i for a^i since, as you will see, we will rarely have to look at the coordinates of a^i.

Definition 5.5.2. Let A be a set of parameters. We say that a^I is *A-indiscernible*, or *indiscernible over A*, iff, for every $n\in\omega$, if $i_0<\ldots<i_n$ and $j_0<\ldots<j_n$, then

$$\operatorname{tp}(a^{i_0},a^{i_1},\ldots,a^{i_n}/A)=\operatorname{tp}(a^{j_0},a^{j_1},\ldots,a^{j_n}/A)$$

We also say just *indiscernible* instead of "$\emptyset$-indiscernible".

[16] In fact, counterexamples have been announced, but their status is unclear.

Remark 5.5.3. Almost by definition, an I-sequence is indiscernible over A if and only if the type over A of $a^{i_0}, \dots, a^{i_n}$ only depends on the order type of $i_0, \dots, i_n$ inside I, that is, on $\operatorname{qftp}^I(i_0, \dots, i_n/\emptyset)$.[17]

In the definition above, if I is finite then of course we only need to look at the $n < |I|$. Or only at the $n < |I| - 1$, if you want. At any rate, indiscernible sequences are typically interesting when I is infinite, usually with no maximum. Some of the things below also make sense for finite I. Anyway, if you want, assume that I is always an infinite linear order.

Example 5.5.4. In ACF_p, let a^I be a sequence of A-transcendental elements which are algebraically independent over A. Then a^I is I-indiscernible.

Non-Example 5.5.5. In ACF_0, let a^ω be a sequence of elements with $a^0 \notin \{0,1\}$ and $a^1 = a^0 \cdot a^0$. Then a^ω is not $\emptyset$-indiscernible (hence, for all A it is not A-indiscernible).

Proof. If it was, by indiscernibility we would have[18] $a^2 = a^0 \cdot a^0$ and $a^0 \cdot a^0 = a^2 = a^1 \cdot a^1 = a^0 \cdot a^0 \cdot a^0 \cdot a^0$. Since $a^0 \neq 0$, we must have $a^0 \cdot a^0 = 1$. Since $a^0 \neq 1$, we must have $a^0 = -1$. Again by indiscernibility, a^ω is constantly -1. From $a^1 = a^0 \cdot a^0$ we get $-1 = (-1) \cdot (-1)$, a contradiction. $\square$

Example 5.5.6. In DLO, a sequence $(a^i)_{i<\omega}$, with $|a^i| = 1$ is A-indiscernible if and only if

1. all a_i have the same cut in A (possibly degenerate, that is possibly they are all equal to a fixed $a \in A$), and
2. the sequence is either
   (a) constant,
   (b) increasing, or
   (c) decreasing.

Of course, the collection of types of finite pieces of an indiscernible sequence deserves a name. We define this for arbitrary sequences; in general, it will not be a complete type.

Definition 5.5.7. Let a^I be an I-sequence[19] of tuples of the same length. The *Ehrenfeucht–Mostowski type* $\mathrm{EM}(a^I/A)$ of a^I over A is the set of formulas

$$\mathrm{EM}(a^I/A) := \{\varphi(x^0, \dots, x^n) \in L(A) \mid n < \omega, \forall i_0 < \dots < i_n \in I, \vDash \varphi(a^{i_0}, \dots, a^{i_n})\}$$

So $\mathrm{EM}(a^I/A)$ is the set of those formulas over A which are true in all finite pieces of a^I *provided they are enumerated increasingly.*

Remark 5.5.8.

[17] "Why are we taking as I a linear order? What happens if we take a different structure? And give a similar definition?" If I is a set with no structure, the answer to this will appear in due course in this course. People have also equipped at I with different structures, see [Sco15].

[18] Superscripts denote indices in the sequence, and *not* multiplicative powers.

[19] Not necessarily indiscernible.

1. If a^I is arbitrary, then $\mathrm{EM}(a^I/A)$ may as well be empty (up to deductive closure[20]). Example: take T the theory of infinite sets, let A be infinite, $I = |A|$, and let a^I be some enumeration of A where every point appears infinitely often.

2. On the other hand, if a^I is A-indiscernible and I is infinite, then $\mathrm{EM}(a^I/A)$ is complete type in ω (tuples of) variables, namely, it coincides with $\mathrm{tp}(a^{<\omega}/A)$. Note that this is a type in ω variables regardless of whether I is ω or another infinite linear order. This is not a bug, but a feature: it allows us to compare indiscernible sequences indexed over different linear orders.

3. Not *all* elements of $S_\omega(A)$ are the Ehrenfeucht–Mostowski types of some A-indiscernible sequence (of tuples), see for instance Non-Example 5.5.5 above.

Sometimes, we have some infinite a^I such that, for a fixed $L(A)$-formula $\varphi(x^0, \ldots, x^n)$ and all $i_0 < \ldots < i_n \in I$, we have $\vDash \varphi(a^{i_0}, \ldots, a^{i_n})$, and we want to produce an A-indiscernible J-sequence, with J an arbitrary infinite linear order, with the analogous property. Note that $\varphi(x^0, \ldots, x^n) \in \mathrm{EM}(a^I/A)$. So we want an A-indiscernible b^J with $\mathrm{EM}(b^J/A) \supseteq \mathrm{EM}(a^I/A)$. The fact that these always exist is the content of what [TZ12] calls the *Standard Lemma*, also known in the literature as "Ramsey and compactness", ominously telling us how it will be proven.

Fact 5.5.9 (Ramsey's theorem)**.** Let $k, r \in \omega \setminus \{0\}$. Denote by $X^{[k]}$ the set of subsets of X of size k. If X is infinite, then for any function $c \colon X^{[k]} \to r$, there is an infinite $H \subseteq X$ such that $c \upharpoonright H^{[k]}$ is constant.

Usually (and suggestively), c is called a *colouring*, and H a *monochromatic* set (*homogeneous* is also used). Ramsey's theorem can be proven in a number of ways, for example by induction or by using the tensor product of ultrafilters, but we will not see the proof here. If you have never seen this theorem before, here is a typical easy application: every sequence of reals has a subsequence which is either strictly decreasing, constant, or strictly increasing. To prove it, you colour $\{m, n\}$ with three colours, one for each sign of $a_m - a_n$, say where $m < n$. Then you restrict your sequence to a monochromatic set.

Lemma 5.5.10 (Standard Lemma)**.** Let I, J be infinite linear orders, with J small, and $A \subset^+ \mathfrak{U}$. For any a^I, there is an A-indiscernible b^J with $\mathrm{EM}(b^J/A) \supseteq \mathrm{EM}(a^I/A)$.

Proof. Let $\pi(x^{<\omega}) \coloneqq \mathrm{EM}(a^I/A)$. Denote by $\pi(x^J)$ the following set of formulas: for every n, choose $j_0^n < \ldots < j_n^n \in J$, let $\pi(x^{\leq n})$ be the restriction of $\pi(x^{<\omega})$ to the first n tuples of variables, and substitute $y^{j_i^n}$ for x^i inside it; take the union of all these as $n \in \omega$ varies.[21] By saturation of $\mathfrak{U}$, it is enough to show consistency of $\Phi(y^J) \coloneqq \pi(y^J) \cup \Psi(y^J)$, where $\Psi(y^J)$ says that y^J is A-indiscernible. Namely:

$$\Psi(y^J) \coloneqq \{\varphi(y^{i_0}, \ldots, y^{i_n}) \leftrightarrow \varphi(y^{j_0}, \ldots, y^{j_n}) \mid n < \omega, \varphi \in L(A), i_0 < \ldots < i_n \in J, j_0 < \ldots < j_n \in J\}$$

[20] That is, it may consist only of those $\varphi(x)$ with $\vDash \forall x\, \varphi(x)$.

[21] If J was for example an infinite ordinal, we could have just said "take $\pi(y^{<\omega})$". But note that J may not contain any copy of ω in general: for example, take as J the negative integers.

By compactness, it is enough to show that every finite subset Φ_0 of $\Phi(y^J)$ is consistent. Such a finite subset will only be able to mention a finite subsequence y of y^J, a finite subset $\Theta(y)$ of $\pi(y^J)$, and, up to enlarging Φ_0, there will be a finite set of formulas Δ such that $\Phi_0(y)$ says that the increasing tuples from y cannot be distinguished by formulas $\varphi(z^{j_0}, \ldots, z^{j_n}) \in \Delta$ (they are Δ-*indiscernible*). Let k be maximum such that there is some $\varphi(z^{j_0}, \ldots, z^{j_{k-1}}) \in \Delta$. Let $r := 2^{|\Delta|}$, and colour the k-element subsets of the original sequence a^I as follows: list each k-element set increasingly, that is, as $a^{i_0}, \ldots, a^{i_{k-1}}$ with $i_0 < \ldots < i_{k-1}$; colour it with the set of those $\varphi \in \Delta$ such that $\vDash \varphi(a^{i_0}, \ldots, a^{i_{k-1}})$. By Ramsey's Theorem, there is an infinite $I_0 \subseteq I$ such that $c \restriction a^{I_0}$ is monochromatic, that is, a^{I_0} is Δ-indiscernible. Since a^{I_0} is a subsequence of a^I, clearly it also satisfies $\Theta(y)$ which, remember, was obtained from a finite piece of $\mathrm{EM}(a^I/A)$ by a change of variables. Therefore, any subsequence of a^{I_0} of the correct length will witness that that $\Phi_0(y)$ is consistent, and we are done. $\square$

Remark 5.5.11. The assumption that J is infinite is not important: if you want a finite one, you can first build an infinite one and then trim it. On the other hand, the assumption that I is infinite is crucial: otherwise, by starting with $I = (2, 4)$, we would violate Non-Example 5.5.5.[22]

Going back to the programme sketched at the start of our section, indiscernible sequences are the promised sequences of "points that look all the same". Now we deal with the second ingredient, that is, the "enveloping" part.

Definition 5.5.12. A function $f\colon M^n \to M$ is *definable* iff its graph is a definable set. We say *A-definable* iff we only allow parameters from A in a formula defining the graph of f.

If $f(x)$ is a definable function, say its graph is defined by $\varphi(x, y)$, it is common to write $y = f(x)$ in place of $\varphi(x, y)$. More generally, one usually abuses the notation and pretends that f is an actual function symbol of L, by writing e.g. $\psi(f(x), z)$, which of course is, formally, an abbreviation for $\forall y\ \varphi(x, y) \to \psi(y, z)$. Of course, for several purposes we may as well just add functions symbols to the language.

Remark 5.5.13. If we expand the language to L' by naming every $\emptyset$-definable function by an actual function symbol, then $\mathrm{dcl}(A)$ in the sense of L is the same as (the domain of) the structure generated by A in the sense of L'.

One may also consider functions $f\colon M^n \to M^k$; trivially, their graph will be definable if and only if each of the k components of f is definable. So they may be identified with tuples of definable functions.

Definition 5.5.14. We say that that T has *definable Skolem functions* iff for every formula $\varphi(x, y)$ over $\emptyset$ with $|x| = 1$ there is an $\emptyset$-definable function f such that

$$T \vdash \forall y\ ((\exists x\ \varphi(x, y)) \to \varphi(f(y), y)) \tag{5.1}$$

[22] This is more important than it may seem: the fact that certain formulas $\varphi(x, y)$ display certain patterns on some *infinite* set of tuples is a way to say that T is in a sense "wild". Finite restrictions of these patterns are usually easy to find even when T is the theory of infinite sets.

The fact that $|x| = 1$ is not a real restriction: if T has definable Skolem functions, then you can easily prove by induction that (5.1) also holds when x is a tuple, with f now a tuple of definable functions.

Having definable Skolem functions is clearly preserved by Morleyising, and it is also easily shown that it is preserved by naming parameters. Before even looking at examples, let us make an addition to our list of easy but important facts. Note the similarity between the definition above and Henkin constructions: after all, what is a function symbol, if not a "constant symbol which depends on a tuple of arguments"? Since Henkin theories have models where all elements are the interpretation of a constant, the following is not surprising.

Remark 5.5.15. If T has definable Skolem functions, then the isolated types are dense in every $S_n(A)$, hence T has prime models over every set. Even better, for every A we have $\operatorname{dcl}(A) \preceq \mathfrak{U}$, as follows easily from the Tarski–Vaught test.

Non-Example 5.5.16. In ACF_0, it is not very difficult to show that $\operatorname{dcl}(\emptyset)$ is isomorphic to $\mathbb{Q}$, which is notoriously not algebraically closed. Therefore, ACF_0 does not have definable Skolem functions.

Non-Example 5.5.17. DLO does not have definable Skolem functions, nor does any of its expansions by constants.

Proof. Assume towards a contradiction that expanding by naming parameters from A grants definable Skolem functions. Look at the formula $\exists x\; x > y$. Pick any $c > A$, and let $d \coloneqq f(c)$. Let $d' > d$ be arbitrary. Clearly, $\operatorname{tp}(d'/Ac) = \operatorname{tp}(d/Ac)$, a contradiction, since $d' \not\models x = f(c)$. □

Example 5.5.18. We left our dear old friend $\mathcal{R}$ all alone in a corner since page 20. By now it's time to tell you that $\operatorname{Th}(\mathcal{R})$ is called DOAG, is the theory of nontrivial divisible ordered abelian groups, and eliminates quantifiers in L_{oag}, see [vdD98a, Corollary 1.7.8]. If we add a constant for any nonzero point, say positive, call it 1, then the resulting theory has definable Skolem functions[23]. The idea is to argue by induction on the dimension, starting by taking things like midpoints of intervals, or adding 1 to a to find a point in $(a, +\infty)$, see [vdD98a, Proposition 6.1.2].

We can use Remark 5.5.15 to build a model M "around" an indiscernible I-sequence in such a way that M has at least as many automorphism as $(I, <)$.

Definition 5.5.19. Suppose that T has definable Skolem functions and a^I is an indiscernible with $|a^i| = 1$. We call $M \coloneqq \operatorname{dcl}(a^I)$ the *Ehrenfeucht–Mostowski model* with *spine* a^I, or the *Skolem hull* of a^I.

Proposition 5.5.20. If T has definable Skolem functions and M is a Ehrenfeucht–Mostowski model with spine a^I, then for every $f \in \operatorname{Aut}((I, <))$ there is $\tilde{f} \in \operatorname{Aut}(M)$ with $\tilde{f}(a^i) = a^{f(i)}$.

Proof. By definition, every $b \in M$ is of the form $g(a^{i_0}, \ldots, a^{i_n})$, for some $\emptyset$-definable function g. Set $\tilde{f}(b) \coloneqq g(a^{f(i_0)}, \ldots, a^{f(i_n)})$. We need to check that $\tilde{f}$ is well-defined, because in general b may also be represented as $h(a^{j_0}, \ldots, a^{j_m})$,

[23] And even something stronger called *definable choice*, which, if you have seen the T^{eq} construction, is essentially definable Skolem functions for T^{eq}.

for a different $\emptyset$-definable h. Because a^I is indiscernible, this is not really a problem: since $f \in \mathrm{Aut}(I,<)$, we have

$$\mathrm{qftp}^I(i_0,\dots,i_n,j_0,\dots,j_m) = \mathrm{qftp}^I(f(i_0),\dots,f(i_n),f(j_0),\dots,f(j_m))$$

and by indiscernibility $\vDash g(a^{i_0},\dots,a^{i_n}) = h(a^{j_0},\dots,a^{j_m})$ if and only if

$$\vDash g(a^{f(i_0)},\dots,a^{f(i_n)}) = h(a^{f(j_0)},\dots,a^{f(j_m)})$$

The argument above shows that $\tilde{f}$ preserves and reflects the formula $g(x) = h(y)$, where x and y suitable tuples of variables. A similar argument, replacing $g(x) = h(y)$ with formulas such as $R(g_0(x^0),\dots,g_\ell(x^\ell))$, shows that $\tilde{f}$ is indeed an automorphism. □

It turns out that we *can* expand a theory to one having definable Skolem functions, but some care is needed. As having Skolem functions is preserved by Morleyising, if T does not have definable Skolem functions and we want to add them, then we *must* change the definable sets. In the case of DOAG, we got away with just changing the $\emptyset$-definable ones. In the case of DLO, we proved above that naming constants is not enough, which means that we also need to change the class of $\mathfrak{U}$-definable ones.[24]

Proposition 5.5.21. For every, possibly incomplete, L-theory T, there are $L' \supseteq L$ with $|L'| = |L|$ and a possibly incomplete L'-theory $T' \supseteq T$ with definable Skolem functions such that every $M \vDash T$ can be expanded to $M' \vDash T'$.

The construction above[25] is called *skolemisation*.

Proof. Let $L_0 \coloneqq L$, and inductively, for every $L_i(\emptyset)$-formula $\varphi(x,y)$ with $|x| = 1$, add an $|y|$-ary function symbol f_φ to L_i, call the resulting language L_{i+1}, and let T_{i+1} be the union of T_i together with all axioms $\forall y\ ((\exists x\ \varphi(x,y)) \to \varphi(f_\varphi(y),y))$, for $\varphi(x,y)$ as above. Let $T' \coloneqq \bigcup_{i<\omega} T_i$. Given $M \vDash T$, we inductively expand M_i to an L_{i+1}-structure M_{i+1} by setting $f_\varphi(b)$ to be an arbitrary witness to $\exists x\ \varphi(x,b)$ if one exists, and as an arbitrary element otherwise. By repeating ω times we obtain the required expansion M' of M, and prove that T' (is consistent and) has the required properties. □

This construction allows us to build models of arbitrary theories with certain properties by first passing to a skolemisation. For example, skolemising may be used to deduce from Proposition 5.5.20 the following.

Corollary 5.5.22. For every theory T and every linear order $(I,<)$ there is $M \vDash T$ containing an indiscernible a^I such that, for every $f \in \mathrm{Aut}((I,<))$, there is $\tilde{f} \in \mathrm{Aut}(M)$ with $\tilde{f}(a^i) = a^{f(i)}$.

Proof. Apply Proposition 5.5.21 to T, obtaining $L' \supseteq L$ and $T' \supseteq T$ with definable Skolem functions. In a monster model of a completion of T', let a^I be L'-indiscernible, let M' be its Skolem hull, and let $M \coloneqq M' \upharpoonright L$. The conclusion follows from Proposition 5.5.20, by observing that

[24] Note that if some expansion by constants gives us definable Skolem functions, then $|L|$ constants will suffice: we only need finitely many constants for every formula over $\emptyset$.

[25] Or below, in the proof, if you prefer. Well, I wrote this in a footnote, which I guess makes it "above" again (or "in the next page", depending on the version of these notes).

1. every L'-indiscernible sequence is L-indiscernible, and
2. every L'-automorphism is an L-automorphism. $\square$

Let me stress this again: taking a skolemisation is not a "mostly harmless" expansion, like Morleyising[26] or naming constants[27], so we *really* need to take the reduct to L: the properties of $\mathrm{Th}(M')$ and $\mathrm{Th}(M)$ can be *very* different. If you prefer, skolemising M is a highly non-canonical construction; it depends on several choices, and M' is not even determined up to elementary equivalence: for example, if $\vDash \forall x, y\ \varphi(x, y) \to \psi(x, y)$, we may have one M' where $\forall y\ f_\varphi(y) = f_\psi(y)$ holds and one where it fails. So be aware that, if you skolemise something, you are playing with fire.

Now, playing with fire may be dangerous, but the world is full of barbecues and fire jugglers. We probably will not have the time to learn to use flaming bolas (but see [She90, Section VIII.2]), but let us at least grill something.

Theorem 5.5.23. For every $\kappa \geq |L|$, there is $M \vDash T$ with $|M| = \kappa$ and such that, for every $A \subseteq M$, the model M realises at most $|A| + |L|$ types over A.

Note that since $|L|$ is infinite by convention, by usual type-counting tricks[28] it does not matter whether with "types" we mean "1-types" or "n-types for every n".

Proof. Use Proposition 5.5.21 to skolemise T to T' and work in a monster model of a completion of T'. Let a^κ be an L'-indiscernible κ-sequence, and let M' be the Skolem hull of a^κ. Since $|L| = |L'|$, if we show that the conclusion holds for the L'-structure M', then it will *a fortiori* hold for $M \coloneqq M' \restriction L$, because every type in $L(A)$ can be completed to a type in $L'(A)$, hence there are at least as many $L'(A)$-types as there are $L(A)$-types.

Therefore, we may assume that T has definable Skolem functions. Take as M an Ehrenfeucht–Mostowski model with spine a^κ indexed on κ. Let $A \subseteq M = \mathrm{dcl}(a^\kappa)$. Because if $b \in \mathrm{dcl}(B)$ then there is a finite $B_0 \subseteq B$ with $b \in \mathrm{dcl}(B_0)$, there is $A' \subseteq a^\kappa$ with $A \subseteq \mathrm{dcl}(A')$ and $|A'| \leq |A|$. Up to deductive closure, types over B are the same as types over $\mathrm{dcl}(B)$, so it enough to prove the conclusion when A is included in the spine. Assume this is the case.

Claim 5.5.24. For every $n \in \omega$, there are at most $|A| + \aleph_0$ many n-types over A which are realised in the spine.

Proof of the Claim. Since A is included in the spine, for some $J \subseteq \kappa$ we may write $A = \{a^j \mid j \in J\}$. For fixed n, we need to count the possibilities for $\mathrm{tp}(a^{i_0}, \dots, a^{i_{n-1}}/A)$, where $i_0, \dots, i_{n-1} \in \kappa$. By indiscernibility, this is determined by which inequalities hold between the different i_k (finitely many choices), and by the (possibly degenerate) cuts of the i_k in J. Since κ is well-ordered, so is J. But the right part of a cut in a well-ordered set must either be empty or have a minimum, hence there are at most $|J| \cdot 2 + 1$ possibly degenerate cuts in J. It follows that there are at most $|A| + \aleph_0$ possibilities for $\mathrm{tp}(a^{i_0}, \dots, a^{i_{n-1}}/A)$. $\square_{\text{CLAIM}}$

[26]. . . , which anyway may change the notion of substructure,. . .

[27]. . . , which anyway may break things like ω-categoricity,. . .

[28]Cf. the proof of Proposition 4.4.5.

By construction, every element of M is of the form $f(a^{i_0},\dots,a^{i_{n-1}})$, for a suitable $\emptyset$-definable function f and $a^{i_0},\dots,a^{i_{n-1}}$ a finite tuple from the spine. By indiscernibility, $\mathrm{tp}(f(a^{i_0},\dots,a^{i_{n-1}})/A)$ only depends on f, for which there are at most $|L|$ choices, and on $\mathrm{tp}(a^{i_0},\dots,a^{i_{n-1}}/A)$, which by the Claim can only be chosen in $|A|+\aleph_0$ ways, and the conclusion follows. $\square$

While we are on the theme of indiscernible sequences, let us also talk about *totally indiscernible* sequences.

Definition 5.5.25. We call a^I *totally indiscernible* over A, or an *indiscernible set* over A iff every permutation of a^I is indiscernible over A.

Equivalently, an I-sequence is indiscernible over A if and only if the type over A of $a^{i_0},\dots,a^{i_n}$ only depends on the quantifier-free type $\mathrm{qftp}^{(I\restriction\{=\})}(i_0,\dots,i_n/\emptyset)$, where I is seen as a set with no extra structure (i.e. we forget the order on I).

Sometimes, the only totally indiscernible sequences are the constant ones. This happens for example in DLO.

Definition 5.5.26. A *partitioned formula* $\varphi(x;y)$ is formula $\varphi(x,y)$ together with an ordered partition of its free variables into two parts. We call the variables in the first part x *object* variables, those in the second part y *parameter* variables.

The notation is usually abused and we just say that $\varphi(x;y)$ is a formula. Observe that a partitioned formula can be though of as a family of subsets of $\mathfrak{U}^{|x|}$ parameterised (possibly with repetitions) by $\mathfrak{U}^{|y|}$. In other words, $\varphi(x;y)$ induces a *definable family* of definable sets $\{\varphi(x;b)\mid b\in\mathfrak{U}^{|y|}\}$.

Definition 5.5.27. A partitioned formula $\varphi(x;y)$ has the *order property* (OP) if and only if there are sequences $(a^i)_{i<\omega}$ in $\mathfrak{U}^{|x|}$ and $(b^j)_{j<\omega}$ in $\mathfrak{U}^{|y|}$ such that $\vDash\varphi(a^i;b^j)\iff i<j$. A theory has OP (or *is* OP) iff some partitioned formula has OP. If T does not have OP, we say that T is (or *has*) NOP.

Note that such $(a^i)_{i<\omega}$ and $(b^j)_{j<\omega}$ are not guaranteed to exist in *every model*. Nevertheless, if $k<\omega$, then the existence of $(a^i)_{i<k}$ and $(b^j)_{j<k}$ with similar properties is expressible by a sentence. Hence, whether $\varphi(x;y)$ has OP or not may be checked on an arbitrary model, provided that we check for every $k<\omega$, and not for ω directly. If you want to check for ω directly, you need to do so on an ω-saturated model.

Remark 5.5.28. If there is a partitioned formula with OP, then there is one over $\emptyset$: enlarge x or y, then append the needed parameters to each a_i and b_j.

Proposition 5.5.29. The following are equivalent.

1. T is NOP.
2. There are no $\varphi(x;y)$ with $|x|=|y|$ and $(c^k)_{k<\omega}$ in $\mathfrak{U}^{|x|}$ such that $\vDash\varphi(c^k;c^{k'})\iff k<k'$.
3. For every $n\in\omega$, every indiscernible sequence of n-tuples is totally indiscernible.

Proof. The implication $1 \Rightarrow 2$ is trivial, and for $2 \Rightarrow 1$, if $\psi(t; w)$ and $(a^i)_{i<\omega}$, $(b^j)_{j<\omega}$ witness OP, it is sufficient to take $x = t^0w^0$, $y = t^1w^1$, then set $\varphi(x; y) = \varphi(t^0w^0; t^1w^1) \coloneqq \psi(t^0; w^1)$ and $c^k \coloneqq a^kb^k$.

To prove $3 \Rightarrow 2$, suppose there are such $\varphi(x; y)$ and $(c^k)_{k<\omega}$. By the Standard Lemma, there is an indiscernible $(d^\ell)_{\ell<\omega}$ such that $\mathrm{EM}(c^{<\omega}/\emptyset) \subseteq \mathrm{EM}(d^{<\omega}/\emptyset)$. By construction, $\vDash \varphi(d^0; d^1) \wedge \neg\varphi(d^1; d^0)$, hence $d^{<\omega}$ is indiscernible but not totally indiscernible.

Let us finish by proving $2 \Rightarrow 3$. Let c^I be indiscernible but not totally indiscernible. By the Standard Lemma, we may assume $I = \omega$. This means that, for some bijection $f : \omega \to \omega$ and some formula $\psi(x^0, \dots, x^n)$ over $\emptyset$ we have

$$\vDash \psi(c^0, \dots, c^n) \wedge \neg\psi(c^{f(0)}, \dots, c^{f(n)})$$

Since $c^{<\omega}$ is indiscernible, by shifting $c^{f(0)}, \dots, c^{f(n)}$ backwards, we find $\sigma \in S_{n+1}$, that is, a permutation of $\{0, \dots, n\}$, such that

$$\vDash \psi(c^0, \dots, c^n) \wedge \neg\psi(c^{\sigma(0)}, \dots, c^{\sigma(n)})$$

Claim 5.5.30. By changing $\psi(x^0, \dots, x^n)$, we may assume that σ is a transposition of two consecutive elements.

Proof of the Claim. Every element of S_{n+1} can be written as a product of transpositions permuting two consecutive elements.[29] Write σ in this fashion, say as $\sigma = \delta_\ell \cdot \dots \cdot \delta_0$, and for $i \le \ell$ let $\sigma_i \coloneqq \delta_i \cdot \dots \cdot \delta_0$. By assumption, there exists i such that

$$\vDash \psi(c^0, \dots, c^n) \wedge \neg\psi(c^{\sigma_i(0)}, \dots, c^{\sigma_i(n)})$$

Let i be minimal with the property above. If $i = 0$, we are done. Otherwise, just permute the variables of ψ according to σ_{i-1}. $\square_{\text{CLAIM}}$

By the claim, we may assume that there is $r < n$ such that

$$\vDash \psi(c^0, \dots, c^n) \wedge \neg\psi(c^0, \dots, c^{r-1}, c^{r+1}, c^r, c^{r+2}, \dots, c^n)$$

We prove that $\varphi(x; y) \coloneqq \psi(c^0, \dots, c^{r-1}, x, y, c^{r+2}, \dots, c^n)$ has OP, which is enough by Remark 5.5.28. By the Standard Lemma, there is an indiscernible $\mathbb{Q}$-sequence $d^{\mathbb{Q}}$ with the same Ehrenfeucht–Mostowski type as $c^{<\omega}$ and, up to an automorphism of $\mathfrak{U}$, we may assume that for $i \in \omega$ we have $d^i = c^i$. To conclude, just choose your favourite increasing sequence $(j_m)_{m<\omega}$ in $(r-1, r+2) \cap \mathbb{Q}$, and observe that, by construction, $\varphi(d^{j_{m_0}}, d^{j_{m_1}}) \iff m_0 < m_1$. $\square$

Exercise 5.5.31. Prove the following.

1. The theory of infinite sets is NOP.
2. DLO has OP.
3. T_{rg} has OP.

You may want to go through the theories introduced so far and try get a feeling for which have OP and which do not. Don't worry if you don't see a quick way to prove that a certain theory is NOP: will see soon that this property has several characterisations.

[29] Start by moving $\sigma(n)$ to the end, one place at a time, by permuting consecutive elements. Then apply induction.

Chapter 6

Having few types

6.1 Counting types

We saw that, for countable theories, having few types over $\emptyset$ has very special consequences. In this chapter we will see that, if we count types over arbitrary sets, then there are "few" types —namely, the bare minimum— if and only if there is a reason for this, if and only if there are *several* reasons for this. In order to make type-counting easier, we introduce *local* type spaces.

Definition 6.1.1. Let $\varphi(x; y)$ be a partitioned formula.

1. $\varphi^*(y; x)$ is obtained from $\varphi(x; y)$ by reversing the order of the partition: the formula is the same, but y is the tuple of object variables.
2. An *instance* of $\varphi(x; y)$ is a formula of the form $\varphi(x; b)$.
3. $S_\varphi(A)$ is the space of *φ-types over A*: maximal consistent sets of instances of φ and $\neg\varphi$ with parameters from A.
4. If $\kappa \geq |L|$, we define

$$f_{T,\varphi}(\kappa) := \sup\{|S_\varphi(M)| \mid M \vDash T, |M| = \kappa\}$$
$$f_T(\kappa) := \sup\{|S_1(M)| \mid M \vDash T, |M| = \kappa\}$$

Remark 6.1.2. By usual counting tricks, in the definition of f_T, instead of just S_1, we may equivalently take all the S_n at once. On the other hand, the supremum *must* be taken over all models of size κ. Some of them may simply have not enough of the "right" parameters to make the size of type space grow.

As you probably expect, if $\varphi(x; y)$ is a partitioned formula, then natural restriction map $S_x(A) \to S_\varphi(A)$ is continuous. It is also possible to consider finite sets Δ of partitioned formulas, all with the same partition,[1] and to talk of Δ-types.[2] In fact, $S_x(A)$ may be written as the inverse limit of all the $S_\Delta(A)$ for Δ as above along these maps. But let's actually start counting.

[1] Note that, if needed, one may always add extra parameter variables which are not necessarily used in every formula of Δ.

[2] Although if you have at least 2 parameters, or 2 $\emptyset$-definable elements, you can code boolean combinations of instances of formulas from finite set with boolean combinations of instances of a single formula. The trick is adding parameters to do case distinctions, e.g. "$\theta(x; yt) := (t = 0 \wedge \varphi(x; y)) \vee (t = 1 \wedge \psi(x; y))$".

Lemma 6.1.3. Let $\kappa \geq |L|$. Then $\kappa \leq f_T(\kappa) \leq 2^\kappa$.

Proof. For the first inequality look at realised types. For the second one, note that each type over A yields, injectively, a function $L(A) \to \{0, 1\}$. ☐

Quite remarkably, it will turn out that if there are many types over arbitrarily large set, then the culprit *must* be some $\varphi(x; y)$ with the order property: in one direction, we will get a lower bound on the size of $S_x(A)$ by looking at $S_\varphi(A)$, which is the reason we introduced it. Intuitively, the reason is that linear orders may have many cuts, and each will give us a different type. To make this precise, we introduce the following function on infinite cardinals.

Definition 6.1.4. If κ is an infinite cardinal, we define

$$\operatorname{ded}\kappa \coloneqq \sup\{\lambda \mid \text{there is a linear order of size } \kappa \text{ with } \lambda \text{ cuts}\}$$

Remark 6.1.5.

1. Every cut $L \sqcup R$ is determined by L, so there are no more cuts than subsets, hence $\operatorname{ded}\kappa \leq 2^\kappa$.
2. Notoriously, $\mathbb{Q}$ is dense in $\mathbb{R}$, hence $\operatorname{ded}\aleph_0 = 2^{\aleph_0}$.
3. Since we may always append a copy of $\mathbb{Q}$ to an infinite linear order without changing its cardinality, $\operatorname{ded}\kappa$ is always at least $2^{\aleph_0}$.
4. We may equivalently define $\operatorname{ded}\kappa$ by just looking at DLO's, instead of all linear orders. To see this, if $(I, <)$ is a linear order of size κ, for every pair $a, b \in I \cup \{\pm\infty\}$ such that $(a, b) = \emptyset$, insert a copy of $\mathbb{Q}$ between a and b, obtaining a DLO $J \supseteq I$. We are inserting at most $\kappa \cdot \aleph_0$ new points, so $|J| = \kappa$. As for the number of cuts in J, note that we are at most introducing $2^{\aleph_0}$ new cuts in κ places, hence I and J have the same number of cuts.
5. We may equivalently define $\operatorname{ded}\kappa$ as

   $$\sup\{\lambda \mid \text{there is a linear order of size } \lambda \text{ with a dense subset of size } \kappa\}$$

   In fact, if I has size κ and λ cuts, we may first replace each point of I with a copy of $\mathbb{Q}$, obtaining $I' \vDash \mathsf{DLO}$ of the same size and the same number of cuts, and then filling each cut with one element returns a linear order of size λ in which I' is dense. Conversely, if J has size λ and $I \subseteq J$ is dense (of size κ), then different points of J have different cut in I.

Lemma 6.1.6. $\kappa < \operatorname{ded}\kappa$.

Proof. Let μ be minimum with $2^\mu > \kappa$. Look at the tree $2^{<\mu}$ with the lexicographic ordering, induced by the convention that $0 <$ undefined < 1. By assumption $|2^{<\mu}| \leq \kappa$. But every branch in $2^{<\mu}$ yields a different cut, and there are 2^μ branches. ☐

Here are some facts on $\operatorname{ded}\kappa$ that we will neither prove nor need, but you may find interesting. By [Mit72], if κ has uncountable cofinality, in a cardinal preserving forcing extension $\operatorname{ded}\kappa < 2^\kappa$, and by [CKS16] it is consistent to have $\operatorname{ded}\kappa < (\operatorname{ded}\kappa)^{\aleph_0}$ for certain κ. Moreover, by [CS16], for any κ, we have $2^\kappa \le \operatorname{ded}(\operatorname{ded}(\operatorname{ded}(\operatorname{ded}(\kappa))))$.[3]

Another thing we will not prove is that, by [Kei76], if $|L| = \aleph_0$, then $f_T(\kappa)$ can only be one of these:

$$\kappa \qquad \kappa + 2^{\aleph_0} \qquad \kappa^{\aleph_0} \qquad \operatorname{ded}\kappa \qquad (\operatorname{ded}\kappa)^{\aleph_0} \qquad 2^\kappa$$

If you haven't done this already, it is a good idea to go back through these notes, e.g. to Section 4.3, and to compute cardinalities of $S_1(M)$ in different theories. After which, I recommend you try to solve the following exercise.

Exercise 6.1.7. For each of the six functions above, find some T in a countable L having that function as f_T.

6.2 The order property

The order property OP, introduced in Definition 5.5.27, will play a crucial role in the whole chapter. We defined it using ω, but it is easy to see that we may produce similar patterns with other linear orders.

Remark 6.2.1. By compactness and saturation of $\mathfrak{U}$, if $\varphi(x;y)$ has OP and I is any small linear order then, for $i \in I$, there are $a_i \in \mathfrak{U}^{|x|}$ and $b_i \in \mathfrak{U}^{|y|}$ such that $\vDash \varphi(a_i;b_j) \iff i < j$.

NOP (partitioned) formulas are closed under several constructions.

Lemma 6.2.2. Let $\varphi(x;y)$ and $\psi(x;z)$ be NOP, where y and z are allowed to share variables. Then:

1. If $y = uv$ and $c \in \mathfrak{U}^{|v|}$ then $\varphi(x;uc)$ is NOP.
2. $\varphi^*(y;x)$ is NOP.
3. Boolean combinations of φ, ψ, partitioned as $\theta(x;yz)$, are NOP.

Proof. The first part follows very easily from the definitions, and the second one from applying Remark 6.2.1 to $\varphi(x;y)$ with I the (negative) integers. The fact that being NOP is preserved under taking negations is similarly proven, so it is enough to show that $\theta(x;yz) \coloneqq \varphi(x;y) \lor \psi(x;z)$ is NOP. Suppose θ has OP, witnessed by $(a_i)_{i<\omega}$ and $(b_ic_i)_{i<\omega}$ such that $\vDash \varphi(a_i;b_j) \lor \psi(a_i;c_j) \iff i < j$. Colour $\{i,j\} \in [\omega]^2$, with $i < j$, white if $\vDash \varphi(a_i;b_j)$ and black if $\vDash \psi(a_i;c_j)$. By Ramsey's Theorem there is an infinite $I \subseteq \omega$ such that, for all $i < j$ both in I, the colour of $\{i,j\}$ is always white or always black. In the first case φ has OP, in the second case ψ does. □

Proposition 6.2.3. If $\varphi(x;y)$ has OP and $\kappa \ge |L|$, then $f_{\varphi,T}(\kappa) \ge \operatorname{ded}\kappa$. In particular, if T has OP then $f_T(\kappa) \ge \operatorname{ded}\kappa$.

Proof. Choose $I \vDash \mathsf{DLO}$ of size κ, take $(a_i)_{i\in I}$, $(b_i)_{i\in I}$, given by Remark 6.2.1, contained in some $M \vDash T$ of size κ. For each cut $C = L \sqcup R$ in I, define

$$\Phi_C(x) = \{\neg\varphi(x;b_j) \mid j \in L\} \cup \{\varphi(x;b_j) \mid j \in R\}$$

[3] According to [Che21], "[Shelah's] other superpower is the ability to discover number 4 where it has absolutely no reason to be."

Since $I \vDash \mathsf{DLO}$, every finite subset of Φ_C is realised by some a_i. Therefore, Φ_C is consistent, hence can be completed to $p_C \in S_\varphi(M)$. But $C \mapsto p_C$ is injective: if $j \in L_C \setminus L_{C'}$ then p_C and $p_{C'}$ disagree on $\varphi(x, b_j)$. The "in particular" part follows by completing to elements of $S_x(M)$ and invoking Remark 6.1.2. □

Hence, as promised, if φ has the order property, then there are many φ-types. In fact, the converse is true, where "many" just means "more than the bare minimum". We will prove this by using the following combinatorial result.

Theorem 6.2.4 (Erdős–Makkai)**.** Suppose B is infinite, and let $\mathcal{F} \subseteq \mathscr{P}(B)$ be a family of subsets of B of size $|\mathcal{F}| > |B|$. For $i \in \omega$, there are $b_i \in B$ and $S_i \in \mathcal{F}$ such that, either

1. for all $i, j \in \omega$ we have $b_i \in S_j \iff j < i$, or
2. for all $i, j \in \omega$ we have $b_i \in S_j \iff i < j$.

Proof. Since there are at most $|B|$ pairs of finite subsets of B, we can build $\mathcal{F}' \subseteq \mathcal{F}$ with $|\mathcal{F}'| = |B|$ satisfying:

for all finite $B_0, B_1 \subseteq B$, if there is $S \in \mathcal{F}$ with $B_0 \subseteq S$ and $B_1 \subseteq S^\complement$, then there is such an S in $\mathcal{F}'$.

Since there are at most $|B|$ Boolean combinations of elements of $\mathcal{F}'$, there is $S_* \in \mathcal{F}$ which is not such a Boolean combination.

Build by induction $(b'_i)_{i<\omega}$ in S_*, $(b''_i)_{i<\omega}$ in $S_*^\complement$, and $(S_i)_{i<\omega}$ in $\mathcal{F}'$ such that, for all $n \in \omega$,

(a) $\{b'_0, \dots, b'_n\} \subseteq S_n$,

(b) $\{b''_0, \dots, b''_n\} \subseteq S_n^\complement$, and

(c) for all $i < n$ we have $b'_n \in S_i \iff b''_n \in S_i$.

The base step is trivial. For the induction step:

Claim 6.2.5. There are $b'_n \in S_*$ and $b''_n \in S_*^\complement$ such that for all $i < n$ we have $b'_n \in S_i \iff b''_n \in S_i$.

Proof of the Claim. Suppose not, and fix $b \in S_*$. For every $i < n$, define S_i^b to be S_i if $b \in S_i$ and $S_i^\complement$ otherwise. Let $S^b := \bigcap_{i<n} S_i^b$. If there is $c \in S^b \cap S_*^\complement$ we can take $b'_n := b$ and $b''_n := c$; since we are assuming these things do not exist, we have $S^b \subseteq S_*$. Hence $S_* = \bigcup_{b \in S_*} S^b$. But this union is finite, since there are only 2^n possibilities for S^b. So S_* is a Boolean combination of the S_i, against choice of S_*. □ CLAIM

This gives us b'_n, b''_n satisfying (c). By choice of $\mathcal{F}'$ there is $S_n \in \mathcal{F}'$ satisfying (a) and (b). By Ramsey's Theorem, up to passing to an infinite $I \subseteq \omega$, either:

1. for all $j < i$ we have $b'_i \in S_j$, or
2. for all $j < i$ we have $b'_i \notin S_j$.

In the first case, set $b_i := b''_i$ and obtain 1 from the conclusion. In the second case obtain 2 by setting $b_i := b'_{i+1}$. □

Corollary 6.2.6. If $|S_\varphi(B)| > |B|$ for some infinite B, then $\varphi(x;y)$ has OP.

Proof. By Erdős–Makkai applied to the family of subsets of B

$$\mathcal{F} \coloneqq \{\{b \in B \mid \varphi(a,b)\} \mid a \in \mathfrak{U}^{|x|}\}$$

which has the same size as $S_\varphi(B)$ because whether $\vDash \varphi(a,b)$ only depends on $\operatorname{tp}_\varphi(a/B)$. Depending on cases 1 or 2 in the conclusion of Erdős–Makkai, we get OP for either φ or φ^*, and conclude by Lemma 6.2.2. □

Definition 6.2.7. Let κ be an infinite cardinal. A theory is *κ-stable* iff, for all A with $|A| = \kappa$, we have $|S_1(A)| = \kappa$. A theory is *stable* iff it is κ-stable for some $\kappa \geq |L|$.

Corollary 6.2.8. A theory is stable if and only if it is NOP.

Proof. Left to right is Proposition 6.2.3. As for right to left, by Corollary 6.2.6, if T is NOP we have $f_{\varphi,T}(\kappa) = \kappa$. But every $p(x) \in S_x(A)$ is determined by the collection, of its restrictions to instances of the various $\varphi(x;y) \in L(\emptyset)$, that is, by the function mapping $\varphi(x;y) \mapsto p \upharpoonright \varphi$. Therefore, in a NOP theory we have $f_T(\kappa) \leq \kappa^{|L|}$. To conclude, choose your favourite $\kappa \geq |L|$ with the property that $\kappa^{|L|} = \kappa$, for example $2^{|L|}$. □

For this reason, iff $\varphi(x;y)$ is NOP, we will say that $\varphi(x;y)$ is *stable*, and we call $\varphi(x;y)$ *unstable* iff it has OP.

6.3 Local ranks

As promised, we have shown that there are many types if and only if there is a good reason for it. In fact, there are at least two more equivalently good reasons to have many types, to which the rest of the chapter is devoted.

In this section, we look at a rank which will give us a "quantitative" version of stability. The idea is the following. In the proof of Proposition 6.2.3, we obtained many types by following the branches of a tree. For example, in DLO, we can use instances of the formula $\varphi(x;y) \coloneqq x < y$, which clearly has OP, to build the tree in Figure 6.1. The children of each node partition their parent into two classes,[4] and we are able to complete branches to pairwise inconsistent partial φ-types. The idea behind the rank we are about to introduce is to measure the height of the tallest tree we can build this way.

Definition 6.3.1 (Shelah's local 2-rank)**.** Fix a partitioned formula $\varphi(x;y)$. We inductively define the rank of a small partial type[5] $\theta(x)$ as follows.

- $R_\varphi(\theta(x)) \geq 0$ iff $\theta(x)$ is consistent, and $R_\varphi(\theta(x)) = -\infty$ otherwise.
- $R_\varphi(\theta(x)) \geq n+1$ iff there is $b \in \mathfrak{U}^{|y|}$ with

$$R_\varphi(\theta(x) \wedge \varphi(x,b)) \geq n \text{ and } R_\varphi(\theta(x) \wedge \neg\varphi(x,b)) \geq n$$

- $R_\varphi(\theta(x)) = n$ iff $R_\varphi(\theta(x)) \geq n$ and $R_\varphi(\theta(x)) \not\geq n+1$. Iff for all $n \in \omega$ we have $R_\varphi(\theta(x)) \geq n$, we write $R_\varphi(\theta(x)) = \infty$.

[4] By taking conjunctions. Of course $x \geq 1/4$ does not imply $x < 1/2$.
[5] As usual, parameters are allowed.

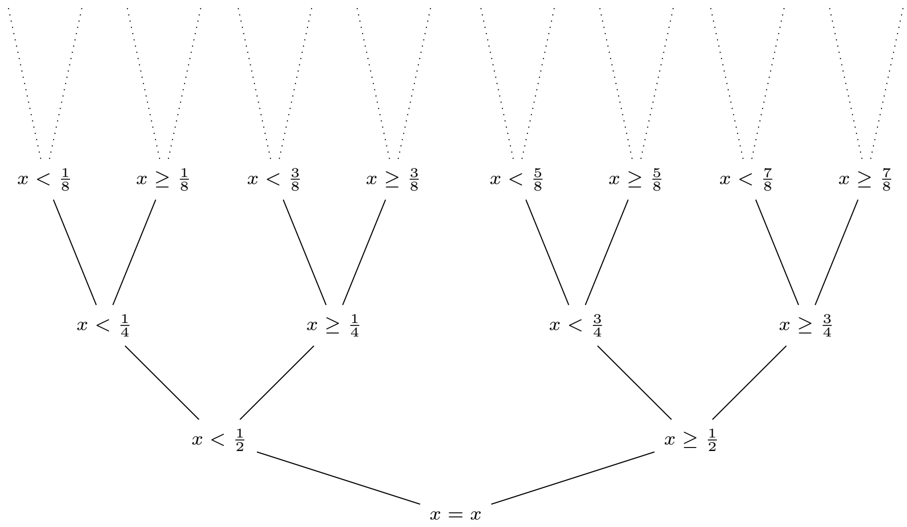

Figure 6.1: A binary tree of instances of $\varphi(x; y) \coloneqq x < y$ and of its negation.

The following exercise, besides being a good way to get familiar with the rank R_φ, is at the heart of the next section.

Exercise 6.3.2. If $\theta(x; y)$ is a formula, for all $n \in \omega$, the set $\{y \mid R_\varphi(\theta(x, y)) \geq n\}$ is definable.[6]

The fact that "R_φ gives us a quantitative version of stability" is made precise in the statement below.

Proposition 6.3.3. $\varphi(x; y)$ is stable if and only if $R_\varphi(x = x)$ is finite.

Proof. If φ is unstable, use Remark 6.2.1 with $I = [0, 1]$. So both $\varphi(x, b_{1/2})$ and $\neg\varphi(x, b_{1/2})$ contain densely many a_i. Keep splitting on the diadic rationals.

If $R_\varphi(x = x) = \infty$, then by compactness there is a tree of parameters $B = (b_\eta \mid \eta \in 2^{<\omega})$ such that for every $\eta \in 2^\omega$ this set is consistent:

$$\{\varphi(x; b_{\eta \restriction i}) \mid \eta(i) = 0\} \cup \{\neg\varphi(x; b_{\eta \restriction i}) \mid \eta(i) = 1\}$$

Complete each to an element of $S_\varphi(B)$, which therefore has size $> |B|$, then invoke Corollary 6.2.6. $\square$

6.4 Definable types

Definition 6.4.1. Let $\varphi(x; y)$ be a partitioned formula.

1. We say that $p(x) \in S_\varphi(B)$ is *A-definable* iff there is $\psi(y) \in L(A)$ such that, for all $b \in B$,
$$\varphi(x; b) \in p \iff \vDash \psi(b)$$
2. We say that $p(x) \in S_x(B)$ is *A-definable* iff every $p \restriction \varphi$ is. We say that it is *definable* iff it is B-definable.

[6] Hint: you just need to use induction and say that certain formulas are consistent.

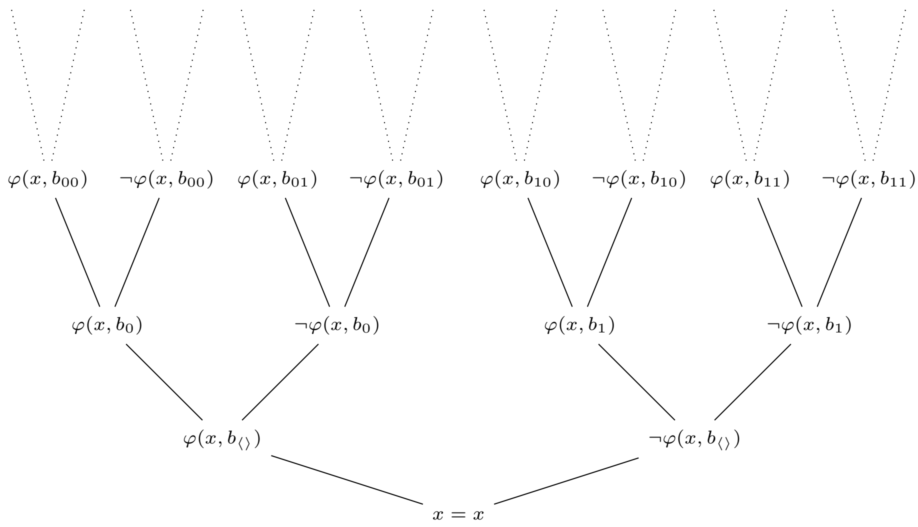

Figure 6.2: The same tree as in Figure 6.1, but with more general labels.

3. We say that *φ-types are uniformly definable* iff there is $\psi(y; z)$ such that: for every B with $|B| \geq 2$, for every $p \in S_\varphi(B)$, there is $c \in B$ such that p is defined by $\psi(y; c)$.

Example 6.4.2. In DLO, examples of definable types are $\mathrm{tp}(+\infty/\mathbb{Q})$ and $\mathrm{tp}(0^+/\mathbb{Q})$, while $\mathrm{tp}(\sqrt{2}/\mathbb{Q})$ is not definable.

Proposition 6.4.3. If $\varphi(x; y)$ is stable, then φ-types are uniformly definable.

Proof. Let $p \in S_\varphi(B)$. Define $p_0 = \emptyset$ and, inductively, if there is $p_{i+1} \subseteq p$ *obtained by adding only one formula to* p_i such that $R_\varphi(p_{i+1}) < R_\varphi(p_i)$, choose it. After at most $R_\varphi(x = x)$ steps, we have to stop, say at p_m. Because $R_\varphi(x = x)$ does not depend on p, modulo tricks (repeating parameters, case-distinctions done with parameters,[7] etc), such p_m may be written as instances of the same formula, uniformly across p. Use Exercise 6.3.2 to define

$$\psi(y) := \text{“}R_\varphi(p_m(x) \wedge \varphi(x, y)) = R_\varphi(p_m)\text{”}$$

Again, this $\psi(y)$ has parameters which depend on p, but besides that it is uniform in p. We show that $\psi(y)$ defines p.

- If $\varphi(x, b) \in p$, then $R_\varphi(p_m(x) \wedge \varphi(x, b)) = R_\varphi(p_m)$ by definition of p_m.
- If $\neg\varphi(x, b) \in p$, then $R_\varphi(p_m(x) \wedge \neg\varphi(x, b)) = R_\varphi(p_m(x))$ again by definition of p_m. But by definition of R_φ, then $p_m(x) \wedge \varphi(x, b)$ must have smaller rank, otherwise $R_\varphi(p_m)$ would go up. □

Let us put everything we know about stable formulas together.

Theorem 6.4.4. The following are equivalent for $\varphi(x; y)$.

[7] This uses $|B| \geq 2$; the trick is enlarging the tuple of parameter variables in order to write things like $(t = 0 \wedge (\ldots)) \vee (t = 1 \wedge (\ldots))$.

(a) φ is NOP.

(b) $R_\varphi(x = x) < \infty$.

(c) All φ-types are uniformly definable.

(d) All φ-types over models are definable.

(e) If $M \models T$ has size $\kappa \geq |L|$, then $|S_\varphi(M)| \leq \kappa$.

(f) There is $\kappa \geq |L|$ such that $f_{\varphi,T}(\kappa) < \operatorname{ded} \kappa$.

Proof. We already saw (a) $\Leftrightarrow$ (b) and (f) $\Rightarrow$ (a) $\Rightarrow$ (c). But (c) $\Rightarrow$ (d) is obvious and (e) $\Rightarrow$ (f) holds because $\kappa < \operatorname{ded} \kappa$, so we are left with (d) $\Rightarrow$ (e). Just count defining formulas: there are at most $\kappa + |L|$ of them. □

Before making a list of properties equivalent to stability of T, let us add another one to the list. It says that "all *externally definable* subsets are definable".

Definition 6.4.5. A set A is *stably embedded* iff, for all $n \geq 1$ and all $\mathfrak{U}$-definable sets $X \subseteq \mathfrak{U}^n$, there is an A-definable Y such that $X \cap A^n = Y \cap A^n$.

Remark 6.4.6. Let $X = \varphi(\mathfrak{U}, c) \subseteq \mathfrak{U}^n$. By spelling out definitions, we see that there is an A-definable Y such that $X \cap A^n = Y \cap A^n$ if and only if $\operatorname{tp}_\varphi(c/A)$ is definable. In particular, a theory is stable if and only if every set is stably embedded.

Theorem 6.4.7. Let T be a complete theory. The following are equivalent.

1. T is NOP.
2. There are no $(c_i \mid i < \omega)$ and φ with $\varphi(c_i, c_j) \iff i < j$.
3. Every indiscernible sequence is totally indiscernible.
4. All types over models are definable.
5. Every $A \subseteq \mathfrak{U}$ is stably embedded.
6. All formulas $\varphi(x; y)$ with x a single variable are stable.
7. $f_T(\kappa) \leq \kappa^{|L|}$.
8. $\exists \kappa \geq |L|\ f_T(\kappa) = \kappa$, that is, T is stable.
9. $\exists \kappa \geq |L|\ f_T(\kappa) < \operatorname{ded} \kappa$.

Proof. We have already seen 1 $\Leftrightarrow$ 2 $\Leftrightarrow$ 3 in Proposition 5.5.29, while 1 $\Leftrightarrow$ 4 follows from (a) $\Leftrightarrow$ (d) in the previous theorem, and 1 $\Leftrightarrow$ 5 is Remark 6.4.6. Moreover, 1 $\Rightarrow$ 6 is trivial, and 6 $\Rightarrow$ 7 is due to the fact that by Remark 6.1.2 it does not matter if we count types in one or several variables. For 7 $\Rightarrow$ 8 we take $\kappa = 2^{|L|}$. Finally, 8 $\Rightarrow$ 9 because $\kappa < \operatorname{ded} \kappa$, and 9 $\Rightarrow$ 1 by (a) $\Leftrightarrow$ (f) in the previous theorem. □

It is *not* enough to check definability of types on one model to get stability. For example, all types over the ordered field $\mathbb{R}$ are uniformly definable [MS94], but its theory is clearly unstable. Another instance of this phenomenon can be found [Del89] in the field $\mathbb{Q}_p$. Anyway, this is also true for the reduct $(\mathbb{R}, <)$, which you may prove as an exercise.

Exercise 6.4.8. In DLO, all φ-types over $\mathbb{R}$ are uniformly definable.

Chapter 7

Having very few models

7.1 Morley rank

Part of the idea behind Morley rank was hinted in Spoiler 4.3.11: the "simplest" types we can find are the isolated ones, to which we want to assign rank 0. Next, there are those types which are isolated amongst the nonisolated one, which will have rank 1. Inductively, the types of rank $n+1$ are the isolated ones amongst those of rank larger than n. In other words, we are looking at the *Cantor rank* of the points of type space.

Anyway, Morley rank is not exactly this.[1] The idea is that in $S_x(A)$ there may be types which are isolated "by mistake", that is, because A does not have "enough" parameters. For example, in ω-categorical countable theories, if A is finite then every element of $S_x(A)$ is isolated, yet we may have reasons to regard some types over A as being, in a sense, "less" isolated than other types: for instance, those realised in A will have a unique extension to any $B \supseteq A$, which will still be isolated, while other types will have at least one extension in $S_x(B)$ which is not isolated (see Example 7.1.5). To fix this, the original approach of Morley in [Mor65] was to look at preimages under all the restriction maps $S_x(B) \to S_x(A)$, and only consider a point to be "really" isolated if *all* of its extensions to B are isolated. Instead of doing this, we take advantage of the fact that we are working in a monster model. In fact, we will first define Morley rank for formulas, and obtain from it a definition for types.[2]

Definition 7.1.1. Let $\varphi(x) \in L(\mathfrak{U})$. Its *Morley rank* is either an ordinal, ∞, or $-\infty$, and is defined as follows.[3]

- $\mathrm{RM}(\varphi(x)) \geq 0$ iff $\varphi(x)$ is consistent, and $\mathrm{RM}(\varphi(x)) = -\infty$ otherwise.
- $\mathrm{RM}(\varphi(x)) \geq \alpha + 1$ iff there is a family $\{\psi_i(x) \mid i < \omega\}$ of pairwise inconsistent $L(\mathfrak{U})$-formulas, each of which implies $\varphi(x)$ and has $\mathrm{RM}(\psi_i(x)) \geq \alpha$.

[1] Otherwise, we would have just called it "Cantor rank", no?

[2] There are other possible ranks that one may put on type spaces, but not all of them make sense on formulas.

[3] For some reason, Morley rank tends to be denoted by RM instead of MR. I suspect that this is due to the abundance of literature on Morley rank written in French. If anyone has more precise information, my email address is at page vi.

- If λ is a limit ordinal, then $\mathrm{RM}(\varphi(x)) \geq \lambda$ iff for all $\alpha < \lambda$ we have $\mathrm{RM}(\varphi(x)) \geq \alpha$.
- $\mathrm{RM}(\varphi(x)) = \alpha$ iff $\mathrm{RM}(\varphi(x)) \geq \alpha$ and $\mathrm{RM}(\varphi(x)) \not\geq \alpha + 1$.
- Iff for all ordinals α we have $\mathrm{RM}(\varphi(x)) \geq \alpha$, we write $\mathrm{RM}(\varphi(x)) = \infty$.[4]
- Suppose that for some ordinal α we have $\mathrm{RM}(\varphi(x)) = \alpha$. We define the *Morley degree* $\mathrm{DM}(\varphi(x))$ as the maximum $n \in \omega$ such that there is a family $\{\psi_i(x) \mid i < n\}$ of pairwise inconsistent $L(\mathfrak{U})$-formulas, each of which has Morley rank α and implies $\varphi(x)$.
- If $p(x) \in S_x(A)$, its *Morley rank* is defined as

$$\mathrm{RM}(p(x)) \coloneqq \min\{\mathrm{RM}(\varphi(x)) \mid \varphi(x) \in L(A), p(x) \vdash \varphi(x)\}$$

- Suppose that for some ordinal α we have $\mathrm{RM}(p(x)) = \alpha$. The *Morley degree* of $p(x)$ is defined as

$$\mathrm{DM}(p(x)) \coloneqq \min\{\mathrm{DM}(\varphi(x)) \mid \varphi(x) \in L(A), p(x) \vdash \varphi(x), \mathrm{RM}(\varphi(x)) = \alpha\}$$

Note that, if $p(x) \in S_x(A)$, in order to compute its Morley rank, we need to look at formulas with parameters outside of A. All that is needed about $\mathfrak{U}$ in order for this to work is ω-saturation: that is, if $M \supseteq A$ is ω-saturated, then we may compute $\mathrm{RM}(p(x))$ by replacing $\mathfrak{U}$ by M in Definition 7.1.1.[5] This follows from the following exercises.

Exercise 7.1.2. Prove that, in the definition of $\mathrm{RM}(\varphi(x)) \geq \alpha + 1$, instead of requiring the existence of an infinite family of ψ_i, we may require the existence of arbitrarily large finite families with the same properties.

Exercise 7.1.3. Let $\varphi(x; y) \in L(\emptyset)$. Prove that, if $a, b \in \mathfrak{U}^{|y|}$ and $\mathrm{tp}(a/\emptyset) = \mathrm{tp}(b/\emptyset)$, then $\mathrm{RM}(\varphi(x; a)) = \mathrm{RM}(\varphi(x; b))$ and $\mathrm{DM}(\varphi(x; a)) = \mathrm{DM}(\varphi(x; b))$. In particular, if we change the parameters in a formula or type along an elementary map, Morley rank and degree do not change.

The definition of Morley rank we gave above may seem a bit removed from the introductory explanation in terms of isolated points. It is not. This is due to the following exercise, together with the fact that finite subspaces of Hausdorff spaces are discrete.

Exercise 7.1.4. If $\mathrm{RM}(\varphi(x) \vee \psi(x)) \geq \alpha$, then $\mathrm{RM}(\varphi(x)) \geq \alpha$ or $\mathrm{RM}(\psi(x)) \geq \alpha$.

Example 7.1.5. Let T be the theory of a generic equivalence relation; we looked at its types in Section 4.3.5. For every A, the realised types in $S_1(A)$ have Morley rank 0, generic types of equivalence classes represented in A have Morley rank 1, and the generic type over A has Morley rank 2. In this theory, we can also see that Morley rank may be different from Cantor rank. It is easy to see that T is ω-categorical, hence, if A is finite, then every point of $S_1(A)$ is isolated, hence has Cantor rank 0.

[4] Some authors prefer to say that, in this case, $\mathrm{RM}(\varphi(x))$ does not exist.

[5] It follows from this fact that, if M is ω-saturated, then the Morley rank of every $p \in S_x(M)$ equals its Cantor rank, and that having no type of rank ∞ over any A is the same as every $S_x(A)$ not containing a perfect set.

Example 7.1.6. Let T be the theory of an equivalence relation with exactly n equivalence classes, all infinite. Then the formula $x = x$, with $|x| = 1$ (hence the unique element of $S_1(\emptyset)$), has Morley rank 1 and Morley degree n.

Exercise 7.1.7. Let $|x| = 1$. For which cardinals κ, in the theory of κ generic equivalence relations, $\mathrm{RM}(x = x)$ is an ordinal?

Exercise 7.1.8. Let $|x| = 1$. Compute $\mathrm{RM}(x = x)$ in T_{2^ω}.

Exercise 7.1.9. Show that, in DLO, we have $\mathrm{RM}(x = x) = \infty$.

Proposition 7.1.10.

1. If $q(x) \supseteq p(x)$, then $\mathrm{RM}(q(x)) \le \mathrm{RM}(p(x))$, and if equality holds and $\mathrm{RM}(p(x))$ is an ordinal then $\mathrm{DM}(q(x)) \le \mathrm{DM}(p(x))$.
2. Let $p(x) \in S_x(A)$ and α an ordinal. If $\mathrm{RM}(p) = \alpha$, then there is a finite $A_0 \subseteq A$ such that $\mathrm{RM}(p \restriction A_0) = \alpha$.
3. If $p(x) \in S_x(A)$ has ordinal Morley rank and $B \supseteq A$, then there is $q(x) \in S_x(B)$ with $q(x) \supseteq p(x)$ and $\mathrm{RM}(p(x)) = \mathrm{RM}(q(x))$.

Proof. The first point is immediate from the definition of Morley rank. For the second one, if $\mathrm{RM}(p(x)) = \alpha$, then this is witnessed by a formula $\varphi(x) \in L(A)$. If A_0 is the set of parameters appearing in $\varphi(x)$, clearly $\mathrm{RM}(p \restriction A_0) \le \alpha$. But $(p \restriction A_0) \subseteq p$, so we conclude by the previous point. The last point will be proven once we show that this set of formulas is consistent

$$p(x) \cup \{\neg\varphi(x) \mid \varphi(x) \in L(B), \mathrm{RM}(\varphi(x)) < \mathrm{RM}(p(x))\}$$

If not, then by compactness there are $\psi(x) \in p(x)$ and $\varphi_i(x) \in L(B)$ with $\mathrm{RM}(\varphi_i(x)) < \mathrm{RM}(p(x))$ such that $\psi(x) \vdash \bigvee_{i<n} \varphi_i(x)$. But $\mathrm{RM}(\psi(x)) \ge \mathrm{RM}(p(x))$ by definition, hence by Exercise 7.1.4 there is $i < n$ such that $\mathrm{RM}\,\varphi_i(x) \ge \mathrm{RM}\,p(x)$, a contradiction. □

Observe that, contrary to what happens in the first point of the proposition above, if $p(x) \in S_x(A)$, $q(x, y) \in S_{xy}(A)$, and $p(x) \subseteq q(x, y)$, then $\mathrm{RM}(q) \ge \mathrm{RM}(p)$. The point is that when we regard the formulas of p as formulas in (x, y), we are working in a larger space, c.f. Remark 0.2.12.

Exercise 7.1.11. Let α be an ordinal. Then $\mathrm{RM}(\varphi(x)) = \alpha$ if and only if the set below is finite and nonempty.

$$\{p(x) \in [\varphi(x)] \subseteq S_x(\mathfrak{U}) \mid \mathrm{RM}(p(x)) \ge \alpha\}$$

Moreover, if this is the case, then the cardinality of the set above equals $\mathrm{DM}(p)$.

Lemma 7.1.12. If the Morley rank of a formula, or type, is at least $(|L| + |S_{<\omega}(\emptyset)|)^+$, then it is automatically ∞.

Proof. By Exercise 7.1.3, the number of possible ranks is at most $|L| + |S_{<\omega}(\emptyset)|$. But the definition of Morley rank, together with a tiny bit of transfinite induction, shows that if $\alpha < \beta$ are ordinals and there is a formula of Morley rank β, then there is one of Morley rank α. Therefore, there is no gap in the possible ordinal Morley ranks, the conclusion for formulas follows, and so does that for types. □

Here is a standard application of Morley rank; we will see some of its consequences in Chapter 8.

Remark 7.1.13. Suppose that G is a definable group, that is, a definable set together with definable functions $\cdot\colon G^2 \to G$ and ${}^{-1}\colon G \to G$ making it into a group. If $\mathrm{RM}(G)$ is an ordinal, then G has the descending chain condition on definable subgroups. In fact, if $H < G$ has infinite index, by looking at the cosets of H we realise that H must have lower Morley rank than G. Similarly, if H has finite index, it must have lower Morley degree. An infinite descending chain of proper subgroups would therefore yield an infinite descending sequence of pairs (α, n), with α an ordinal and $n \in \omega$, ordered lexicographically, which is utter nonsense.

7.2 Totally transcendental theories and prime models

We will see in this section that theories where all types have ordinal Morley rank are *very* stable, and have prime models over every set. So they clearly have a right to a special name.

Definition 7.2.1. We call T *totally transcendental* iff there is no $\varphi(x)$ with $\mathrm{RM}(\varphi(x)) = \infty$.

Equivalently, T is totally transcendental if and only if every type has ordinal Morley rank.

Theorem 7.2.2. Let T be totally transcendental. The following hold.

1. T is κ-stable for every $\kappa \geq |L|$.
2. If $p(x) \in [\varphi(x)]$ is of minimal Morley rank amongst the points of $[\varphi(x)]$, then $p(x)$ is isolated. In particular, over every A the isolated types are dense, hence T has prime models over every set.

Proof. Fix $p(x) \in S_x(A)$. By assumption, there are $\varphi(x) \in p(x)$ and an ordinal α such that $\mathrm{RM}(p(x)) = \mathrm{RM}(\varphi(x)) = \alpha$. By Exercise 7.1.11, there are only finitely many types in $S_x(\mathfrak{U})$ of rank α containing $\varphi(x)$, hence, by taking restrictions, there are only finitely many types in $S_x(A)$ of rank α containing $\varphi(x)$, say $p(x) = p_0(x), \dots, p_m(x)$. Finite subspaces of Hausdorff spaces are discrete, so we can find $\psi_p(x)$ which implies $\varphi(x)$ and such that $p(x)$ is the only element of $[\psi_p(x)]$ of rank α. By definition of $\mathrm{RM}(p)$, we must have $\mathrm{RM}(\psi_p(x)) = \alpha$. By construction, using the fact that the types in $[\psi_p]$ have rank at most α, the map $p(x) \mapsto \psi_p(x)$ is injective, but $\psi_p(x)$ is an $L(A)$-formula, which can only be chosen in $|L| + |A|$ ways, proving the first part of the conclusion.

For the second one, let $p(x) \in [\varphi(x)]$ have minimal Morley rank amongst the points of $[\varphi(x)] \subseteq S_x(A)$. Since T is totally transcendental, the Morley rank of $\varphi(x)$, hence also that of $p(x)$, is an ordinal, so we may consider the formula $\psi_p(x)$ constructed above. Therefore, $p(x)$ is the only element of $[\psi_p(x)]$ of rank α. Because α was the minimal rank of points of $[\varphi(x)]$, and $[\psi_p(x)] \subseteq [\varphi(x)]$, it follows that $[\psi_p(x)]$ isolates $p(x)$, and we conclude by Theorem 5.3.11. $\square$

Corollary 7.2.3. T is totally transcendental if and only if, whenever $L_0 \subseteq L$ is countable, the restriction of T to L_0 is ω-stable. In particular, if $|L| = \aleph_0$, the following are equivalent.

1. T is totally transcendental.
2. T is κ-stable for all κ.
3. T is ω-stable.

Proof. The only thing we still need to prove is right to left: since the Morley rank of a formula can only go down if we pass to a reduct, Theorem 7.2.2 will then supply us with left to right, and with the "in particular" statement.

By Lemma 7.1.12 there is an ordinal α such that

$$\mathrm{RM}(\varphi(x)) \geq \alpha \Longrightarrow \mathrm{RM}(\varphi(x)) \geq \alpha + 1 \tag{7.1}$$

If T is not totally transcendental, then there is some $\varphi(x)$ with $\mathrm{RM}(\varphi(x)) = \infty$, and (7.1) allows us to carry out the following construction. Because $\mathrm{RM}(\varphi(x)) \geq \alpha + 1$, we can find $\varphi_0(x)$ and $\varphi_1(x)$, both of Morley rank $\geq \alpha$, both implying $\varphi(x)$, and with $\varphi(x)_0 \wedge \varphi_1(x)$ inconsistent. But by (7.1), we can iterate this construction, building a binary tree akin to that of Figure 6.2, except it will be labelled with formulas which are not necessarily all instances of the same partitioned formula. The set A of parameters appearing in these formulas is clearly countable, and so is the sublanguage $L_0 \subseteq L$ consisting of the symbols appearing in the tree. As all complete binary trees of infinite height worth their salt, our tree has $2^{\aleph_0}$ branches, which we may complete to pairwise distinct elements of $S_x^{L_0}(A)$, proving that $T \upharpoonright L_0$ is not ω-stable. $\square$

Definition 7.2.4. Let β be an ordinal, and $(A_i)_{i<\beta}$ be a chain of subsets indexed on β, that is, if $i < j < \beta$ then $A_i \subseteq A_j$. We call the chain *continuous* iff, whenever $\lambda < \beta$ is limit, we have $A_\lambda \coloneqq \bigcup_{i<\lambda} A_i$.

Lemma 7.2.5. Let T be totally transcendental and $(A_i)_{i<\beta}$ be a chain.

1. Suppose that $p_i \in S_x(A_i)$, and that if $i < j$ then $p_i \subseteq p_j$. Then, the Morley rank and degree of p_i are eventually constant. Moreover, $\bigcup_{i<\beta} p_i(x) \in S_x(\bigcup_{i<\beta} A_i)$ has Morley rank and degree equal to eventual rank and degree of the p_i.
2. If $p_0(x) \in S_x(A_0)$ is isolated, then there are isolated $p_i \in S_x(A_i)$ such that if $i < j$ then $p_i \subseteq p_j$.

Proof. The first part is immediate from Proposition 7.1.10. For the second part, we inductively build the sequence of the p_i, ensuring that p_i has minimal Morley rank amongst the types that, for every $j < i$, restrict to p_j, and show that such a sequence works. If $\pi \colon S_x(A_{i+1}) \to S_x(A_i)$ is the natural projection, since p_i is isolated $\pi^{-1}(\{p_i\})$ is open nonempty; we take as p_{i+1} a type of minimal Morley rank in $\pi^{-1}(\{p_i\})$, which is isolated by Theorem 7.2.2. For the limit step, up to adding some extra A_i, we may assume our chain to be continuous. Moreover, by the first part, up to trimming the sequence we may assume that the Morley rank of the p_i built so far is constantly α. Let $p_\lambda \coloneqq \bigcup_{i<\lambda} p_i$, and let $\pi \colon S_x(A_\lambda) \to S_x(A_i)$ be the natural projection. Again, each $\pi^{-1}(\{p_i\})$ is open, and it contains p_λ by definition, so by Theorem 7.2.2 it is enough to show that p_λ is of minimal Morley rank in $\pi^{-1}(\{p_i\})$. If not, then there is some nonempty $[\psi(x)] \subseteq \pi^{-1}(\{p_i\})$ with $\mathrm{RM}(\psi(x)) < \alpha$. Let j be such that $i < j < \lambda$ and $\psi(x) \in L(A_j)$. If $\pi_j \colon S_x(A_j) \to S_x(A_i)$ is the natural projection, then, regarding now $[\psi(x)]$ as a subset of $S_x(A_j)$, we have $[\psi(x)] \subseteq \pi_j^{-1}(\{p_i\})$. This contradicts minimality of the Morley rank of p_j. □

Proposition 7.2.6. Let T be totally transcendental, and $(A_i)_{i<\beta}$ be a continuous chain. Then there is a continuous chain $(M_i)_{i<\beta}$ of models of T such that each M_i is prime over A_i.

Proof. We build $M_i \supseteq A_i$ by induction on i, ensuring the following property: for every $i < j$, and every model N, every elementary map $A_j M_i \to N$ can be extended to an elementary map $A_j M_{i+1} \to N$. Using a "only forth" argument, together with continuity of the chain $(M_i)_{i<\beta}$, this ensures that each M_j is indeed prime over A_j.

Start with a model M_0 prime over A_0, which exists by Theorem 7.2.2. If we want continuity, we have no choice but to take unions at limit stages, so we only need to take care successor steps.

List the isolated points of $S_1(A_{i+1}M_i)$ as $(p_k \mid k < \mu)$, for a suitable cardinal μ, and use Lemma 7.2.5 to find a chain $(p_{0,\ell} \mid i < \ell < \beta)$ starting with $p_{0,i+1} = p_0$ and made of isolated types $p_{0,\ell} \in S_1(A_\ell M_i)$. We set $q_0 \coloneqq \bigcup_{\ell<\beta} p_{0,\ell}$, choose $a_0 \vDash q_0$, and inductively build $(a_k \mid k < \mu)$ such that

$$\text{for every } j \text{ with } i < j < \beta \text{ the type } \operatorname{tp}(a_k / A_j M_i a_{<k}) \text{ is isolated} \tag{7.2}$$

In order to do this, inductively, extend p_k to an isolated $p'_k \in S_1(A_{i+1}M_i a_{<k})$, similarly as in the proof of Theorem 5.3.11, then find q_k obtained from p'_k via Lemma 7.2.5 similarly as q_0 was obtained from p_0, that is, as a union of a chain of isolated types $p'_{k,\ell} \in S_1(A_\ell M_i a_{<k})$. Finally, choose $a_k \vDash q_k$. Again as in the proof of Theorem 5.3.11, we iterate this ω times, building $(a_k \mid k < \mu \cdot \omega)$ which still satisfies (7.2), and such that $M_{i+1} \coloneqq M_i\{a_k \mid k < \mu \cdot \omega\}$ is the required model. □

7.3 Countable ω-stable theories

The results in this section are the last technical steps towards proving Morley's Theorem. Hence, from now on we will work in a countable L. By Corollary 7.2.3, we may then say "ω-stable" instead of "totally transcendental".

Theorem 7.3.1. Let L be countable and T be ω-stable. Suppose that $A \subseteq C$, that $\kappa \coloneqq |C|$ is a regular uncountable cardinal, and that $|A| < \kappa$. Then C contains a totally A-indiscernible nonconstant sequence of length κ.

Proof. Let $\mathcal{F}$ be the family of all pairs (B, p) such that

1. $A \subseteq B \subseteq C$,
2. $|B| < \kappa$,
3. $p \in S_1(B)$, and
4. p has κ realisations in C.

By Corollary 7.2.3 $|B| < \kappa$ implies $|S_1(B)| < \kappa$, and since $\kappa > |B|$ is regular there must be κ elements of C with the same type over B, proving that $\mathcal{F}$ is nonempty. Let (B_0, p_0) be an element of $\mathcal{F}$ such that the pair $(\mathrm{RM}(p_0), \mathrm{DM}(p_0))$ is minimal in the lexicographical order.

Claim 7.3.2. For every B with $|B| < \kappa$ and $B_0 \subseteq B \subseteq C$, there is a unique $p_B \in S_1(B)$ such that

- $p_B \supseteq p_0$, and
- $(\mathrm{RM}(p_B), \mathrm{DM}(p_B)) = (\mathrm{RM}(p_0), \mathrm{DM}(p_0))$.

Moreover this p_B satisfies $(B, p_B) \in \mathcal{F}$.

Proof of the Claim. Let $(r, n) \coloneqq (\mathrm{RM}(p_0), \mathrm{DM}(p_0))$. Consider the closed subset $[p_0] \subseteq S_1(B)$. Again by Corollary 7.2.3, it has size $< \kappa$, hence by regularity of κ it must contain some p_B which is realised κ times in C, so $(B, p_B) \in \mathcal{F}$. By definition, we have $p_B \supseteq p_0$, and we are left to show that p_B is the unique $p \in [p_0]$ with $(\mathrm{RM}(p), \mathrm{DM}(p)) = (r, n)$. Note that $\le$ holds for every $p \in [p_0]$ by Proposition 7.1.10, and in particular for p_B. But $(B, p_B) \in \mathcal{F}$, hence by minimality of $(\mathrm{RM}(p_0), \mathrm{DM}(p_0))$ we must have equality. Fix $\varphi(x) \in p_0(x)$ of rank and degree (r, n), and suppose there is another $p \neq p_B$ in $[p_0]$ with rank and degree (r, n). Since $p \neq p_B$, there are $L(B)$-formulas separating them, which, up to taking conjunctions, we may assume to imply $\varphi(x)$ and to have rank and degree (r, n). This implies that $\mathrm{DM}(\varphi(x)) \ge 2n$, a contradiction. $\square_{\text{CLAIM}}$

Now build a κ-sequence a_κ as follows. Start by choosing $a_0 \in C$ such that $a_0 \vDash p_0$. Inductively, define $B_i \coloneqq B_0 \cup \{a_j \mid j < i\}$, and set p_i to be the type p_{B_i} given by the claim. Since $(B_i, p_i) \in \mathcal{F}$, there are κ realisations of p_i in C; set a_i to be any such realisation not in $B_0 a_{<i}$.

If we prove that a_κ is B_0-indiscernible, we are done: it will in particular be A-indiscernible, and totally so by stability and Proposition 5.5.29. Hence we prove by induction on n that if $i_0 < \ldots < i_n$ and $j_0 < \ldots < j_n$ are ordinals in κ, then $a_{i_0}, \ldots, a_{i_n} \equiv_{B_0} a_{j_0}, \ldots, a_{j_n}$. All elements of a_κ have the same 1-type, namely p_0, giving us the case $n = 0$ of the induction. Moreover, if $i < j$, since $p_j \supseteq (p_j \upharpoonright B_i) \supseteq p_0$, we have that $(p_j \upharpoonright B_i)$ still has rank and degree equal to (r, n), by the "uniqueness" part in the claim it must equal p_i; in other words, $p_j \supseteq p_i$. Set $B \coloneqq B_0 a_{i_0}, \ldots, a_{i_n}$ and $B' \coloneqq B_0 a_{j_0}, \ldots, a_{j_n}$. By construction, $a_{i_{n+1}}$ (respectively, $a_{j_{n+1}}$) realises the unique p_B (respectively, $p_{B'}$) given by the claim. By inductive hypothesis, the map $f \colon B \to B'$ fixing B_0 pointwise and sending a_{i_ℓ} to a_{j_ℓ} is elementary. If we change the parameters in p_B according to f, by Exercise 7.1.3 we obtain a type over B' of the same Morley rank and degree, which still extends p_0 because $f \upharpoonright B_0 = \mathrm{id}_{B_0}$; by the Claim, this type must be $p_{B'}$, and we are done. $\square$

Corollary 7.3.3. If we are in the assumptions of Theorem 7.3.1, except κ is possibly not regular, then for every $\mu < \kappa$ the set C contains a totally A-indiscernible nonconstant sequence of length μ.

Proof. Successor infinite cardinals are regular and cofinal in every singular cardinal. $\square$

Corollary 7.3.4. Let L be countable and T be ω-stable. Suppose there is $N \vDash T$ with $|N| > \aleph_0$ which is not saturated. Then there are $M \preceq N$ with $|M| = \aleph_0$ and $A \subseteq M$ such that

1. $M \setminus A$ contains a (totally) A-indiscernible nonconstant sequence a^ω, and
2. some $q(x) \in S_1(A)$ is omitted in M.

Proof. If N is not saturated, there must be $B \subseteq N$ with $|B| < |N|$ and some $p(x) \in S_1(B)$ omitted in N. Use Corollary 7.3.3 to get a B-indiscernible nonconstant sequence a^ω, indexed on ω, in $N \setminus B$. By Löwenheim–Skolem there is a countable M_0 with $a^\omega \subseteq M_0 \preceq N$. Clearly, M_0 still omits $p(x)$, hence for every $m \in M_0$ there is $\varphi_m(x) \in p(x)$ such that $m \vDash \neg\varphi_m(x)$. If we collect all the parameters of the $\varphi_m(x)$ in a (countable) subset A_1 of B, then by construction no $m \in M_0$ realises $p \upharpoonright A_1$. Take a countable $M_1 \preceq N$ with $M_1 \supseteq M_0 A_1$. Repeat this construction ω times, obtaining chains $(M_i)_{i<\omega}$ of models and $(A_i)_{i<\omega}$ of sets such that $A_i \subseteq M_i \cap B$ and M_i omits $p \upharpoonright A_{i+1}$. Take $M \coloneqq \bigcup_{i<\omega} M_i$ and $A \coloneqq \bigcup_{i<\omega} A_i$. By construction $q(x) \coloneqq p(x) \upharpoonright A$ is omitted in M. Moreover, $a^\omega \subseteq M_0 \subseteq M$; since $A \subseteq B$, and a^ω is B-indiscernible, it is in particular A-indiscernible, and by construction $a^\omega \cap A \subseteq a^\omega \cap B = \emptyset$. $\square$

Proposition 7.3.5. If $|L| = \aleph_0$ and T is ω-stable, then for every $\kappa > \aleph_0$ there is an $\aleph_1$-saturated model of size κ.

Proof. By Corollary 7.2.3 T is κ-stable. Hence we can start with any M_0 of size κ, and build a continuous elementary chain $(M_i \mid i < \omega_1)$ such that each M_{i+1} has size κ and realises all types over M_i. Then we just take the union of this chain, and use regularity of ω_1 to prove $\aleph_1$-saturation. $\square$

7.4 Morley's Theorem

We put the pieces together and prove that if $|L| = \aleph_0$, then a theory is κ-categorical for every uncountable κ if and only if it is κ-categorical for some uncountable κ. Of course, the fact that we introduced Morley rank in this chapter is not a coincidence.

Proposition 7.4.1. If $|L| = \aleph_0 < \kappa$ and T is κ-categorical, then T is ω-stable.

Proof. Otherwise there is a countable A such that $S_1(A)$ is uncountable, hence we may find $B \supseteq A$ of size $\aleph_1$ whose elements have pairwise distinct types over A. By Löwenheim–Skolem, there is $N \supseteq B$ of size κ. On the other hand, applying Theorem 5.5.23 to T_A gives us an $M \supseteq A$ of size κ which realises at most $\aleph_0$ types over A, and which cannot therefore be isomorphic to N. $\square$

Theorem 7.4.2. Suppose that $|L| = \aleph_0$ and T is ω-stable. If there is an uncountable model which is not saturated, then for every $\kappa > \aleph_0$ there is a model of size κ which is not $\aleph_1$-saturated.

Proof. Let M, A, a^ω and $q(x) \in S_1(A)$ be given by Corollary 7.3.4. Using the Standard Lemma and an automorphism in $\mathrm{Aut}(\mathfrak{U}/A)$, we may extend a^ω to a (totally) A-indiscernible sequence a^κ. For each $i < \kappa$, set $A_i \coloneqq Aa_{<i}$. Clearly, the A_i form a continuous chain, hence by Proposition 7.2.6 there is a

continuous chain $(M_i)_{i<\kappa}$ of models of T such that each M_i is prime over A_i. Since $M_\kappa \coloneqq \bigcup_{i<\kappa} M_i$ has size κ, if we prove that M_κ omits $q(x)$, which is over the countable set A, then it cannot be $\aleph_1$-saturated.

Of course, it is enough to prove by induction on i that every M_i omits $q(x)$. Recall that M_i is prime over $A_i = Aa_{<i}$. For $i < \omega$, since $M \supseteq Aa_{<i}$ omits q, so does M_i, by primality. At limit stages the conclusion is obvious from the inductive hypothesis, so we only need to show that if $i \geq \omega$ and M_i omits q then so does M_{i+1}. Since a^κ is *totally* A-indiscernible, and $i \geq \omega$, we may extend id_A to an elementary map sending $a_{\leq i}$ to $a_{<i}$. This induces an embedding of M_{i+1} into M_i, hence, if M_{i+1} realises $q(x)$, then so does M_i. □

Theorem 7.4.3. Let L be countable and $\kappa > \aleph_0$. If T is κ-categorical, then every uncountable model of T is saturated. In particular, T is κ-categorical for some uncountable κ if and only if T is κ-categorical for every uncountable κ.

Proof. By Proposition 7.4.1 T is ω-stable, hence if the conclusion fails, by Theorem 7.4.2 there is model of size κ which is not $\aleph_1$-saturated. By Proposition 7.3.5 T has an $\aleph_1$-saturated model of size κ, contradicting κ-categoricity. The "in particular" part is then immediate from Theorem 4.5.4. □

Chapter 8

A taste of definable groups

Fix, as usual, a complete T and a monster $\mathfrak{U} \vDash T$. As we saw in Remark 7.1.13, a definable group is nothing but a definable set G, together with definable functions $\cdot\colon G^2 \to G$ and ${}^{-1}\colon G \to G$ making it into a group. For example, if T is the complete theory of a field K, then usual matrix groups such as GL_n, SL_n, etc are definable groups.

In this chapter, we will develop some basics of the theory of definable groups in ω-stable theories, and use these techniques to prove Macintyre's Theorem, that the only infinite totally transcendental fields are the algebraically closed ones. In order to do this, we will need to develop some further model-theoretic tools. But first, let us see what we can immediately harvest from the descending chain condition.

8.1 Consequences of the descending chain condition

By Remark 7.1.13, if G is a definable group of ordinal Morley rank, then it satisfies the DCC on definable subgroups. This has a lot of consequences.

Proposition 8.1.1. Let G be a definable group with the DCC on definable subgroups. The following facts hold.

1. Every definable injective homomorphism $G \to G$ is surjective.

2. If $\{H_i \mid i \in I\}$ is a family of definable subgroups, then there is a finite $I_0 \subseteq I$ such that $\bigcap_{i\in I} H_i = \bigcap_{i\in I_0} H_i$.

3. The centraliser of any (not necessarily definable) $A \subseteq G(\mathfrak{U})$ is definable.

Proof.

1. Any counterexample $f\colon G \to G$ yields a violation of the DCC by considering $G \supsetneq f(G) \supsetneq f^2(G) \supsetneq \dots$

2. Otherwise there is an infinite sequence $(i_n)_{n\in\omega}$ such that $\bigcap_{j<n} H_{i_j} \supsetneq \bigcap_{j<n+1} H_{i_j}$, again violating the DCC.

3. The centraliser of a single element a is definable by the formula $x\cdot a = a\cdot x$. Apply the previous point to the family of centralisers of elements of A. □

This has an important consequence on stabilisers of types under a certain natural action. Before stating it, let us introduce some commonly used notation.

Notation 8.1.2. If X is a definable set, say defined by $\varphi(x)$, we write $S_X(A)$ for the subspace $[\varphi(x)]$ of $S_x(A)$.

Definition 8.1.3. For $p \in S_G(M)$ and $g \in G(M)$, we define $g \cdot p \coloneqq \{\varphi(x) \in L(M) \mid \varphi(g \cdot x) \in p(x)\}$. The *stabiliser* of p is $\operatorname{Stab}(p) \coloneqq \{g \in G(M) \mid g \cdot p = p\}$.

In other words, $a \vDash p$ if and only if $g \cdot a \vDash g \cdot p$. Of course, this is the *left* stabiliser, and everything we are going to say also goes through, *mutatis mutandis*, for the *right* stabiliser, whose definition is left to the reader.

Remark 8.1.4. Recall that a type $p(x) \in S_x(A)$ is definable iff for every $\varphi(x; y)$ the set $\{b \in A \mid \varphi(x; b) \in p(x)\}$ equals the set of solutions in A of some $\psi(y) \in L(A)$. If $A = M$ is a model,[1] and $\psi_0(y), \psi_1(y) \in L(M)$ are as above, then ψ_0 and ψ_1 define the same subset of $M^{|y|}$, and since M is a model this implies $\vDash \forall y\ \psi_0(y) \leftrightarrow \psi_1(y)$.

Notation 8.1.5. If $A = M$ is a model and $p \in S_x(M)$ is a definable type, we denote such a $\psi(y)$ with $(d_p\varphi)(y)$.

Theorem 8.1.6. Let G be a definable group in a totally transcendental T. For every $p \in S_G(M)$, the stabiliser $\operatorname{Stab}(p)$ is definable.

Proof. Define

$$\operatorname{Stab}^{\varphi}(p) \coloneqq \big\{g \in G(M) \;\big|\; \forall h \in G(M)\ \big(\varphi(h \cdot x) \in p(x) \iff \varphi(h \cdot g \cdot x) \in p(x)\big)\big\}$$

It is easy to show that $\operatorname{Stab}(p) = \bigcap_{\varphi(x) \in p(x)} \operatorname{Stab}^{\varphi}(p)$; therefore, by total transcendence and Proposition 8.1.1, it is enough to show that every $\operatorname{Stab}^{\varphi}(p)$ is a definable subgroup. The proof that it is a subgroup is again easy[2]; as for definability, recall that totally transcendental theories are stable, hence the type $p(x)$ is definable. If $\psi(x; y)$ is the formula $\varphi(y \cdot x)$, consider $(d_p\psi)(y)$; by definition, if $h \in G(M)$, then $\vDash d_p\psi(h) \iff \varphi(h \cdot x) \in p(x)$. Therefore, $\operatorname{Stab}^{\varphi}(p)$ is defined by the formula $\theta(y) \coloneqq \forall z\ (d_p\psi(z) \leftrightarrow d_p\psi(z \cdot y))$. □

8.2 Interpretability

Some constructions, e.g. projective space,[3] are usually carried out by using quotients. When considering "definable quotients", we speak of *interpretable* sets: a set is interpretable iff it is the quotient of a definable set by a definable equivalence relation. We can also speak of interpretable structures.

Definition 8.2.1. Let M_0 be an L_0-structure and M_1 be an L_1-structure. We say that M_1 is *interpretable* in M_0 iff there are

1. some n and some L_0-definable $X \subseteq M_0^n$
2. an L_0-definable equivalence relation E on X

[1] Without this assumption, the conclusion is in general false.
[2] Hint: use that the definition requires something to happen for *all* $h \in G(M)$.
[3] ...which by the way, since we are talking about groups, is where elliptic curves live,...

3. for every $s \in L_1$, an L_0-definable X_s in some $M_0^{n \cdot m_s}$, for a suitable m_s,[4]

such that every X_s is E-equivariant and X/E, with the L_1-structure induced by the X_s, is isomorphic to M_1.

It is possible to lift this to the level of theories. It is also possible to define a structure, called M^{eq}, and best viewed as a structure in multi-sorted logic, such that interpretability in M is the same as definability in M^{eq}. But we do not have much time left, so I will refer you to the literature for that, and leave you with these two exercises we will need later.

Exercise 8.2.2. Let K be a field, and F an algebraic extension of K with $\dim(F/K)$ finite. Then F is interpretable in K.

Exercise 8.2.3. If $\mathrm{Th}(M)$ is totally transcendental (resp. stable) and M interprets N, then $\mathrm{Th}(N)$ is totally transcendental (resp. stable).

8.3 Some forking calculus in disguise

In a longer course, this chapter would have come only after another one developing the theory of *forking* in stable theories. Given the name of this chapter, forking will only be served as an appetiser, but you should be aware that it is a crucial tool in the analysis of stable theories and its applications, that it allows to define an *independence relation*, known as *nonforking independence*, and that in this section you are learning something about it, although only in special cases, and without even seeing its definition.[5]

Notation 8.3.1. Write $\mathrm{RM}(a/A)$ for $\mathrm{RM}(\mathrm{tp}(a/A))$.

Proposition 8.3.2. If $a \in \mathrm{acl}(Ab)$ then $\mathrm{RM}(ab/A) = \mathrm{RM}(b/A)$.

Proof. The inequality $\geq$ is easy to prove, and does not even need the assumption $a \in \mathrm{acl}(Ab)$. For the other inequality, let $\alpha := \mathrm{RM}(b/A)$, and start by choosing some $\varphi(x, y) \in \mathrm{tp}_{xy}(ab/A)$ witnessing $a \in \mathrm{acl}(Ab)$. Up to adding conjuncts to φ, we may further assume that

1. $\mathrm{RM}(\exists x\ \varphi(x, y)) = \alpha$ (just take a conjunction with a formula in y of rank α from $\mathrm{tp}(b/A)$), and
2. every $\varphi(x, b')$ has finitely many solutions (just take a conjunction with a suitable $\exists^{\leq n} t\ \varphi(t, y)$; this conjunction is still in $\mathrm{tp}(b/A)$, and it cannot lower the Morley rank above, which is already minimum by definition of Morley rank of a type).

We prove that $\mathrm{RM}(\varphi(x, y)) \leq \alpha$. By Exercise 7.1.11, it is enough to show that the following subset of $S_{xy}(\mathfrak{U})$ is finite

$$[\varphi(x, y)] \cap \bigcap_{\mathrm{RM}(\psi(x,y))<\alpha} [\neg\psi(x, y)]$$

[4] It depends on whether s is a constant, function, or relation symbol, and on its arity. I have not written precisely who m_s is, but it should be clear if you read the rest of the definition.

[5] Ok, I guess I should at least say that in a totally transcendental theory, if $A \subseteq B$, then $q \in S_x(B)$ is a nonforking extension of $p \in S_x(A)$ if and only if $q \supseteq p$ and $\mathrm{RM}(q) = \mathrm{RM}(p)$.

We prove finiteness of the larger set

$$[\Phi(x,y)] \coloneqq [\varphi(x,y)] \cap \bigcap_{\mathrm{RM}(\psi(y))<\alpha} [\neg\psi(y)]$$

Let $[\Psi(y)]$ be[6] the projection of $[\Phi(x,y)]$ to $S_y(\mathfrak{U})$, and note that $[\Psi(y)] \subseteq [\exists x\ \varphi(x,y)]$. By definition of $[\Phi(x,y)]$, the set $[\Psi(y)] \subseteq S_y(\mathfrak{U})$ only contains types of Morley rank at least α. Again by Exercise 7.1.11, the fact that $\mathrm{RM}(\exists x\ \varphi(x,y)) = \alpha$ implies that $[\Psi(y)]$ is finite. Moreover, since every $\varphi(x,b')$ is finite, so is $[\Psi(y)] \cap [\varphi(x,y)] \subseteq S_{xy}(\mathfrak{U})$. But this set contains $[\Phi(x,y)]$. $\square$

This has the following important consequence. Note that the case of definable bijections is almost immediate from the definition of Morley rank.

Exercise 8.3.3. If X, Y are definable sets and $f\colon X(\mathfrak{U}) \to Y(\mathfrak{U})$ is a definable finite-to-one[7] function, then $\mathrm{RM}(X) = \mathrm{RM}(Y)$.

By now you are probably convinced that a definable subset of $M^{|x|}$, say a $\emptyset$-definable one, is not "really" a subset of $M^{|x|}$, but rather the equivalence class modulo $\mathrm{ED}(M)$ of a formula defining it; surely we can look at the set it defines in M, but we may also look at the set it defines in elementary extensions (see also Footnote 16 at page 11).

It turns out that something similar is true of definable types: after all, they are defined through definable sets, so why can't we just evaluate that definable set in an elementary extension and see what happens? This allows us to get a canonical extension to bigger parameter sets, defined as follows.

Definition 8.3.4. Let $B \supseteq M$, and let $p \in S_x(M)$ be a definable type. We define $p \mid B$ as

$$(p \mid B)(x) \coloneqq \{\varphi(x;b) \mid \varphi(x;y) \in L, b \in B^{|y|}, b \vDash (d_p\varphi)(y)\}$$

Exercise 8.3.5. 1. Check that $(p \mid B) \in S_x(B)$ and $(p \mid B) \supseteq p$.

2. Check that whether $p \in S_x(M)$ is definable does not depend on whether we work on T or in $\mathrm{ED}(M)$. In particular, in the definition of $p \mid B$, we may equivalently take $\varphi(x;y) \in L(M)$ instead of $\varphi(x;y) \in L$.

Proposition 8.3.6 (Forking symmetry over models)**.** Suppose that T is stable and $p(x), q(y) \in S(M)$. If $a \vDash p(x)$ and $b \vDash q(y) \mid Ma$, then $a \vDash p(y) \mid Mb$.

Proof. Start with $a_0 \coloneqq a \vDash p(x)$ and $b_0 \coloneqq b \vDash (q \mid Ma)(y)$, and inductively choose $a_{i+1} \vDash p \mid Ma_{\le i}b_{\le i}$ and $b_{i+1} \vDash q \mid Ma_{\le i+1}b_{\le i}$. Note immediately that, by construction, both b_0 and b_1 realise $q \mid Ma_0$, hence for every $\varphi(x,y) \in L(M)$

$$\vDash \varphi(a_0,b_0) \iff \vDash \varphi(a_0,b_1) \tag{8.1}$$

Claim 8.3.7. The sequence of $|a|+|b|$-tuples $(a_ib_i)_{i<\omega}$ is M-indiscernible.

[6] Note that $\Psi(y)$ is still a closed set. You can prove this syntactically, or by recalling that continuous functions from a compact to an Hausdorff space are closed.

[7] You may wonder whether the size of the fibers needs to be uniformly bounded. That is why I have written $X(\mathfrak{U})$ and not just X (compactness!).

Proof of the Claim. By definition, for $i < 2 \le n$ we have $a_n \models \varphi(x, a_i b_i) \iff \models (d_p\varphi)(a_i b_i)$. By (8.1) we have $a_0 b_0 \equiv_M a_0 b_1$. From this, the choice of b_1, and the fact that by construction $a_0 \equiv_M a_1$, it follows that $a_0 b_0 \equiv_M a_1 b_1$. Hence $\models (d_p\varphi)(a_0 b_0) \iff \models (d_p\varphi)(a_1 b_1)$, so $a_0 b_0 a_n \equiv_M a_1 b_1 a_n$. By a similar argument, $a_0 b_0 a_n b_n \equiv_M a_1 b_1 a_n b_n$. From here it is a matter of induction and not getting the indices wrong. $\square_{\text{CLAIM}}$

By the claim and stability, $(a_i b_i)_{i<\omega}$ is *totally* M-indiscernible, and in particular $a_0 b_0 a_1 b_1 \equiv_M a_1 b_1 a_0 b_0$. From this and the fact that, by construction, $a_1 \models p \mid M b_0$, it follows that $a_0 \models p \mid M b_1$. By Exercise 8.3.5, for every $\varphi(x, y) \in L(M)$ we have $\models \varphi(a_0, b_1) \iff \models (d_p\varphi)(b_1)$, and because b_1 and b_0 both realise q, and $d_p\varphi \in L(M)$, we have $\models (d_p\varphi)(b_1) \iff \models (d_p\varphi)(b_0)$. Hence $\models \varphi(a_0, b_1) \iff \models (d_p\varphi)(b_0)$, which together with (8.1) gives us the conclusion. $\square$

Exercise 8.3.8. Let M be ω-saturated. If $p \in S_x(M)$ and $B \supseteq M$, then $\mathrm{RM}(p \mid B) = \mathrm{RM}(p)$.

Remark 8.3.9. In fact, in the previous exercise the saturation assumption is not necessary. One can prove the conclusion just assuming that M is a model, but the proof gets more involved/requires developing a bit more forking calculus. For this reason, below we state a lot of things only for ω-saturated M (it will anyway suffice for our purposes), but be aware that this assumption can be dropped as soon as you know how to drop it from Exercise 8.3.8.

8.4 The connected component

Definition 8.4.1. Let G be a definable group. The *connected component* $G^0(\mathfrak{U})$ is the intersection of all $\mathfrak{U}$-definable subgroups of $G(\mathfrak{U})$ of finite index. We say that G is *connected* iff $G(\mathfrak{U}) = G^0(\mathfrak{U})$.

In general, there can be infinitely many finite index definable subgroups of G, hence G^0 is not guaranteed to be of finite index, nor to be definable.

Proposition 8.4.2. Let G be a definable group.

1. $G^0(\mathfrak{U})$ is normal; in fact, it is definably characteristic, that is, it is fixed setwise by every definable automorphism of $G(\mathfrak{U})$.

2. If G has the DCC on definable subgroups, then $G^0(\mathfrak{U})$ has itself finite index, and is $\emptyset$-definable.

Proof. Clearly, every definable automorphism (and in particular every conjugation!) induces a permutation of the definable finite index subgroups of $G(\mathfrak{U})$; since $G^0(\mathfrak{U})$ is their intersection, it must be fixed setwise, proving the first part. For the second part, Proposition 8.1.1 allows us to write $G^0(\mathfrak{U})$ as a finite intersection of finite index, definable groups. This immediately tells us that $G^0(\mathfrak{U})$ is definable, and of finite index. To prove it is $\emptyset$-definable, by Proposition 4.7.4 we only need to show that every $f \in \mathrm{Aut}(\mathfrak{U}/\emptyset)$ fixes $G^0(\mathfrak{U})$ setwise. But every such f, as in the case of definable automorphisms, induces a permutation of the definable finite index subgroups of $G(\mathfrak{U})$, and we conclude as above. $\square$

Remark 8.4.3. If $G^0(\mathfrak{U})$ is $\emptyset$-definable, then $G(\mathfrak{U}) = G^0(\mathfrak{U})$ if and only if for some/all $M \vDash T$ we have $G(M) = G^0(M)$, where the latter denotes the intersection of all M-definable finite index subgroups of $G(M)$.

If T is totally transcendental, then by Theorem 8.1.6 stabilisers of types are definable, and by the previous proposition so is G^0. We may therefore compute their Morley ranks and compare them. We will make use of the following fact.

Exercise 8.4.4. Let T be totally transcendental, and G a definable group. Suppose that $p \in S_G(M)$ and $\psi(x)$ defines $\mathrm{Stab}(p)$ in M. If $N \succ M$, then $\psi(x)$ also defines $\mathrm{Stab}(p \mid N)$.[8]

Proposition 8.4.5. Let T be totally transcendental and let M be ω-saturated.[9] If $p \in S_G(M)$, then

1. $\mathrm{Stab}(p) \subseteq G^0(M)$, and

2. $\textsc{rm}(\mathrm{Stab}(p)) \leq \textsc{rm}(p)$.

From this proof, we occasionally drop the multiplication symbol, that is, we write e.g. ab instead of $a \cdot b$.[10]

Proof. Since G^0 is definable and of finite index, every type over M must choose one of its finitely many cosets. Hence, if $\varphi(x)$ is a formula defining G^0, there must be $b \in G(M)$ such that $\varphi(b^{-1}x) \in p(x)$. Fix $a \in \mathrm{Stab}(p)$, and note that $\varphi(b^{-1}ax) \in p(x)$ by definition. Therefore, whenever $c \in \mathfrak{U}$ is such that $c \vDash p$, we have that $b^{-1}ac$ and $b^{-1}c$ both belong to $G^0(\mathfrak{U})$. Hence, so does $(b^{-1}c)^{-1}b^{-1}ac$, which equals $c^{-1}ac$. But $G^0(\mathfrak{U})$ is normal, hence $a \in G^0(\mathfrak{U}) \cap M = G^0(M)$.

For the second point, suppose that $\psi(x)$ defines $\mathrm{Stab}(p)$, and let $q(x) \in [\psi(x)] \subseteq S_G(M)$ be such that $\textsc{rm}(q(x)) = \textsc{rm}(\psi(x))$. Let $a \vDash p$, then take $b \vDash q \mid Ma$.

Claim 8.4.6. $\textsc{rm}(\mathrm{Stab}(p)) \leq \textsc{rm}(b \cdot a/M)$.

Proof of the Claim. By choice of b and Exercise 8.3.8 , we have $\textsc{rm}(\mathrm{Stab}(p)) = \textsc{rm}(b/M) = \textsc{rm}(b/Ma)$. Because Morley rank is preserved by definable bijections, such as $x \mapsto x \cdot a$, we have $\textsc{rm}(b/Ma) = \textsc{rm}(b \cdot a/Ma)$, and by Proposition 7.1.10 $\textsc{rm}(b \cdot a/Ma) \leq \textsc{rm}(b \cdot a/M)$. □ CLAIM

By Proposition 8.3.6, we also have $a \vDash p \mid Mb$; since, by Exercise 8.4.4, ψ still defines the stabiliser of $p \mid N$ in any $N \succ M$, and in particular in those N containing b, we find that b stabilises $\mathrm{tp}(a/Mb)$, that is, $\mathrm{tp}(b \cdot a/Mb) = \mathrm{tp}(a/Mb)$. In particular $\mathrm{tp}(b \cdot a/M) = \mathrm{tp}(a/M) = p$, and we conclude by the claim. □

[8]Hint: look at the proof of Theorem 8.1.6.
[9]See Remark 8.3.9.
[10]It should be clear from context whether ab is the product of two elements of a definable group or the concatenation of two tuples.

8.5 Generic types

Definition 8.5.1. Let T be totally transcendental, $M \vDash T$, and G a definable group. We call $p \in S_G(M)$ *generic* iff $\mathrm{RM}(p) = \mathrm{RM}(G)$.

Proposition 8.5.2. Let T be totally transcendental, G a definable group, and $M \vDash T$ be ω-saturated.[11].

1. If $p \in S_G(M)$ is generic and $g \in G(M)$, then $g \cdot p$ is also generic.
2. The following are equivalent.
   (a) $p \in S_G(M)$ is generic.
   (b) $\mathrm{Stab}(p)$ has finite index.
   (c) $\mathrm{Stab}(p) = G^0(M)$.
3. There is a unique generic $p \in S_G(M)$ if and only if G is connected.

Proof. 1. This follows from Proposition 8.3.2 and the fact that $a \vDash p$ if and only if $g \cdot a \vDash g \cdot p$.

2. If p is generic, by the previous point and the fact that there can be only finitely many types of maximum Morley rank, the set $\{g \cdot p \mid g \in G(M)\}$ is finite, say equal to $\{g_0 \cdot p, \ldots, g_n \cdot p\}$. Then the index of $\mathrm{Stab}(p)$ is at most $n+1$, since if $a \vDash p$ and $g \in G(M)$ there must be $i \leq n$ with $g \cdot a \equiv_M g_i \cdot a$, that is, with $g^{-1}g_i \in \mathrm{Stab}(p)$, proving (2a) $\Rightarrow$ (2b).

   For (2b) $\Rightarrow$ (2c), if the subgroup $\mathrm{Stab}(p)$ has finite index, since it is definable it must contain G^0, but the other inclusion is always true by Proposition 8.4.5.

   Finally, (2c) $\Rightarrow$ (2a) follows from Proposition 8.4.5 and the fact that, since G^0 has finite index, we have $\mathrm{RM}(G^0) = \mathrm{RM}(G)$.

3. By the first point, if $p \in S_G(M)$ is generic and $g \in G(M)$, then so is $g \cdot p$. Hence, if there is a unique generic type, it is stabilised by the whole of G, and since $\mathrm{Stab}(p) = G^0$ we have left to right.

   Right to left, suppose that $p, q \in S_G(M)$ are generic types, and that $a \vDash p$ and $b \vDash q \mid Ma$. Take some $N \succ M$ with $b \in N$, and let $a_1 \vDash p \mid N$. By Proposition 8.3.6, both a and a_1 realise $p \mid Mb$, so $\mathrm{tp}(a, b/M) = \mathrm{tp}(a_1, b/M)$. Now, $p \mid N$ is still generic by Exercise 8.3.8, and connectedness of G does not depend on M, so by the previous point $\mathrm{Stab}(p \mid N) = G(N)$, and it follows that $b \cdot a_1 \vDash p \mid N$, hence $b \cdot a_1 \vDash p$. From this and the fact that $a, b \equiv_M a_1, b$ it follows that $b \cdot a \vDash p$. If we argue symmetrically, using right stabilisers instead of left stabilisers, we also find out that $b \cdot a \vDash q$, hence $p = q$. □

We will not prove (nor use) it, but you may like to know the following fact.

Fact 8.5.3. If G is a totally transcendental group, and $X \subseteq G$ is definable and generic, that is, of maximal Morley rank, then it is *syndetic*, that is, finitely many translates of X cover G.

[11] Again, see Remark 8.3.9.

8.6 Totally transcendental fields

We conclude the course by characterising the totally transcendental fields. Extra structure is allowed, that is, the language may be larger than the language of fields. We will need the following two facts from Galois theory.

Fact 8.6.1. Every perfect[12] field with no proper Galois extension is algebraically closed.

Fact 8.6.2. Let F/K be a Galois extension of degree n.

1. If $n = p = \mathrm{char}(K) > 0$, then there is $a \in K$ such that the minimal polynomial of F/K is $X^p + X - a$.

2. If K contains all n-th roots of unity, $\mathrm{Char}(K)$ is either 0 or coprime with n, and $\mathrm{Gal}(F/K)$ is cyclic, then there is $a \in K$ such that the minimal polynomial of F/K is $X^n - a$.

Lemma 8.6.3. If K is an infinite totally transcendental field and $n > 1$, then $K^n = K$, that is, the map $x \mapsto x^n$ is surjective. If $p = \mathrm{char}(K) > 0$, then so is the map $x \mapsto x^p + x$.

Proof. Observe immediately that whether a field K satisfies the conclusion or not is written in $\mathrm{Th}(K)$. Hence, it is enough to prove the conclusion for whichever model of $\mathrm{Th}(K)$ we please, say $\mathfrak{U}$.

If $a \neq 0$, then multiplication by a is a definable automorphism of the additive group $(K(\mathfrak{U}), +)$, hence by Proposition 8.4.2 it fixes $K^0(\mathfrak{U})$ setwise. This implies that $K^0(\mathfrak{U})$ is an ideal, and since $K(\mathfrak{U})$ is a field it must be either $\{0\}$ or $K(\mathfrak{U})$. Since $K^0(\mathfrak{U})$ has finite index and $K(\mathfrak{U})$ is infinite, it follows that the additive group $K(\mathfrak{U})$ is connected. Therefore, by Proposition 8.5.2 there is a unique type in $S_K(\mathfrak{U})$ of Morley rank $\mathrm{RM}(K)$. Clearly, $\mathrm{RM}(K) = \mathrm{RM}(K \setminus \{0\})$, and the unique generic type must entail $x \neq 0$. Again by Proposition 8.5.2, it follows that the multiplicative group $(K^\times(\mathfrak{U}), \cdot)$ is also connected.

Fix an ω-saturated[13] $M \prec^+ \mathfrak{U}$, let $p \in S_{K^\times}(M)$ be the unique generic type, and let $a \vDash p$. Since a^n and a are interalgebraic over M, by Proposition 8.3.2 we also have $a^n \vDash p$. Therefore, $p(x) \vdash x \in (K^\times)^n$. It follows that $(K^\times)^n$ has maximal Morley rank. This is the image of $x \mapsto x^n$, an endomorphism of $K^\times$, hence it is a subgroup of finite index, and since $(K^\times, \cdot)$ is connected we have $K^n = K$.

If $p = \mathrm{char}(K) > 0$, then $x \mapsto x^p + x$ is an endomorphism of $(K, +)$, the elements a and $a^p + a$ are interalgebraic over M, and we can argue as above. □

Lemma 8.6.4. Let $n > 1$ and K be an infinite totally transcendental field. If, for every $m \leq n$, the field K contains all m-th roots of unity, then K has no Galois extension of degree n.

Proof. Suppose that n is minimal such that there is some K as above which is a counterexample, as witnessed by some Galois extension $F \supseteq K$ of degree n. If q is a prime dividing n, by basic group theory there is a subgroup of $\mathrm{Gal}(F/K)$

[12] Recall that K is *perfect* iff either $\mathrm{char}(K) = 0$ or $x \mapsto x^{\mathrm{char}\, K}$ is surjective (equivalently, an automorphism).

[13] I am once again asking for your support in looking at Remark 8.3.9.

of degree q, and by [Lan02, Theorem VI.1.8] there is E such that $K \subseteq E \subseteq F$ and $E \subseteq F$ is a Galois extension of degree q.

But then, by Exercise 8.2.2 and Exercise 8.2.3, E is still totally transcendental; since, for every $m \leq n$, the field K contains all m-th roots of unity, so does $E \supseteq K$, and in particular this holds for every $m \leq q$. Therefore, E is another counterexample, and minimality of n yields $E = K$, hence $q = n$. By Fact 8.6.2, depending on whether $q = \mathrm{char}(K)$ or not, the minimal polynomial of F over K is of the form $X^q + X - a$ or $X^q - a$. By Lemma 8.6.3, these polynomials are reducible over K, a contradiction. □

Theorem 8.6.5 (Macintyre)**.** Every totally transcendental infinite field, possibly with extra structure, is algebraically closed.

Proof. Let K be infinite and totally transcendental. By Lemma 8.6.3, K is perfect. If n is minimal such that ζ is a primitive $n+1$-th root of unity not in K, then $K(\zeta)$ is a Galois extension of degree at most n, contradicting Lemma 8.6.4. It follows that K contains all roots of unity, and again by Lemma 8.6.4 K has no Galois extension, so we conclude by Fact 8.6.1. □

Cherlin and Shelah have shown that, in fact, every superstable field is algebraically closed, where a theory is *superstable* iff it is κ-stable for every sufficiently large κ; we know by Theorem 7.2.2 that totally transcendental theories are superstable, but the converse is false: if you go through your list of standard examples, you should be able to find pretty soon a superstable theory which is not ω-stable, and a stable theory which is not superstable.

You may wonder if the above can be generalised to stable fields. The answer is negative: separably closed fields (with no extra structure) are always stable, but they are not always algebraically closed, see [TZ12, Example 8.6.7]. It is still unknown whether these are the only examples.

Conjecture 8.6.6 (Stable fields conjecture)**.** Every infinite stable field is separably closed.

The conjecture is still open; Scanlon has recently suggested that a possible counterexample could be the field $\mathbb{C}(t)$.

8.7 An alternate ending

Some weeks before writing this chapter, I held a poll among the attendees of the course these notes grew out of, asking whether they preferred fields or groups. You probably already guessed who won; had the outcome be different, this chapter would have contained a proof of the following theorem. It also uses the machinery of generic types, and you can read a proof in [Mar02, Section 7.2].

Theorem 8.7.1 (Reineke)**.** Let G be an infinite totally transcendental group.

1. If G has no proper definable infinite subgroup, then G is abelian, and either G is divisible (not necessarily torsion-free), or there is a prime p such that every element has order p.

2. If $\mathrm{RM}(G) = 1$, then G^0 is abelian. In particular, G is abelian-by-finite.

Finally, I should say that several proofs in this chapter are not as conceptual as they could be[14], since giving more enlightening ones would have required, as I hinted some sections ago, the developing of more machinery than we would have had time to go through, namely, stability theory and forking calculus.

If you are curious, luckily there is no shortage of literature about it. Some common sources are [Bal88, Bue96, Pil83, Poi00, She90, TZ12]. Some applications of stability theory are in [Bou98, MMP17, Pil96, Poi01, Wag97].

People have also applied ideas from stability theory to wider settings; see for example [Cas11, Kim14, Sim15, vdD98b, Wag00]. This has resulted in the developing of a wide array of *dividing lines*, classifying first-order theories according to which combinatorial patterns they display (such as OP) and which consequences follow from omitting them. A very nice map of most of them can be found at `http://forkinganddividing.com`.

[14] And some statements are not optimal. Did I already mention Remark 8.3.9?

Bibliography

[Bal88] John T. Baldwin. *Fundamentals of Stability Theory*, volume 12 of *Perspectives in Mathematical Logic*. Springer, 1988.

[Bou98] E. (ed.) Bouscaren. *Model Theory and Algebraic Geometry*. Lecture Notes in Mathematics. Springer, 1998.

[Bue96] Steven Buechler. *Essential Stability Theory*, volume 4 of *Perspectives in Mathematical Logic*. Cambridge University Press, 1996.

[Cas11] Enrique Casanovas. *Simple Theories and Hyperimaginaries*, volume 39 of *Lecture Notes in Logic*. Cambridge University Press, 2011.

[Che19] Artem Chernikov. Lecture notes on stability theory. `https://www.ams.org/open-math-notes/omn-view-listing?listingId=110792`, version of 1st May 2019.

[Che21] Artem Chernikov. Saharon shelah and number 4, a thread. `https://twitter.com/archernikov/status/1363159234242899970`, 2021.

[CK90] Chen Chung Chang and Howard Jerome Keisler. *Model Theory (3rd edition)*. Elsevier, 1990. 3rd edition.

[CKS16] Artem Chernikov, Itay Kaplan, and Saharon Shelah. On non-forking spectra. *Journal of the European Mathematical Society*, 18(12):2821–2848, 2016.

[CS16] Artem Chernikov and Saharon Shelah. On the number of Dedekind cuts and two-cardinal models of dependent theories. *Journal of the Institute of Mathematics of Jussieu*, 15(4):771–784, October 2016. Publisher: Cambridge University Press.

[Del89] Françoise Delon. Définissabilité avec paramètres extérieurs dans Q_p et R. *Proceedings of the American Mathematical Society*, 106(1):193–198, 1989.

[DST19] Max Dickmann, Niels Schwartz, and Marcus Tressl. *Spectral Spaces*. New Mathematical Monographs. Cambridge University Press, Cambridge, 2019.

[EF95] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Springer Monographs in Mathematics. Springer-Verlag, 1995.

[End01] Herbert B. Enderton. *A Mathematical Introduction to Logic*. Harcourt Academic Press, 2001. 2nd edition.

[Fuh62] Gebhard Fuhrken. Bemerkung zu Einer Arbeit E. Engelers. *Mathematical Logic Quarterly*, 8(3-4):277–279, 1962.

[Gil68] Robert Gilmer. A Note on the Algebraic Closure of a Field. *The American Mathematical Monthly*, 75(10):1101–1102, 1968. Publisher: Mathematical Association of America.

[Hay19] Levon Haykazyan. Spaces of types in positive model theory. *The Journal of Symbolic Logic*, 84(02):833–848, June 2019.

[HL19] Martin Hils and François Loeser. *A First Journey through Logic*, volume 89 of *Student Mathematical Library*. AMS, 2019.

[Hod93] Wilfrid Hodges. *Model Theory*, volume 42 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1993.

[Kam22] Mark Kamsma. Type space functors and interpretations in positive logic. *arXiv:2005.03376 [math]*, March 2022. arXiv: 2005.03376.

[Kei76] H. Jerome Keisler. Six classes of theories*. *Journal of the Australian Mathematical Society*, 21(3):257–266, May 1976. Publisher: Cambridge University Press.

[Kim14] Byunghan Kim. *Simplicity Theory*, volume 53 of *Oxford Logic Guides*. Oxford University Press, 2014.

[Kir19] Jonathan Kirby. *An Invitation to Model Theory*. Cambridge University Press, 2019.

[Lan02] Serge Lang. *Algebra*, volume 211 of *Graduate Texts in Mathematics*. Springer, 2002. Revised Third Edition.

[Mar02] David Marker. *Model Theory: an Introduction*, volume 217 of *Graduate Texts in Mathematics*. Springer, 2002.

[Mit72] William Mitchell. Aronszajn trees and the independence of the transfer property. *Annals of Mathematical Logic*, 5(1):21–46, December 1972.

[MMP17] David Marker, Margit Messmer, and Anand Pillay. *Model Theory of Fields*. Lecture Notes in Logic. Cambridge University Press, Cambridge, 2017.

[Mor65] Michael Morley. Categoricity in Power. *Transactions of the American Mathematical Society*, 114(2):514–538, 1965. Publisher: American Mathematical Society.

[MS94] David Marker and Charles Steinhorn. Definable Types in O-minimal Theories. *Journal of Symbolic Logic*, 59:185–198, 1994.

[MT03] Annalisa Marcja and Carlo Toffalori. *A Guide to Classical and Modern Model Theory*, volume 19 of *Trends in logic*. Springer, 2003.

[Pil83] Anand Pillay. *An Introduction to Stabiliy Theory*, volume 8 of *Oxford Logic Guides*. Oxford University Press, 1983.

[Pil96] Anand Pillay. *Geometric Stability Theory*, volume 32 of *Oxford Logic Guides*. Oxford University Press, 1996.

[Poi00] Bruno Poizat. *A Course in Model Theory*. Universitext. Springer, 2000.

[Poi01] Bruno Poizat. *Stable Groups*, volume 87 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2001. Translated from the 1987 French original.

[Pre88] Mike Prest. *Model Theory and Modules*, volume 130 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1988.

[Sac72] Gerald E. Sacks. *Saturated Model Theory*. Reading, Mass., W. A. Benjamin, 1972.

[Sco15] Lynn Scow. Indiscernibles, EM-Types, and Ramsey Classes of Trees. *Notre Dame Journal of Formal Logic*, 56(3):429–447, January 2015. Publisher: Duke University Press.

[She90] Saharon Shelah. *Classification Theory*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North Holland, 1990.

[Sim15] Pierre Simon. *A guide to NIP theories*, volume 44 of *Lecture Notes in Logic*. Cambridge University Press, 2015.

[TZ12] Katrin Tent and Martin Ziegler. *A Course in Model Theory*, volume 40 of *Lecture Notes in Logic*. Cambridge University Press, 2012.

[vdD98a] Lou van den Dries. *Tame Topology and O-minimal Structures*, volume 248 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1998.

[vdD98b] Lou van den Dries. *Tame Topology and O-minimal Structures*, volume 248 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1998.

[Wag97] F. Wagner. *Stable Groups*. London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge, 1997.

[Wag00] Frank Olaf Wagner. *Simple Theories*, volume 503 of *Mathematics and Its Applications*. Kluwer Academic Publishers, 2000.