

225A: Model Theory

Nir Elber

Fall 2023

CONTENTS

How strange to actually have to see the path of your journey in order to make it.

—Neal Shusterman, [Shu16]

Contents	2
1 Compactness	4
1.1 August 24	4
1.1.1 Logistics	4
1.1.2 Languages and Structures	4
1.1.3 Formulae	5
1.2 August 29	7
1.2.1 Theories	7
1.2.2 Definable Sets	8
1.2.3 The Compactness Theorem	9
1.3 August 31	10
1.3.1 Proof of Compactness	11
1.4 September 5	13
1.4.1 Completing the proof of Theorem 1.32	13
1.5 September 7	16
1.5.1 Ultrafilters	16
1.5.2 Compactness via Ultraproducts	18
1.6 September 12	21
1.6.1 Elementary Equivalence	21
1.7 September 14	24
1.7.1 The Löwenheim–Skolem Theorem	24
1.8 September 19	26
1.8.1 An Example of the Back-and-Forth Method	26
1.8.2 Dense Linear Orders Without Endpoints	28
2 Eliminating Quantifiers	30
2.1 September 21	30
2.1.1 Algebraically Closed Fields	30
2.2 September 26	33

2.2.1	A Taste of Types	33
2.2.2	Back to Algebraically Closed Fields	35
2.3	September 28	36
2.3.1	Ehrenfeucht–Fraïssé Games	36
2.4	October 3	38
2.4.1	Real Closed Fields	39
2.5	October 5	41
2.5.1	Quantifier Elimination via Back-and-Forth	41
2.5.2	Back to Real Closed Fields	42
2.6	October 10	43
2.6.1	Back to Back to Real Closed Fields	43
2.6.2	ω -Minimality	44
2.7	October 12	45
2.7.1	More on ω -Minimality	45
2.8	October 17	48
2.8.1	The Cell Decomposition Theorem	48
3	Types	51
3.1	October 26	51
3.1.1	Introducing Types	51
3.1.2	Types with Parameters	52
3.1.3	Automorphisms	54
3.2	October 31	54
3.2.1	Partial Elementary Embeddings	54
3.2.2	Back to Automorphisms	55
3.3	November 2	56
3.3.1	The Stone Topology	56
3.3.2	Isolated Types	58
3.4	November 7	59
3.4.1	Types of Theories	60
3.4.2	Type Omitting	61
3.5	November 9	62
3.5.1	Prime and Atomic Models	62
3.6	November 14	64
3.6.1	Homogeneous Models	65
3.6.2	Saturated Models	66
3.7	November 16	67
3.7.1	Countable Saturated Models	67
3.7.2	Universal Models	68
3.8	November 21	69
3.8.1	Miscellaneous Saturated Models	69
3.8.2	Construction of Saturated Models	70
3.9	November 28	71
3.9.1	Primer on Indiscernibles	71
3.9.2	Introducing Indiscernibles	72
3.10	November 30	73
3.10.1	Skolem Functions	73
3.10.2	Building Automorphisms	74
	Bibliography	77
	List of Definitions	78

THEME 1

COMPACTNESS

That something so small could be so beautiful.

—Anthony Doerr, [Doe14]

1.1 August 24

It begins.

1.1.1 Logistics

Here are some logistical notes.

- There is a [bCourses](#).
- We will use [Mar02].
- Professor Montalbán and Scanlon will teach the course jointly.
- There will be a midterm (in-class on the 19th of October) and final exam (take-home over three days).
- There are suggested but technically ungraded exercises. They are helpful.
- We will assume basic first-order logic, and examples will be taken from a few other areas of mathematics.
- This is a graduate class. It will be pretty fast.

We are studying model theory, which is the study of models and theories. Our main two theorems are the Compactness theorems and results on admitting types. We will use these results again and again.

1.1.2 Languages and Structures

Let's review chapter 1 of [Mar02]. Here is a language.

Definition 1.1 (language). A language \mathcal{L} consists of the sets \mathcal{F} , \mathcal{R} , and \mathcal{C} of symbols. Here, \mathcal{F} are functions, \mathcal{R} are relations, and \mathcal{C} are constants. Notably, there is an arity function $n: (\mathcal{F} \cup \mathcal{R}) \rightarrow \mathbb{N}$.

Concretely, fix a language $\mathcal{L} = (\mathcal{F}, \mathcal{R}, \mathcal{C})$. If $f \in \mathcal{F}$ and $n(f) = 3$, then we say that f has arity three; the analogous statement holds for relations.

We will often abbreviate a language to just a long tuple. For example, the notation $(\mathbb{N}, 0, 1, +, \leq)$ has the domain \mathbb{N} and constants 0 and 1 and function $+$ and relation \leq , even though the notation has not made it obvious what any of these things are.

So far we only have the prototype of data. Here is the data.

Definition 1.2 (structure). Fix a language \mathcal{L} . Then an \mathcal{L} -structure \mathcal{M} consists of the following data.

- Domain: a nonempty set M .
- Functions: for each $f \in \mathcal{F}$, there is a function $f^{\mathcal{M}}: M^{n(f)} \rightarrow M$.
- Relations: for each $R \in \mathcal{R}$, there is a relation $R^{\mathcal{M}} \subseteq M^{n(R)}$.
- Constants: for each $c \in \mathcal{C}$, there is a constant $c^{\mathcal{M}} \in M$.

The various $(-)^{\mathcal{M}}$ data are called *interpretations*.

Example 1.3. Consider the language \mathcal{L} with the constants 0 and 1 and operations $+$ and \times . Then \mathbb{N} is an \mathcal{L} -structure, in the obvious way.

In general, algebra provides many examples of languages.

We would like to relate our structures.

Definition 1.4 (homomorphism, embedding, isomorphism). Fix a language \mathcal{L} . An \mathcal{L} -homomorphism $\eta: \mathcal{M} \rightarrow \mathcal{N}$ of \mathcal{L} -structures \mathcal{M} and \mathcal{N} is a one-to-one map $\eta: M \rightarrow N$ preserving the interpretations, as follows.

- Functions: for each $f \in \mathcal{F}$, we have $\eta \circ f^{\mathcal{M}} = f^{\mathcal{N}} \circ \eta^{n(f)}$.
- Relations: for each $R \in \mathcal{R}$, if $\overline{m} \in R^{\mathcal{M}}$, then $\eta^{n(R)}(\overline{m}) \in R^{\mathcal{N}}$.
- Constants: for each $c \in \mathcal{C}$, we have $\eta(c^{\mathcal{M}}) = c^{\mathcal{N}}$.

If $\eta: M \rightarrow N$ is one-to-one and the relations condition is an equivalence, then η is an \mathcal{L} -embedding. If $\eta: M \rightarrow N$ is the identity $M \subseteq N$, then we say that \mathcal{M} is an \mathcal{L} -substructure. In addition, if η is onto, then η is an \mathcal{L} -isomorphism.

Explicitly, being a substructure means that the functions and relations are restricted appropriately, and the constants remain the same.

Example 1.5. In the language of groups, subgroups make substructures. A similar sentence holds for other algebraic structures.

1.1.3 Formulae

Thus far we have described a vocabulary: the language provides the data for us to manipulate. We now discuss how to “speak” in this language.

Definition 1.6 (term). Let \mathcal{L} be a language. The set of \mathcal{L} -terms is the smallest set \mathcal{T} satisfying the following.

- Constants: for each $c \in \mathcal{C}$, we have $c \in \mathcal{T}$.
- Variables: $x_i \in \mathcal{T}$ for each $i \in \mathbb{N}$. Notably, we have only countably many variables.
- Functions: if $t_1, \dots, t_n \in \mathcal{T}$ where $n = n(f)$ for some $f \in \mathcal{F}$, then $f(t_1, \dots, t_n) \in \mathcal{T}$.

Given an \mathcal{L} -structure \mathcal{M} and term $t \in \mathcal{T}$ with variables x_1, \dots, x_n and elements $a_1, \dots, a_n \in M$, we define $t^{\mathcal{M}}(\bar{a})$ in the obvious way.

Terms are basically just nouns. We would now like to put them into sentences.

Definition 1.7 (atomic formula). The set of *atomic \mathcal{L} -formulae* is the set of expressions of one of the following forms.

- Equality: $t_1 = t_2$ for any \mathcal{L} -terms t_1 and t_2 .
- Relations: $R(t_1, \dots, t_n)$ for any n -ary relation R and \mathcal{L} -terms t_1, \dots, t_n .

Definition 1.8 (formula). The set of \mathcal{L} -formulae is the smallest set satisfying the following.

- Any atomic \mathcal{L} -formula φ is an \mathcal{L} -formula.
- For any \mathcal{L} -formulae φ and ψ , then $\neg\varphi$ and $\varphi \wedge \psi$ and $\varphi \vee \psi$ are \mathcal{L} -formulae.
- For any variable v_i for $i \in \mathbb{N}$, then $\exists v_i \varphi$ is an \mathcal{L} -formula.

One can then define the shorthand " $\varphi \rightarrow \psi$ " for $\neg\varphi \vee \psi$ and " $\forall v_i \varphi$ " for $\neg\exists v_i \neg\varphi$.

Now that we can talk about our structure, we would like to know if we are making sense.

Definition 1.9 (free variable). Fix a language \mathcal{L} . A variable v in a formula φ is *free* if and only if it is not bound to any quantifier $\exists v$ or $\forall v$. If φ has free variables contained in the variables x_1, \dots, x_n , we write $\varphi(x_1, \dots, x_n)$.

This definition is vague because we have not said what "bound" means, but it is rather obnoxious to explain what it is rigorously, so we will not bother.

Definition 1.10 (sentence). Fix a language \mathcal{L} . An \mathcal{L} -formula with no free variables is a *sentence*.

Definition 1.11 (truth). Fix an \mathcal{L} -structure \mathcal{M} . Further, fix an \mathcal{L} -formula $\varphi(x_1, \dots, x_n)$ and a tuple $\bar{a} \in M^n$. Then we define *truth* as $\mathcal{M} \models \varphi(\bar{a})$ to mean that φ is true upon plugging in \bar{a} , where our definition is inductive on atomic formulae as follows.

- $\mathcal{M} \models (t_1 = t_2)(\bar{a})$ if and only if $t_1^{\mathcal{M}}(\bar{a}) = t_2^{\mathcal{M}}(\bar{a})$.
- $\mathcal{M} \models R(t_1, \dots, t_n)$ if and only if $(t_1^{\mathcal{M}}(\bar{a}), \dots, t_n^{\mathcal{M}}(\bar{a})) \in R^{\mathcal{M}}$.

We define truth inductively on formulae now as follows.

- $\mathcal{M} \models (\varphi \wedge \psi)(\bar{a})$ if and only if $\mathcal{M} \models \varphi(\bar{a})$ and $\mathcal{M} \models \psi(\bar{a})$.
- $\mathcal{M} \models (\varphi \vee \psi)(\bar{a})$ if and only if $\mathcal{M} \models \varphi(\bar{a})$ or $\mathcal{M} \models \psi(\bar{a})$.
- $\mathcal{M} \models \neg\varphi(\bar{a})$ if and only if we do not have $\mathcal{M} \models \varphi(\bar{a})$.
- $\mathcal{M} \models \exists v\varphi(\bar{a}, v)$ if and only if there exists $b \in M$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$.

In this case, we say that \mathcal{M} *satisfies, models, etc.* $\varphi(\bar{a})$ and so on.

Here is our first result of substance.

Proposition 1.12. Fix a language \mathcal{L} and an \mathcal{L} -embedding $\eta: \mathcal{M} \rightarrow \mathcal{N}$. Further, fix a quantifier-free formula φ and $\bar{a} \in M^n$. Then $\mathcal{M} \models \varphi(\bar{a})$ if and only if $\mathcal{N} \models \varphi(\bar{a})$.

Proof. Induction on φ . Roughly speaking, the point is that the interpretations are the same before and after. ■

Remark 1.13. If we allow variables, the statement is false. For example, $(\mathbb{N}, 0, \leq)$ embeds into $(\mathbb{Z}, 0, \leq)$, but $\forall x(0 \leq x)$ is true in the first formula while false in the second.

In the case of isomorphism, we can say more.

Proposition 1.14. Fix a language \mathcal{L} and an \mathcal{L} -isomorphism $\eta: \mathcal{M} \rightarrow \mathcal{N}$. Further, fix any formula φ with free variables x_1, \dots, x_n and a tuple $\bar{a} \in M^n$. Then $\mathcal{M} \models \varphi(\bar{a})$ if and only if $\mathcal{N} \models \varphi(f(\bar{a}))$.

Proof. Induction on φ . The point is that the definition of truth is the same before and after η . ■

1.2 August 29

We continue with the speed run of first-order logic. The goal for today is to state the Compactness theorem.

1.2.1 Theories

Now that we have a notion of truth, it will be helpful to keep track of which sentences exactly we want to be true.

Definition 1.15 (theory). Fix an \mathcal{L} -structure \mathcal{M} . Then the *theory* $\text{Th}(\mathcal{M})$ of \mathcal{M} is the set of all sentences φ such that $\mathcal{M} \models \varphi$.

The theory is essentially all that first-order logic can see.

Definition 1.16 (elementary equivalence). Fix \mathcal{L} -structures \mathcal{M} and \mathcal{N} . Then we say that \mathcal{M} and \mathcal{N} , written $\mathcal{M} \equiv \mathcal{N}$, are *elementarily equivalent* if and only if $\text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$.

Example 1.17. It happens that $(\mathbb{Q}, +) \equiv (\mathbb{R}, +)$ but are not isomorphic because they have different cardinalities.

Example 1.18. Let s denote the successor function. It happens that $(\mathbb{Z}, s) \equiv (\mathbb{Q}, s)$, but one can show that they are not isomorphic.

This notion is different from isomorphism, but it is related.

Lemma 1.19. Fix \mathcal{L} -structures \mathcal{M} and \mathcal{N} . If $\mathcal{M} \cong \mathcal{N}$, then $\mathcal{M} \equiv \mathcal{N}$.

Proof. This is the content of Proposition 1.14 upon unraveling the definitions. ■

Going in the other direction, we might start with some sentences we want to be true and then look for the corresponding models.

Definition 1.20 (theory). Fix a language \mathcal{L} . Then an \mathcal{L} -theory T is a set of \mathcal{L} -sentences. For an \mathcal{L} -structure \mathcal{M} , we say that \mathcal{M} *models* T , written $\mathcal{M} \models T$, if and only if $\mathcal{M} \models \varphi$ for all $\varphi \in T$. We let $\text{Mod}(T)$ denote the class of all models \mathcal{M} of T , and we call it an *elementary class*.

Example 1.21. The class of all groups arises from the language $\{e, \cdot\}$ with some sentences to make a theory. However, the class of torsion groups is not an elementary class.

We want might want to understand what sentences follow from a given theory.

Definition 1.22. Fix a language \mathcal{L} and theory T . Then we say that T *logically implies* a sentence φ , written $T \models \varphi$, if and only if any \mathcal{L} -structure \mathcal{M} modelling T has $\mathcal{M} \models \varphi$.

Remark 1.23. Gödel's completeness theorem shows that $T \models \varphi$ if and only if there is a "proof" of φ from T . We will not use the notion of proof so much, though its proof is similar to the proof of compactness, which we will show.

1.2.2 Definable Sets

We will want the following notion.

Definition 1.24 (definable). Fix an \mathcal{L} -structure \mathcal{M} and subset $B \subseteq M$. Then a subset $X \subseteq M^n$ is *B-definable* if and only if there is a formula $\varphi(v_1, \dots, v_n, w_1, \dots, w_k)$ and tuple $\bar{b} \in B^k$ such that $\bar{a} \in X$ if and only if $\mathcal{M} \models \varphi(\bar{a}, \bar{b})$. The tuple \bar{b} might be called the *parameters*. We may abbreviate M -definable to simply *definable*.

Example 1.25. Any finite set is definable by using the parameters to list out the elements.

Example 1.26. Work with $\mathcal{M} := (\mathbb{Z}, \leq)$. Then $X = \mathbb{N}$ is $\{0\}$ -definable by $\varphi(x, 0)$ where $\varphi(x, y)$ is given by $y \leq x$. However, \mathbb{N} is not \emptyset -definable, as shown by the following proposition.

Proposition 1.27. Fix an \mathcal{L} -structure \mathcal{M} and subset $A \subseteq M$. Further, suppose $X \subseteq M^n$ is A -definable. For any automorphism $\sigma: \mathcal{M} \rightarrow \mathcal{M}$ fixing A pointwise must fix X (not necessarily pointwise).

Proof. Suppose $\varphi(\bar{v}, \bar{w})$ defines X with the parameters $\bar{a} \in A^\bullet$. Then $\bar{x} \in X$ if and only if $\mathcal{M} \models \varphi(\bar{x}, \bar{a})$, but then $\mathcal{M} \models \varphi(\sigma(\bar{x}), \sigma(\bar{a}))$, so $\mathcal{M} \models \varphi(\sigma(\bar{x}), \bar{a})$ so $\sigma(\bar{x}) \in X$. For the converse, use the inverse automorphism σ^{-1} . ■

To further explain Example 1.26, we see that there are automorphisms of \mathbb{Z} (namely, by shifting) which do not fix \mathbb{N} , so \mathbb{N} cannot be \emptyset -definable.

Example 1.28. Work with $\mathcal{M} := (\{1A, 1B, 2A, 2B\}, \leq)$ with partial ordering given by the number. The set $X := \{1A, 1B\}$ is \emptyset -definable by $\varphi(x)$ given by $\exists y((x \neq y) \wedge (x \leq y))$. However, there is an automorphism of our model swapping $1A$ with $1B$ and $2A$ with $2B$, but this automorphism does not fix X pointwise.

1.2.3 The Compactness Theorem

To state compactness, we want a few definitions.

Definition 1.29 (satisfiable). Fix a language \mathcal{L} and theory T . Then T is *satisfiable* if and only if it has a model \mathcal{M} .

With a notion of proof, one can show that being satisfiable means that there is no proof of \perp , but we will avoid a discussion of proofs in this course.

Definition 1.30 (finitely satisfiable). Fix a language \mathcal{L} and theory T . Then T is *finitely satisfiable* if and only if any finite subset of T is satisfiable.

Of course, being satisfiable implies being finitely satisfiable; the converse will be true but is far from obvious. The following example explains why this is strange.

Example 1.31. Consider the natural numbers $\mathcal{N} = (\mathbb{N}, 0, 1, +, \times, \leq)$ and $\mathcal{N}_c := (\mathbb{N}, 0, 1, +, \times, \leq, c)$, where c is some constant symbol, and set

$$T := \text{Th}(\mathcal{N}) \cup \left\{ c \geq \underbrace{1 + 1 + \cdots + 1}_n : n \in \mathbb{N} \right\}.$$

Then T is finitely satisfiable by \mathcal{N} because, for any finite subset of T , the sentences $c \geq 1 + 1 + \cdots + 1$ will have to be bounded in length in our finite subset, so we simply find some c large enough in \mathcal{N} . However, \mathcal{N} does not model T !

Anyway, here is our statement.

Theorem 1.32 (compactness). Fix a language \mathcal{L} and theory T . If T is finitely satisfiable, then T is satisfiable. Furthermore, T has a model \mathcal{M} with cardinality at most $|\mathcal{L}| + \aleph_0$.

Remark 1.33. In particular, the theory T of Example 1.31 has a model \mathcal{N}' , which is going to look very strange. To begin, there is a segment

$$0 < 1 < 2 < \dots$$

But there is now an element c larger than any natural, which produces $c + 1, c + 2, c + 3, \dots$. But also any nonzero element has a predecessor, so we have elements $c - 1, c - 2, c - 3, \dots$. Further, any natural number is either odd or even, so there is also either $(c - 1)/2$ or $c/2$ sitting between the initial piece of \mathbb{N} and the c piece with \mathbb{Z} added everywhere. In fact, a similar argument holds to produce an element approximately equal to qc for any rational $q \in \mathbb{Q}$.

Remark 1.34. One can of course always make our model larger. For example, suppose we have a theory T with an infinite model. If we want a model with cardinality at least \mathbb{R} , we add constants $\{c_r : r \in \mathbb{R}\}$ to our language and add in the sentences

$$\{c_r \neq c_s : \text{distinct } r, s \in \mathbb{R}\}.$$

This remains finitely satisfiable: these constants merely ask for our model to be larger than any finite set. One can even require the new model to be elementarily equivalent to the previous one.

Here are some applications of compactness.

Corollary 1.35. The class of torsion groups is not elementarily definable in the language $\mathcal{L} = \{e, *\}$ of groups.

Notably, it is not okay to write something like

$$\bigvee_{n \in \mathbb{N}} (\forall g \, g^n = e)$$

to encode any torsion because this statement is infinitely long.

Proof. Suppose the class is elementarily definable. Then we have a theory T such that $\text{Mod}(T)$ consists exactly of all torsion groups. Now the trick is to build a model of T which is not actually a torsion group. For this, we expand our language to $\mathcal{L} = \{e, *, c\}$, and let

$$S := T \cup \left\{ \underbrace{c * c * \dots * c}_n \neq e : n \geq 2 \right\}.$$

For any finite subset of S , we can satisfy S by a torsion group containing an element which is not n -torsion for sufficiently large n ; for example, $\mathbb{Z}/n\mathbb{Z}$ will do.

Thus, by Theorem 1.32, there is a model \mathcal{G} of S , so in particular, \mathcal{G} has an element $g \in G$ with

$$\underbrace{g * g * \dots * g}_n \neq e$$

for all $n \geq 2$ (namely, g is the interpretation of the constant symbol c), so it follows that G is not torsion. However, \mathcal{G} is also a model of T and thus is supposed to be torsion, so we have a contradiction! This completes the proof. ■

1.3 August 31

Professor Scanlon is back. Let's prove the Compactness theorem. We are going to prove 2.5 times.

1.3.1 Proof of Compactness

Recall the statement.

Theorem 1.32 (compactness). Fix a language \mathcal{L} and theory T . If T is finitely satisfiable, then T is satisfiable. Furthermore, T has a model \mathcal{M} with cardinality at most $|\mathcal{L}| + \aleph_0$.

Remark 1.36. This result is special to first-order logic: in some sense, Theorem 1.32 combined with a corollary characterizes first-order logic among various logics. Roughly speaking, one wants to formalize what a logic is with its various structures and sentences should do.

Proof with completeness. We can prove this result fairly quickly given the Completeness theorem. Recall that the Completeness theorem says that any theory fails to be satisfiable if and only if there is a proof of contradiction \perp ; one writes that a theory T proves a sentence φ by $T \vdash \varphi$. We have not discussed how formal proofs work, and we won't discuss it further because this is not a proofs class. Approximately speaking, a formal proof is a list of steps one can use the sentence in T to produce φ syntactically.

Now, suppose that T fails to be satisfiable. Then there is a proof of \perp . But then one can look at the proof, which is necessarily finite in length, and then we pick up any sentence φ occurring in the proof of \perp . But then we have a proof of \perp using only finitely many sentences in T , so T fails to be finitely satisfiable! This completes the proof. ■

Anyway, let's present an actual proof.

Definition 1.37 (witness). Fix a theory T of a language \mathcal{L} . Then T has *witnesses* (or *Henkin constants*) if and only if each formula $\varphi(x)$ in one free variable x has a constant symbol c such that $\exists x \varphi(x) \rightarrow \varphi(c)$ lives in T .

Remark 1.38. If T has witnesses, then $T' \supseteq T$ also has witnesses for any theory T' extending T .

Let's quickly sketch our proof.

1. We will show that if T is finitely satisfiable, then there is an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ and expanded finitely satisfiable \mathcal{L}' -theory $T' \supseteq T$ of \mathcal{L}' such that $|\mathcal{L}'| \leq |\mathcal{L}| + \aleph_0$, and T' has witnesses (as does any extended theory T'' of T').
2. Next, suppose T is a maximally finitely satisfiable theory (i.e., T is finitely satisfiable, and any proper extension $T' \supseteq T$ fails to be finitely satisfiable¹). Then we will show T is complete (i.e., each sentence φ has either $\varphi \in T$ or $\neg\varphi \in T$).
3. From here, we want to extend maintaining being complete: if T is finitely satisfiable, then there is an extended language $\mathcal{L}' \supseteq \mathcal{L}$ of size $|\mathcal{L}'| = |\mathcal{L}| + \aleph_0$ and extended theory T' of T which is complete, finitely satisfiable, and has witnesses. This essentially follows from a Zorn's lemma argument.
4. We are now ready to do our construction. At this point, we may assume that our theory T is finitely satisfiable, complete, and has witnesses. Then we claim that there is a model \mathcal{M} such that $|\mathcal{M}| \leq |\mathcal{L}|$. In fact, the model can be described somewhat explicitly. Take $M := \mathcal{C}/\sim$ where \mathcal{C} is our set of constants, and our equivalence relation \sim is given by $c \sim d$ if and only if $(c = d) \in T$. Notably, constants $c \in \mathcal{C}$ are interpreted as $c^{\mathcal{M}} := [c]$. To interpret functions f , we have $f^{\mathcal{M}}([a_1], \dots, [a_n]) = [d]$ if and only if $(f(a_1, \dots, a_n) = d) \in T$. Lastly, to interpret relations R , we have $R^{\mathcal{M}}([a_1], \dots, [a_n])$ if and only if $(R(a_1, \dots, a_n)) \in T$.

Let's start implementing the details.

¹ Such a thing exists by some sort of Zorn's lemma argument: note that there is a theory containing T which fails to be finitely satisfiable: take the set of all sentences!

Remark 1.39. In logic, the answer to a question is often the question. For example, in step 4, we see that T has a model because T says that it has a model.

Here is our first step.

Lemma 1.40. Fix a finitely satisfiable theory T of a language \mathcal{L} . Then there is an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ and expanded finitely satisfiable \mathcal{L}' -theory $T' \supseteq T$ of \mathcal{L}' such that $|\mathcal{L}'| \leq |\mathcal{L}| + \aleph_0$, and T' has witnesses.

Proof. We would like to just set T' to be T together with new constants providing witnesses for all formulae. But these new constants will make new formulae, so we need to do some kind of induction to go upwards.

With this in mind, we will build an increasing sequence of languages

$$\mathcal{L}_0 := \mathcal{L} \subseteq \mathcal{L}_1 \subseteq \mathcal{L}_2 \subseteq \cdots$$

and theories

$$T_0 := T \subseteq T_1 \subseteq T_2 \subseteq \cdots$$

such that T_n is always a finitely satisfiable \mathcal{L}_n -theory, and each \mathcal{L}_n -formula φ with one free variable has a constant $c \in \mathcal{C}_{\mathcal{L}_n}$ such that $\exists x \varphi(x) \rightarrow \varphi(c)$ lives in T_n . We will then set \mathcal{L}' to be the union of the \mathcal{L}_n and T' to be the union of the T_n , and this will complete the proof.

We have already built the $n = 0$ stage, as above. Then to build \mathcal{L}_{n+1} from \mathcal{L}_n , add in new constant symbols $c_{\varphi(x)}$ for each \mathcal{L}_n -formula $\varphi(x)$ with one free variable; all the functions and relations remain the same. Note \mathcal{L}_{n+1} is now the size of the formulae with one free variable in \mathcal{L}_n , so $|\mathcal{L}_{n+1}| = |\mathcal{L}_n| + \aleph_0$.

As for our theory, let T_{n+1} be T_n plus the sentences $\exists x \varphi(x) \rightarrow \varphi(c_{\varphi(x)})$ for each \mathcal{L}_n -formula $\varphi(x)$ with one free variable. We are now ready to set

$$\mathcal{L}' := \bigcup_{n \in \mathbb{N}} \mathcal{L}_n \quad \text{and} \quad T' := \bigcup_{n \in \mathbb{N}} T_n.$$

We see that \mathcal{L}' then has the right size, and T' has witnesses: for any \mathcal{L}' -formula $\varphi(x)$ with one free variable, note that $\varphi(x)$ has only finitely many symbols, so we can find some fixed level \mathcal{L}_n containing all the symbols used in $\varphi(x)$. But then $\varphi(x)$ has a witness from $T_{n+1} \subseteq T'$, as needed.

It remains to show that T' is finitely satisfiable. It suffices to show that T_n is finitely satisfiable for any $n \in \mathbb{N}$ because any finite set will be contained in some T_n . We show this by induction. For $n = 0$, there is nothing to say. Now suppose T_n is finitely satisfiable, and we show that T_{n+1} is finitely satisfiable.

Fix some finite subset $\Delta \subseteq T_{n+1}$ which we would like to build a model for. Now, Δ will be built by some sentences in T_n plus some new sentences from T_{n+1} . Looking hard at $T_{n+1} \setminus T_n$, we see that we can enumerate $\Delta \setminus T_n$ as some sentences

$$\exists x \psi_k(x) \rightarrow \psi_k(c_k)$$

where $\{\psi_k\}_{k=1}^m$ is some \mathcal{L}_n -formulae in one free variable.

We now begin building our model. Note $\Delta \cap T_n$ is a finite subset of T_n , so it is satisfiable by some model \mathcal{M} . Now, for each k , if there is some $a \in M$ such that $\mathcal{M} \models \psi_k(a)$, set $a := a_{k,i}$; otherwise, set $a_k := m$ for any chosen $m \in M$. (Note structures are nonempty.) We now let \mathcal{M}' be the \mathcal{L}_{n+1} -structure with universe M , interpretations of functions and relations the same as in \mathcal{M} , and all old constant symbols are also all still interpreted the same way. Then for each new constant symbol, we interpret $c_k^{\mathcal{M}'} := a_k$, and each other new constant symbol can also go to m . Now \mathcal{M}' is a model for Δ because it models everything in $\Delta \cap T_n$ for free, and it has satisfied $\Delta \setminus T_{n+1}$ by construction, so we are done. ■

To show the second step, we begin with the following lemma.

Lemma 1.41. Fix a finitely satisfiable theory T of a language \mathcal{L} . For any \mathcal{L} -sentence φ , then either $T \cup \{\varphi\}$ or $T \cup \{\neg \varphi\}$ is finitely satisfiable.

Proof. Suppose that both $T \cup \{\varphi\}$ and $T \cup \{\neg\varphi\}$ both fail to be finitely satisfiable. We will show that T fails to be finitely satisfiable.

Well, we are given finite subsets $\Delta_+ \subseteq T \cup \{\varphi\}$ and $\Delta_- \subseteq T \cup \{\neg\varphi\}$ which are not satisfiable. If Δ_+ fails to contain φ , then Δ_+ is a finite subset of T which is not satisfiable, so T fails to be satisfiable. Thus, we may assume that $\varphi \in \Delta_+$. Analogously, we may assume that $\neg\varphi \in \Delta_-$. Now, $(\Delta_+ \cup \Delta_-) \setminus \{\varphi\}$ and $(\Delta_+ \cup \Delta_-) \setminus \{\neg\varphi\}$ both fail to be satisfiable.

But now suppose for the sake of contradiction that T is finitely satisfiable. Then $(\Delta_+ \cup \Delta_-) \setminus \{\varphi, \neg\varphi\}$ has a model \mathcal{M} . But $\mathcal{M} \models \varphi$ or $\mathcal{M} \models \neg\varphi$, so we see that \mathcal{M} will model at least one of $(\Delta_+ \cup \Delta_-) \setminus \{\varphi\}$ or $(\Delta_+ \cup \Delta_-) \setminus \{\neg\varphi\}$, which is the desired contradiction. ■

The second step now follows from a Zorn's lemma argument.

Lemma 1.42. Fix a maximally finitely satisfiable theory T of a language \mathcal{L} . Then T is complete.

Proof. Let φ be any \mathcal{L} -sentence. Then either $T \cup \{\varphi\}$ or $T \cup \{\neg\varphi\}$ is finitely satisfiable by Lemma 1.41, so by maximality, we may conclude that either $T = T \cup \{\varphi\}$ or $T = T \cup \{\neg\varphi\}$, so either $\varphi \in T$ or $\neg\varphi \in T$, which is what we wanted. ■

Combining the work so far completes the third step.

Lemma 1.43. Fix a finitely satisfiable theory T of a language \mathcal{L} . Then there is an extended language $\mathcal{L}' \supseteq \mathcal{L}$ of size $|\mathcal{L}'| \leq |\mathcal{L}| + \aleph_0$ and extended theory T' of T which is complete, finitely satisfiable, and has witnesses.

Proof. We can prove this using the previous two steps.

1. Lemma 1.40 provides an extended language \mathcal{L}' (of cardinality at most $|\mathcal{L}| + \aleph_0$) and extended theory T' which is finitely satisfiable and has witnesses.
2. We use Zorn's lemma to become maximally finitely satisfiable. Let \mathcal{P} denote the set of finitely satisfiable \mathcal{L}' -theories T'' extending T' which is finitely satisfiable. Containment shows that \mathcal{P} is a partial order, and it's nonempty because $T' \in \mathcal{P}$. Next up, we note that any nonempty chain $\{T_\alpha\}_{\alpha \in \lambda}$ is upper-bounded by

$$\bigcup_{\alpha \in \lambda} T_\alpha,$$

which we can see continues to be finitely satisfiable (any finite set belongs to some T_β for β perhaps large) and thus lives in \mathcal{P} and succeeds to upper-bound our chain. Thus, Zorn's lemma provides a maximally finitely satisfiable theory T'' containing T' , which will be complete by Lemma 1.42. Because T'' contains T' , we continue to have witnesses. ■

1.4 September 5

In this lecture, we will complete our proof of Theorem 1.32.

1.4.1 Completing the proof of Theorem 1.32

Last class, we left off having shown Lemma 1.43, which was the third step of our outline. The last step of the proof is the following lemma.

Lemma 1.44. Fix a language \mathcal{L} with a theory T which is finitely satisfiable, complete, and has witnesses. Then T has a model \mathcal{M} with cardinality at most $|\mathcal{L}|$.

Proof. As we did last class, we go ahead and explicitly describe our model and then show that it makes sense. Take $M := \mathcal{C}/\sim$ where \mathcal{C} is our set of constants, and our equivalence relation \sim is given by $c \sim d$ if and only if $(c = d) \in T$. Notably, constants $c \in \mathcal{C}$ are interpreted as $c^M := [c]$. To interpret functions f , we have $f^M([a_1], \dots, [a_n]) = [d]$ if and only if $(f(a_1, \dots, a_n) = d) \in T$. Lastly, to interpret relations R , we have $R^M([a_1], \dots, [a_n])$ if and only if $(R(a_1, \dots, a_n)) \in T$.

We now check that this makes sense. Note that in the following checks, we are a bit sloppy in differentiating between constants and their equivalence classes in \mathcal{C} when there is unlikely to be problems from doing so.

1. We show that \sim is in fact an equivalence relation on \mathcal{C} . There are the following checks.

- Reflexive: we must show $c = c$ is a sentence in T . Because T is complete, one of $c = c$ or $\neg(c = c)$ is in T . But T is finitely satisfiable, and the sentence $\neg(c = c)$ has no model, so it cannot live in T . So instead $c = c$ lives in T .
- Symmetric: suppose $c \sim c'$ so that $c = c'$ is a sentence in T ; we want to show that $c' = c$ is a sentence in T . Well, by completeness one of $c' = c$ or $\neg(c' = c)$ lives in T . But if we have $\neg(c' = c)$, then the finite theory $\{\neg(c' = c), c = c'\}$ will have no model (symmetry of equality will hold in the model), violating that T is finitely satisfiable. So we must have $c' = c$ instead.
- Transitive: suppose $c \sim c'$ and $c' \sim c''$ so that $c = c'$ and $c' = c''$ are sentences in T . We want to show that $c \sim c''$, or equivalently that $c = c''$ lives in T . Well, by completeness, one of $c = c''$ or $\neg(c = c'')$ lives in T . However, if $\neg(c = c'')$ lives in T , then we note that $\{c = c', c' = c'', \neg(c = c'')\}$ is a subset of T with no model, which is a contradiction. So instead $c = c''$ lives in T .

2. We show that our interpretation of functions makes sense. Fix an n -ary function f . We need to show that $f(a_1, \dots, a_n)$ has a unique interpretation in \mathcal{M} .

- Existence: for constants a_1, \dots, a_n , we show that there is a constant b such that $f(a_1, \dots, a_n) = b$ in T . This holds by having witnesses: let $\varphi(x)$ be the formula $f(a_1, \dots, a_n) = x$, and having witnesses tells us that T contains the sentence

$$\exists x \varphi(x) \rightarrow \varphi(b)$$

for some constant b . We show that T contains the sentence $\varphi(b)$. Otherwise, because T is complete, T will have the sentence $\neg\varphi(b)$, but being finitely satisfiable means that

$$\{\exists x \varphi(x) \rightarrow \varphi(b), \neg\varphi(b)\}$$

must have a model; this is an issue because all models satisfy $\exists x f(a_1, \dots, a_n) = x$ and therefore must satisfy $\varphi(b)$, which is a contradiction to satisfying $\neg\varphi(b)$.

- Uniqueness: for constants a_1, \dots, a_n and a'_1, \dots, a'_n and b and b' such that $a_i \sim a'_i$ for all i and both $f(a_1, \dots, a_n) = b$ and $f(a'_1, \dots, a'_n) = b'$, we must show that actually $b \sim b'$.

Well, by completeness, if $b \sim b'$ is not true, then $\neg(b = b')$ lives in T . Then the theory

$$\{a_1 = a'_1, \dots, a_n = a'_n, f(a_1, \dots, a_n) = b, f(a'_1, \dots, a'_n) = b', \neg(b = b')\}$$

is a subset of T but is not satisfiable (because of how functions work in set theory), which is a contradiction.

3. We show that our interpretation of relations makes sense. Fix an n -ary relation R . Essentially, if we have constants a_1, \dots, a_n and a'_1, \dots, a'_n such that $a_i \sim a'_i$ for each i , then we will have $R(a_1, \dots, a_n) \in T$ if and only if $R(a'_1, \dots, a'_n) \in T$. Because \sim is symmetric as shown above, it suffices to show that $R(a_1, \dots, a_n) \in T$ implies $R(a'_1, \dots, a'_n) \in T$.

Well, T is complete, so if T fails to contain $R(a'_1, \dots, a'_n)$, then it must contain $\neg R(a'_1, \dots, a'_n)$ instead. But then

$$\{a_1 = a'_1, \dots, a_n = a'_n, R(a_1, \dots, a_n), \neg R(a'_1, \dots, a'_n)\}$$

is a finite subset of T with no model because of how relations work in set theory; this is a contradiction.

4. As an intermediate step, before going on to show that $\mathcal{M} \models T$, we show that terms behave: suppose $t(x_1, \dots, x_n)$ is a term. For constants c_1, \dots, c_n, c' , we show that $t(c_1, \dots, c_n) = d$ is in T if and only if $t^{\mathcal{M}}([c_1], \dots, [c_n]) = [d]$.

Let T' be the subset of T with this property. Note that T' contains constants by our first check above. To show that $T' = T$, we suppose that $t_1, \dots, t_m \in T'$ and that f is an m -ary function, and we want to show that $f(t_1, \dots, t_m)$ is in T' . Fix enough constants c_1, \dots, c_n (namely, more than the number of free variables of each t_i). Then we note $t_i^{\mathcal{M}}([c_1], \dots, [c_n]) = [d_i]$ for some $[d_i] \in \mathcal{M}$, which then implies that

$$t_i(c_1, \dots, c_n) = d_i$$

is a sentence in T for each t_i . Now, $f^{\mathcal{M}}(t_1^{\mathcal{M}}, \dots, t_m^{\mathcal{M}})(\bar{c})$ is certainly equal to some constant $[d]$, which is now equivalent to having

$$f(d_1, \dots, d_m) = d$$

in T by the functions check above. Now, the finite satisfiability and completeness of T imply that having the above sentence in T is equivalent to having the sentence

$$f(t_1, \dots, t_m)(\bar{c}) = d$$

in T because T already contains $t_i(\bar{c}) = d_i$ for each i . For example, if T fails to contain $f(t_1, \dots, t_m)(\bar{c})$, then it will contain $\neg f(t_1, \dots, t_m)(\bar{c}) = d$ by completeness, but this contradicts $f(d_1, \dots, d_m) = d$ and $t_i(\bar{c}) = d_i$ for each i and therefore the finite subset with all these sentences is not satisfiable. The reverse implication is similar.

5. We show that \mathcal{M} actually satisfies all sentences in T ; in fact, we will show $T \models \varphi(\bar{a})$ for any φ and \bar{a} if and only if $\mathcal{M} \models \varphi(\bar{a})$. We proceed by induction, starting with atomic formulae.

- Our most basic cases are sentences of the form $c_1 = c_2$ and $R(c_1, \dots, c_n)$ where R is some n -ary relation and c_1, \dots, c_n are constants. These are satisfied by \mathcal{M} basically by construction: the definition of \sim establishes from $c_1 = c_2$ that $c_1 \sim c_2$ and thus $c_1^{\mathcal{M}} = [c_1] = [c_2] = c_2^{\mathcal{M}}$. And $R^{\mathcal{M}}(c_1^{\mathcal{M}}, \dots, c_n^{\mathcal{M}})$ is equivalent to $R(c_1, \dots, c_n) \in T$.
- For any terms t and s and enough constants \bar{a} and \bar{b} , we claim that having $(t = s)(\bar{a}, \bar{b})$ in T implies $\mathcal{M} \models (t = s)(\bar{a}, \bar{b})$. The previous step promises constants c and d such that $t(\bar{a}) = c$ and $s(\bar{b}) = d$ are in T and that this is equivalent to $t^{\mathcal{M}}(\bar{a}) = [c]$ and $s^{\mathcal{M}}(\bar{b}) = [d]$.
Now, $(t = s)(\bar{a}, \bar{b})$ being in T is thus equivalent to having $c = d$ in T by the usual argument using the completeness and finite satisfiability of T . Then having $c = d$ is equivalent to $[c] = [d]$, which is equivalent to $t^{\mathcal{M}}(\bar{a}) = s^{\mathcal{M}}(\bar{b})$, which is equivalent to $\mathcal{M} \models (t = s)(\bar{a}, \bar{b})$.
- For any n -ary relation R and terms t_1, \dots, t_n and enough constants \bar{a} , we claim $R(t_1, \dots, t_n)(\bar{a})$ being in T implies $\mathcal{M} \models R(t_1, \dots, t_n)(\bar{a})$. Well, for each term t_i , the previous step promises us a constant c_i such that $t_i(\bar{a}) = c_i$ is in T and has $t_i^{\mathcal{M}}(\bar{a}) = [c_i]$.
Now, having the sentences $t_i(\bar{a}) = c_i$ for each i implies that $R(t_1, \dots, t_n)(\bar{a})$ lives in T if and only if $R(c_1, \dots, c_n)$ lives in T by the usual argument using the completeness and finite satisfiability of T . But by our relations check, we know that $R(c_1, \dots, c_n)$ lives in T if and only if $R^{\mathcal{M}}([c_1], \dots, [c_n])$ is true, which is equivalent to $R^{\mathcal{M}}(t_1^{\mathcal{M}}(\bar{a}), \dots, t_n^{\mathcal{M}}(\bar{a}))$.

We now build up from atomic formulae. Let F' be the subset of formulae such that $\varphi(\bar{a})$ being in T for some constants \bar{a} if and only if $\mathcal{M} \models \varphi(\bar{a})$. The above checks show that F' contains atomic formulae.

- Suppose $\varphi \in F'$. We show $\neg\varphi \in F'$. Well, $\neg\varphi(\bar{a})$ fails to live in T if and only if $\varphi(\bar{a})$ lives in T (by completeness), which is equivalent to $\mathcal{M} \models \varphi(\bar{a})$, which is equivalent to \mathcal{M} not satisfying $\neg\varphi(\bar{a})$.
- Suppose $\varphi, \psi \in F'$. We show that $\varphi \wedge \psi$. Well, $(\varphi \wedge \psi)(\bar{a})$ lives in T if and only if both $\varphi(\bar{a})$ and $\psi(\bar{a})$ live in T (using the usual argument with the completeness and finite satisfiability of T), which is equivalent to $\mathcal{M} \models \varphi(\bar{a})$ and $\mathcal{M} \models \psi(\bar{a})$, which is equivalent to $\mathcal{M} \models (\varphi \wedge \psi)(\bar{a})$.

- Suppose $\varphi(x) \in F'$. We show that $\exists x \varphi(x) \in F'$. Well, $\mathcal{M} \models (\exists x \varphi(x))(\bar{a})$ if and only if there is $[b] \in M$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$. By hypothesis, this is equivalent to having some constant b such that $\varphi(\bar{a}, b)$ is in T .

Now, if $\varphi(\bar{a}, b)$ is in T for some constant b , then the usual argument with completeness and finite satisfiability requires $(\exists x \varphi(x))(\bar{a})$ to be in T . Conversely, if $(\exists x \varphi(x))(\bar{a})$ is in T , then the fact that T has witnesses implies that there is a constant c such that $\varphi(\bar{a}, c)$ is in T from the usual argument. In particular, the sentence $\exists x \varphi(\bar{a})(x) \rightarrow \varphi(\bar{a})(b)$ belongs to T for some constant b .

The above checks complete the induction on formulae. ■

Theorem 1.32 now follows from combining Lemmas 1.43 and 1.44.

1.5 September 7

In this lecture, we will provide another proof of Theorem 1.32, using ultrafilters.

1.5.1 Ultrafilters

Unsurprisingly, the main character of our story will be ultrafilters.

Definition 1.45 (filter). Fix a set I . Then a *filter* \mathcal{F} on I is a subset of $\mathcal{P}(I)$ satisfying the following.

- (a) $I \in \mathcal{F}$.
- (b) Finite intersection: for $X, Y \in \mathcal{F}$, we have $X \cap Y \in \mathcal{F}$.
- (c) Containment: if $X \in \mathcal{F}$ and $Y \subseteq I$ contains X , then $Y \in \mathcal{F}$ also.

The intuition here is that filters contain “large” subsets of I .

Example 1.46. Fix a set I . Then $\{I\}$ is a filter.

Example 1.47. Fix a set I and a filter \mathcal{F} on I . If $\emptyset \in \mathcal{F}$, then we see that any subset $X \subseteq I$ contains \emptyset and thus must live in \mathcal{F} . Thus, $\mathcal{F} = \mathcal{P}(I)$, which is in fact a filter. We call $\mathcal{P}(I)$ the “trivial filter.”

Example 1.48. More generally, fix any subset $X \subseteq I$. Then $\mathcal{F}_X := \{Y \subseteq I : X \subseteq Y\}$ is a filter.

- (a) Note $X \subseteq I$, so $I \in \mathcal{F}_X$.
- (b) Intersection: if $Y, Z \in \mathcal{F}_X$, then $X \subseteq Y$ and $X \subseteq Z$, so $X \subseteq Y \cap Z$, so $Y \cap Z \in \mathcal{F}_X$.
- (c) Containment: if $Y \in \mathcal{F}_X$, and $Z \subseteq I$ contains Y , then $X \subseteq Y \subseteq Z$, so $Z \in \mathcal{F}_X$.

Example 1.49. Fix a set I , and define $\mathcal{F} \subseteq \mathcal{P}(I)$ by $X \in \mathcal{F}$ if and only if $I \setminus X$ is finite. We check that \mathcal{F} is a filter.

- (a) Note $I \in \mathcal{F}$ because $I \setminus I = \emptyset$ is finite.
- (b) Intersection: if $X, Y \in \mathcal{F}$, then $I \setminus (X \cap Y) = (I \setminus X) \cup (I \setminus Y)$ is finite and thus $X \cap Y \in \mathcal{F}$.
- (c) Containment: if $X \in \mathcal{F}$ and $Y \subseteq I$ contains X , then $I \setminus Y \subseteq I \setminus X$ is finite, so $Y \in \mathcal{F}$.

Ultrafilters are the largest filters.

Definition 1.50 (ultrafilter). Fix a set I . Then an *ultrafilter* \mathcal{F} on I is a nontrivial filter on I such that each subset $X \subseteq I$ has one of $X \in \mathcal{F}$ or $I \setminus X \in \mathcal{F}$.

Example 1.51. Fix a set I and element $a \in I$. Define the “principal ultrafilter”

$$\mathcal{F}_a := \{X \subseteq I : a \in X\}.$$

We show that \mathcal{F}_a is an ultrafilter. Note \mathcal{F}_a is already a filter by Example 1.48. To be ultrafilter, for each $X \subseteq I$, either $a \in X$ or $a \in I \setminus X$, which imply $X \in \mathcal{F}_a$ or $I \setminus X \in \mathcal{F}_a$ respectively.

The following result rigorizes the notion that ultrafilters are the largest filters.

Lemma 1.52. Fix a set I and a filter \mathcal{U} on I . The following are equivalent.

- (a) \mathcal{U} is an ultrafilter.
- (b) \mathcal{U} is maximal among the partially ordered set of nontrivial filters on I , ordered by inclusion.

Proof. We have two implications to show.

- We show (a) implies (b). Suppose \mathcal{U}' is a filter properly containing \mathcal{U} , and we want to show that $\mathcal{U}' = \mathcal{P}(I)$. Well, \mathcal{U}' properly contains \mathcal{U} , so there is some $X \in \mathcal{U}' \setminus \mathcal{U}$. But $X \notin \mathcal{U}$ requires $I \setminus X \in \mathcal{U}$, so $I \setminus X \in \mathcal{U}'$ too, but then

$$\emptyset = X \cap (I \setminus X)$$

lives in \mathcal{U}' . It follows that $\mathcal{U}' = \mathcal{P}(I)$ by Example 1.47.

- We show (b) implies (a). Certainly \mathcal{U} is nontrivial. Now, fix any subset $X \subseteq I$. Suppose $I \setminus X \notin \mathcal{U}$, and we want to show that $X \in \mathcal{U}$. Indeed, consider the filter

$$\mathcal{U}' := \{Y \subseteq I : Y \supseteq X \cap X' \text{ for some } X' \in \mathcal{U}\}.$$

Quickly, we check that \mathcal{U}' is a nontrivial filter containing \mathcal{U} .

- Note $I \supseteq X \cap I$, so $I \in \mathcal{U}'$.
- Intersection: if $Y_1, Y_2 \in \mathcal{U}'$, then find $X_1, X_2 \in \mathcal{U}$ such that $Y_i \supseteq X \cap X_i$ for each i , so $X_1 \cap X_2 \in \mathcal{U}$ implies $Y_1 \cap Y_2 \supseteq X \cap (X_1 \cap X_2)$ and so $Y_1 \cap Y_2 \in \mathcal{U}'$.
- Containment: if $Y \in \mathcal{U}'$ and $Z \subseteq I$ contains Y , then find $X' \in \mathcal{U}$ such that $Y \supseteq X \cap X'$, so $Z \supseteq X \cap X'$, so $Z \in \mathcal{U}'$.
- Contains \mathcal{U} : for each $X' \in \mathcal{U}$, note $X' \supseteq X \cap X'$, so $X' \in \mathcal{U}'$.
- Nontrivial: having $\emptyset \in \mathcal{U}'$ would imply $\emptyset \supseteq X \cap X'$ for some $X' \in \mathcal{U}$, which is equivalent to $X' \subseteq I \setminus X$, so it would follow that $I \setminus X \in \mathcal{U}$, which is a contradiction.

We conclude that $\mathcal{U} = \mathcal{U}'$ by maximality of \mathcal{U} . However, $X \supseteq I \cap X$ forces $X \in \mathcal{U}' = \mathcal{U}$, so we are done. ■

It is important to know that it is relatively easy to build ultrafilters.

Proposition 1.53. Fix a nontrivial filter \mathcal{F} on a set I . Then there exists an ultrafilter \mathcal{U} containing \mathcal{F} .

Proof. Let \mathcal{P} be the set of nontrivial filters containing \mathcal{F} , which we turn into a partially ordered by set by inclusion; note $\mathcal{F} \in \mathcal{P}$, so \mathcal{P} is nonempty. Using Lemma 1.52, we would like to show that \mathcal{P} has a maximal

element, for which we use Zorn's lemma. Fix a nonempty chain $\mathcal{C} \subseteq \mathcal{P}$, which we must upper-bound. We claim that

$$\mathcal{F}_u := \bigcup_{\mathcal{F}' \in \mathcal{C}} \mathcal{F}'$$

is a filter containing \mathcal{F} upper-bounding \mathcal{C} , which will complete the proof. Here are our checks.

- Upper-bounds: for any $\mathcal{F}' \in \mathcal{C}$, we see that $\mathcal{F}' \subseteq \mathcal{F}_u$ by construction.
- Any $\mathcal{F}' \in \mathcal{C}$ contains I , so $I \in \mathcal{F}_u$.
- Intersection: if $X, Y \in \mathcal{F}_u$, then we can find $\mathcal{F}'_X, \mathcal{F}'_Y \in \mathcal{C}$ containing X and Y , respectively. Because \mathcal{C} is a chain, we may find $\mathcal{F}' \in \mathcal{C}$ containing both \mathcal{F}'_X and \mathcal{F}'_Y . Then $X, Y \in \mathcal{F}'$, so $X \cap Y \in \mathcal{F}' \subseteq \mathcal{F}_u$ because \mathcal{F}' is a filter.
- Containment: if $X \in \mathcal{F}_u$ and we have a subset $Y \subseteq I$ containing X , then we find $\mathcal{F}' \in \mathcal{C}$ containing X and find that $Y \in \mathcal{F}' \subseteq \mathcal{F}_u$ because \mathcal{F}' is a filter. ■

1.5.2 Compactness via Ultraproducts

For our application, we will want the notion of an ultraproduct.

Lemma 1.54. Fix a language \mathcal{L} and some \mathcal{L} -structures $\{\mathcal{M}_\alpha\}_{\alpha \in I}$. Now, define an \mathcal{L} -structure \mathcal{M} as follows.

- The universe M is $\prod_{\alpha \in I} M_\alpha$ modded out by the equivalence relation \sim given by $(a_\alpha) \sim (b_\alpha)$ if and only if
$$\{\alpha \in I : a_\alpha = b_\alpha\} \in \mathcal{U}.$$
- Functions are interpreted component-wise.
- For an n -ary relation R , $R^\mathcal{M}((a_{1\alpha}), \dots, (a_{n\alpha}))$ if and only if the set of α such that $R^{M_\alpha}(a_{1\alpha}, \dots, a_{n\alpha})$ is in \mathcal{U} .

Then \mathcal{M} is a well-defined \mathcal{L} -structure.

Proof. Here are our various checks.

- We check that \sim is an equivalence relation.
 - Reflexive: note $(a_\alpha) \sim (a_\alpha)$ because $\{\alpha \in I : a_\alpha = a_\alpha\} = I$ lives in \mathcal{U} .
 - Symmetric: if $(a_\alpha) \sim (b_\alpha)$, then

$$\{\alpha \in I : b_\alpha = a_\alpha\} = \{\alpha \in I : a_\alpha = b_\alpha\},$$

which is in \mathcal{U} by hypothesis.

- Transitive: if $(a_\alpha) \sim (b_\alpha)$ and $(b_\alpha) \sim (c_\alpha)$, then $\{\alpha \in I : a_\alpha = c_\alpha\}$ contains the set

$$\{\alpha \in I : a_\alpha = b_\alpha = c_\alpha\} = \{\alpha \in I : a_\alpha = b_\alpha\} \cap \{\alpha \in I : b_\alpha = c_\alpha\},$$

which lives in \mathcal{U} because \mathcal{U} is a filter.

- We check that interpretation of functions makes sense. Fix an n -ary function f and some elements $(a_{1\alpha}), \dots, (a_{n\alpha})$ and $(b_{1\alpha}), \dots, (b_{n\alpha})$. We must show

$$(f^\mathcal{M}(a_{1\alpha}, \dots, a_{n\alpha})) \sim (f^\mathcal{M}(b_{1\alpha}, \dots, b_{n\alpha})).$$

Well, we note $\{\alpha \in I : f^{\mathcal{M}}(a_{1\alpha}, \dots, a_{n\alpha}) = f^{\mathcal{M}}(b_{1\alpha}, \dots, b_{n\alpha})\}$ contains the set

$$\bigcap_{i=1}^n \{\alpha \in I : a_{i\alpha} = b_{i\alpha}\},$$

which lives in \mathcal{U} because \mathcal{U} is a filter.

- We check that interpretation of relations makes sense. Fix an n -ary function R and some elements $(a_{1\alpha}), \dots, (a_{n\alpha})$ and $(b_{1\alpha}), \dots, (b_{n\alpha})$. We must show

$$R((a_{1\alpha}), \dots, (a_{n\alpha})) \iff R((b_{1\alpha}), \dots, (b_{n\alpha})).$$

Unwrapping the definition of $R^{\mathcal{M}}$, this is equivalent to

$$\{\alpha \in I : R^{M_\alpha}(a_{1\alpha}, \dots, a_{n\alpha})\} \in \mathcal{U} \iff \{\alpha \in I : R^{M_\alpha}(b_{1\alpha}, \dots, b_{n\alpha})\} \in \mathcal{U}.$$

By symmetry, it's enough to show the forward direction, for which we note that the right-hand set contains

$$\{\alpha \in I : R^{M_\alpha}(a_{1\alpha}, \dots, a_{n\alpha})\} \cap \bigcap_{i=1}^n \{\alpha \in I : a_{i\alpha} = b_{i\alpha}\},$$

which lives in \mathcal{U} because \mathcal{U} is a filter. ■

Definition 1.55 (ultraproduct). Fix a language \mathcal{L} and some \mathcal{L} -structures $\{\mathcal{M}_\alpha\}_{\alpha \in I}$. The *ultraproduct* is the \mathcal{L} -structure defined in Lemma 1.54, denoted $\prod_{\alpha \in I} \mathcal{M}_\alpha / \mathcal{U}$ or $\prod_{\mathcal{U}} \mathcal{M}_\alpha$.

We are now ready to begin our proof of Theorem 1.32. We want the following definition.

Definition 1.56 (expansion). Fix a language \mathcal{L} and structure \mathcal{M} . Given a subset $A \subseteq M$, we define the *expansion* \mathcal{L}_A as having the same constants in addition to the constants in A but the same functions and relations.

Remark 1.57. Fix a language \mathcal{L} and structure \mathcal{M} and subset $A \subseteq M$. Then \mathcal{M} is in fact an \mathcal{L}_A -structure, where we interpret the new constants $a \in A$ by $a^{\mathcal{M}} := a$.

Compactness will follow from the result.

Theorem 1.58 (Łoś). Fix a language \mathcal{L} and \mathcal{L} -structures $\{\mathcal{M}_\alpha\}_{\alpha \in I}$. Expand \mathcal{L} to the language $\mathcal{L}' := \mathcal{L}_{\prod_{\alpha \in I} \mathcal{M}_\alpha}$. Now, let \mathcal{U} be an ultrafilter on I so that $\mathcal{M} := \prod_{\mathcal{U}} \mathcal{M}_\alpha$ is an \mathcal{L}' -structure. Then for any \mathcal{L} -formula $\varphi(x_1, \dots, x_n)$ has $\mathcal{M} \models \varphi(a_1^{\mathcal{M}}, \dots, a_n^{\mathcal{M}})$ if and only if

$$\{\alpha \in I : \mathcal{M}_\alpha \models \varphi(a_1, \dots, a_n)\} \in \mathcal{U}.$$

Proof. To see that \mathcal{M} is in fact an \mathcal{L}' -structure, note \mathcal{M} is already an \mathcal{L} -structure, and we may interpret the constant (a_α) of \mathcal{L}' by the corresponding equivalence class in \mathcal{M} . Anyway, the content of the proof is to induct on φ .

- Let c_1 and c_2 be constants. Then $\mathcal{M} \models (c_1 = c_2)$ if and only if $c_1^{\mathcal{M}} = c_2^{\mathcal{M}}$ if and only if the set of α such that $c_1^{M_\alpha} = c_2^{M_\alpha}$ is in \mathcal{U} .
- Let $t(x_1, \dots, x_n)$ be a term and c be a constant. We claim that $\mathcal{M} \models (t = c)(a_1, \dots, a_n)$ if and only if

$$\{\alpha \in I : \mathcal{M}_\alpha \models (t = c)(a_1, \dots, a_n)\} \in \mathcal{U}.$$

This is done by induction on the term t . If t is a constant there is nothing to say. Otherwise, suppose that f is an m -ary function, and we have terms $t_1(x_1, \dots, x_n), \dots, t_m(x_1, \dots, x_n)$. Now, $\mathcal{M} \models (f(t_1, \dots, t_m) = c)(a_1, \dots, a_n)$ if and only if $f^{\mathcal{M}}(t_1^{\mathcal{M}}(\bar{a}), \dots, t_m^{\mathcal{M}}(\bar{a})) = c^{\mathcal{M}}$, which after taking enough intersection is equivalent to having $f^{\mathcal{M}}(c_1^{\mathcal{M}}, \dots, c_m^{\mathcal{M}}) = c^{\mathcal{M}}$ for suitable constants c_\bullet coming from the inductive hypothesis. One can then continue the argument backwards to complete.

- Let $t_1(x_1, \dots, x_n)$ and $t_2(x_1, \dots, x_n)$ be terms. Then $\mathcal{M} \models (t_1 = t_2)(a_1^{\mathcal{M}}, \dots, a_n^{\mathcal{M}})$ if and only if the set of α such that

$$t_1^{\mathcal{M}_\alpha}((a_1^{M_\alpha}), \dots, (a_n^{M_\alpha})) = t_2^{\mathcal{M}_\alpha}((a_1^{M_\alpha}), \dots, (a_n^{M_\alpha}))$$

is contained in \mathcal{U} . Choosing constants c_1 and c_2 suitably as above and using the filter property, this is equivalent to having $c_1^{\mathcal{M}} = c_2^{\mathcal{M}}$, from which we can go backwards to complete the argument.

- The same argument holds for atomic formulae of the form $R(t_1, \dots, t_n)$ where R is an n -ary relation.

We now begin inducting on formulae. Let \mathcal{F}' be the set of desired \mathcal{L}' -formulae. The above checks show that \mathcal{F}' contains atomic formulae.

- Suppose $\varphi, \psi \in \mathcal{F}'$. Then $\mathcal{M} \models (\varphi \wedge \psi)(\bar{a})$ if and only if $\mathcal{M} \models \varphi(\bar{a})$ and $\mathcal{M} \models \psi(\bar{a})$ if and only if

$$\{\alpha \in I : \mathcal{M}_\alpha \models \varphi(\bar{a})\} \cap \{\alpha \in I : \mathcal{M}_\alpha \models \psi(\bar{a})\}$$

lives in \mathcal{U} , which is equivalent to

$$\{\alpha \in I : \mathcal{M}_\alpha \models (\varphi \wedge \psi)(\bar{a})\}$$

by the intersection property of \mathcal{U} .

- Suppose $\varphi \in \mathcal{F}'$. Then $\mathcal{M} \models (\neg\varphi)(\bar{a})$ is false if and only if $\mathcal{M} \models \varphi(\bar{a})$ if and only if

$$\{\alpha \in I : \mathcal{M}_\alpha \models \varphi(\bar{a})\} \in \mathcal{U},$$

which because \mathcal{U} is an ultrafilter is equivalent to

$$I \setminus \{\alpha \in I : \mathcal{M}_\alpha \models \varphi(\bar{a})\} \notin \mathcal{U},$$

from which we can work backwards to complete the argument. (To see the last equivalence, note that each $X \subseteq I$ has exactly one of $X \in \mathcal{U}$ or $I \setminus X \in \mathcal{U}$: at least one is true because \mathcal{U} is an ultrafilter, and at most one is true because both being true requires $\emptyset \in \mathcal{U}$, making \mathcal{U} the trivial filter.)

- Suppose $\varphi(x, \bar{a}) \in \mathcal{F}'$. Then $\mathcal{M} \models (\exists x \varphi(x))(\bar{a})$ if and only if there is some $b \in M$ (i.e., b a constant because we expanded our language) such that $\mathcal{M} \models \varphi(b, \bar{a})$, which is equivalent to

$$\{\alpha \in I : \mathcal{M}_\alpha \models \varphi(b, \bar{a})\} \in \mathcal{U}$$

for some constant b . ■

Corollary 1.59. Let T be a finitely satisfiable \mathcal{L} -theory. Then T is satisfiable.

Proof. We follow [Mar02, Exercise 2.5.20]. We may suppose that T is nonempty. Let I be the set of finite subsets of T , and for each $\Delta \in I$, let \mathcal{M}_Δ be a model for Δ . We have two steps.

1. We define a filter. For each $\varphi \in T$, let $X_\varphi := \{\Delta \in I : \mathcal{M}_\Delta \models \varphi\}$. Then we define

$$D := \{A \in I : A \supseteq X_\varphi \text{ for some } \varphi \in T\}.$$

We show that D is a nontrivial filter on I .

- Note that $\emptyset \notin D$ because this would require that $\emptyset \supseteq X_\varphi$ for some $\varphi \in T$, which is bad because $\mathcal{M}_{\{\varphi\}} \models \varphi$ shows X_φ is nonempty.

- Note any $\varphi \in T$ has $X_\varphi \subseteq I$, so $I \in D$.
 - Intersection: if $A, B \in D$, then find $\varphi, \psi \in T$ such that $X_\varphi \subseteq A$ and $X_\psi \subseteq B$. Then $A \cap B$ contains $X_\varphi \cap X_\psi$, but $X_\varphi \cap X_\psi$ consists of Δ such that \mathcal{M}_Δ models both φ and ψ , which is equivalent to $\mathcal{M} \models \varphi \wedge \psi$, so $X_\varphi \cap X_\psi = X_{\varphi \wedge \psi}$.
 - Containment: if $A \in D$ is contained in $B \subseteq I$, then find $\varphi \in T$ with $A \supseteq X_\varphi$ so that $B \supseteq X_\varphi$ as well.
2. Let \mathcal{U} be an ultrafilter containing D , and let \mathcal{M} be $\prod_{\mathcal{U}} \mathcal{M}_\Delta$. Then for each $\varphi \in T$, we see by Theorem 1.58 that $\mathcal{M} \models \varphi$ if and only if

$$\{\Delta \in I : \mathcal{M}_\Delta \models \varphi\} \in \mathcal{U},$$

which is true by construction of \mathcal{U} . ■

Remark 1.60. Theorem 1.32 was able to bound the size of the model, but the above proof does not. Indeed, the models \mathcal{M}_Δ are potentially large, and \mathcal{M} is approximately the size of all of them multiplied together.

1.6 September 12

We started class by showing that Theorem 1.58 implies the compactness theorem. Professor Scanlon's proof is distinct from the one in my notes, but I have not bothered to record his proof.

1.6.1 Elementary Equivalence

The following notion will be helpful.

Definition 1.61 (theory). Fix a language \mathcal{L} and an \mathcal{L} -structure \mathcal{M} . Then the *theory* $\text{Th}_{\mathcal{L}}(\mathcal{M})$ is the set of sentences φ such that $\mathcal{M} \models \varphi$. For a subset $A \subseteq M$, we may abbreviate $\text{Th}_{\mathcal{L}_A}(\mathcal{M})$ to just $\text{Th}_A(\mathcal{M})$ for brevity.

The following notions are also sometimes helpful.

Definition 1.62 (diagram). Fix a language \mathcal{L} and an \mathcal{L} -structure \mathcal{M} . The *diagram* $\text{Diag}(\mathcal{M})$ is the set φ of atomic \mathcal{L}_M -sentences (in the expanded language \mathcal{L}_M) or negations of atomic sentences such that $\mathcal{M} \models \varphi$. The *elementary diagram* is the theory $\text{Th}_{\mathcal{L}_M}(\mathcal{M}_M)$.

The theory is in some sense everything that a structure can see. As such, we make the following definition.

Definition 1.63 (elementarily equivalent). Fix a language \mathcal{L} . Then two \mathcal{L} -structures \mathcal{M} and \mathcal{N} are *elementarily equivalent*, written $\mathcal{M} \equiv \mathcal{N}$ if and only if $\text{Th}_{\mathcal{L}}(\mathcal{M}) = \text{Th}_{\mathcal{L}}(\mathcal{N})$.

Remark 1.64. In fact, it is enough to merely have $\text{Th}_{\mathcal{L}}(\mathcal{M}) \supseteq \text{Th}_{\mathcal{L}}(\mathcal{N})$. Indeed, suppose for the sake of contradiction that $\text{Th}_{\mathcal{L}}(\mathcal{M}) \not\supseteq \text{Th}_{\mathcal{L}}(\mathcal{N})$. Then there is a sentence φ with $\mathcal{M} \models \varphi$ but \mathcal{N} does not satisfy φ . But then $\mathcal{N} \models \neg\varphi$, so $\mathcal{M} \models \neg\varphi$ too! But this does not make sense because \mathcal{M} cannot satisfy both φ and $\neg\varphi$.

Proposition 1.65. Fix a language \mathcal{L} and isomorphic \mathcal{L} -structures \mathcal{M} and \mathcal{N} . Then \mathcal{M} and \mathcal{N} are elementarily equivalent.

Proof. We show this by induction. Fix an isomorphism $f: \mathcal{M} \rightarrow \mathcal{N}$. We will actually show that $\mathcal{M}_M \equiv \mathcal{N}_M$, where \mathcal{N}_M means \mathcal{M} viewed as an \mathcal{L}_M -structure where the constants $a \in M$ are interpreted as $a^{\mathcal{N}} := f(a)$.

Anyway, we induct on φ .

- Suppose that φ is atomic of the form $t_1(\bar{a}) = t_2(\bar{a})$. If $\mathcal{M}_M \models (t_1(\bar{a}) = t_2(\bar{a}))$, then an induction on terms t shows that

$$t^{\mathcal{N}}(\bar{a}) = f(t^{\mathcal{M}}(\bar{a})).$$

Indeed, if t is a constant term, then this follows directly from f being an isomorphism. Otherwise, t takes the form $g(t_1, \dots, t_n)$ for a function symbol g , and the interpretation of g is also respected by f because it is an isomorphism.

Now, $\mathcal{M}_M \models (t_1(\bar{a}) = t_2(\bar{a}))$ if and only if $t_1^{\mathcal{M}}(\bar{a}) = t_2^{\mathcal{M}}(\bar{a})$, which is equivalent to $\mathcal{N}_M \models (t_1(\bar{a}) = t_2(\bar{a}))$ by passing through f as above.

- Suppose that φ is atomic of the form $R(t_1(\bar{a}), \dots, t_n(\bar{a}))$. Well, $\mathcal{M}_M \models R(t_1(\bar{a}), \dots, t_n(\bar{a}))$ if and only if $(t_1^{\mathcal{M}}(\bar{a}), \dots, t_n^{\mathcal{M}}(\bar{a})) \in R^{\mathcal{M}}$, and then passing everything through f shows that this is equivalent to

$$(t_1^{\mathcal{N}}(\bar{a}), \dots, t_n^{\mathcal{N}}(\bar{a})) \in R^{\mathcal{N}},$$

which is $\mathcal{N}_M \models R(t_1(\bar{a}), \dots, t_n(\bar{a}))$.

- Suppose that φ takes the form $\neg\psi$. Then the usual semantic argument takes care of us.
- Suppose that φ takes the form $\psi \wedge \theta$. Then the usual semantic argument takes care of us.
- Suppose that φ takes the form $\exists x \psi(x)$. Then \mathcal{M}_M models this if and only if there is some $a \in M$ such that $\mathcal{M}_M \models \psi(a)$, but $\psi(a)$ is a perfectly valid sentence in our language because we expanded our constants, so this is equivalent to $\mathcal{N}_M \models \psi(a)$ for some $a \in M$. This last assertion is equivalent to $\mathcal{N}_M \models \exists x \psi(x)$ (the forward direction is clear, and the backward direction is because any $b \in \mathcal{N}$ witnessing takes the form $f(a)$ for some $a \in \mathcal{M}$ because f is a bijection on the universe).

The above induction completes the argument. ■

Proposition 1.65 is a nice result. We might hope for a converse, but it is false in general. There is a converse for finite structures.

Proposition 1.66. Fix a finite language \mathcal{L} and a finite structure \mathcal{M} . Then $\mathcal{M} \equiv \mathcal{N}$ if and only if $\mathcal{M} \cong \mathcal{N}$.

Proof. Say that \mathcal{M} has n elements. Then we build a sentence which asserts that there are exactly n elements x_1, \dots, x_n , and then add on conditions for each m -ary function symbol f what $f(x_{i_1}, \dots, x_{i_m})$ should equal, for each m -ary function symbol R whether $R(x_{i_1}, \dots, x_{i_m})$ should be, and so on.

Let's write this out. The start of this sentence

$$\exists x_1 \cdots \exists x_n \left(\left(\bigwedge_{i \neq j} \neg(x_i = x_j) \right) \wedge \left(\forall y \bigvee_{i=1}^n (y = x_i) \right) \wedge \cdots \right)$$

dictates that any model satisfying this sentence has exactly n elements. (Namely, the first part asserts that the model has at least n elements, and the second bit says that any element equals one of the given n elements.) Next we write in function symbols. Enumerate \mathcal{M} as a_1, \dots, a_n . For each m -ary function symbol f in the language \mathcal{L} , and m elements a_{i_1}, \dots, a_{i_m} of M , we note that $f^{\mathcal{M}}(a_{i_1}, \dots, a_{i_m})$ is some element of M , which by abuse of notation we will write as $a_{\bar{f}(i_1, \dots, i_m)}$. As such, we next tack on the sentence

$$\bigwedge_{m\text{-ary } f} \bigwedge_{1 \leq i_1, \dots, i_m \leq n} \left(f(x_{i_1}, \dots, x_{i_m}) = x_{\bar{f}(i_1, \dots, i_m)} \right).$$

Next up, we interpret constant symbols: by abuse of notation, let $c^{\mathcal{M}}$ be $a_{\bar{c}}$, so we add on the sentence

$$\bigwedge_{c \text{ constant}} (c = x_{\bar{c}}).$$

Lastly, we interpret relations: we need the sentence

$$\bigwedge_{m\text{-ary } R} \bigwedge_{\substack{1 \leq i_1, \dots, i_m \leq n \\ R(a_{i_1}, \dots, a_{i_m})}} R(x_{i_1}, \dots, x_{i_m}).$$

In total, our sentence looks like

$$\begin{aligned} \exists x_1 \cdots \exists x_n & \left(\left(\bigwedge_{i \neq j} \neg(x_i \neq x_j) \right) \wedge \left(\forall y \bigvee_{i=1}^n (y = x_i) \right) \right. \\ & \wedge \bigwedge_{m\text{-ary } f} \bigwedge_{1 \leq i_1, \dots, i_m \leq n} \left(f(x_{i_1}, \dots, x_{i_m}) = x_{\bar{f}(i_1, \dots, i_m)} \right) \\ & \wedge \bigwedge_{c \text{ constant}} (c = x_{\bar{c}}) \\ & \wedge \bigwedge_{m\text{-ary } R} \bigwedge_{\substack{1 \leq i_1, \dots, i_m \leq n \\ R(a_{i_1}, \dots, a_{i_m})}} R(x_{i_1}, \dots, x_{i_m}) \\ & \left. \wedge \bigwedge_{m\text{-ary } R} \bigwedge_{\substack{1 \leq i_1, \dots, i_m \leq n \\ \neg R(a_{i_1}, \dots, a_{i_m})}} \neg R(x_{i_1}, \dots, x_{i_m}) \right). \end{aligned}$$

Let's quickly explain why this works. Notably, \mathcal{M} satisfies the above sentence by taking x_i to be a_i . On the other hand, for any \mathcal{N} which is an \mathcal{L} -structure satisfying the above sentence, the first line dictates that \mathcal{N} must have exactly n elements b_1, \dots, b_n . The second line dictates what $f^{\mathcal{N}}(b_{i_1}, \dots, b_{i_m})$ must equal for each m -ary function symbol f . The third line dictates what $c^{\mathcal{N}}$ for each constant symbol c . Lastly, the last two lines dictate what $R^{\mathcal{N}}(b_{i_1}, \dots, b_{i_m})$ for each m -ary relation symbol R . Thus, we see that we have an isomorphism $\rho: \mathcal{M} \rightarrow \mathcal{N}$ by $a_i \mapsto b_i$.

Writing this out a bit, let's check that ρ preserves function symbols. The other checks are no harder. By construction, we see that

$$\begin{aligned} \rho(f^{\mathcal{M}}(a_{i_1}, \dots, a_{i_m})) &= \rho(a_{\bar{f}(i_1, \dots, i_m)}) \\ &= b_{\bar{f}(i_1, \dots, i_m)} \\ &= f^{\mathcal{N}}(b_{i_1}, \dots, b_{i_m}), \end{aligned}$$

which is what we wanted. Notably, the last equality holds because it was required by our sentence. ■

Remark 1.67. The infinite language case might be an interesting question for the midterm exam. The proof should be quite similar.

Let's verify that infinite structures are not determined by their theories.

Proposition 1.68. Fix a language \mathcal{L} and infinite \mathcal{L} -structure \mathcal{M} . Then there exists an \mathcal{L} -structure \mathcal{N} such that $\mathcal{M} \not\equiv \mathcal{N}$ but $\mathcal{M} \equiv \mathcal{N}$.

Proof. We will choose \mathcal{N} to simply be larger than \mathcal{M} . Choose a cardinal κ strictly larger than $|M|$, and let \mathcal{L}' be an expanded language with κ new constants c_α for each $\alpha \in \kappa$.

We now use compactness to construct \mathcal{N} . Choose the theory T to be

$$\text{Th}_{\mathcal{L}}(\mathcal{M}) \sqcup \{c_\alpha \neq c_\beta : \alpha \neq \beta \text{ for } \alpha, \beta \in \kappa\}.$$

We claim that T is finitely satisfiable. Indeed, for any finite subset Δ , we claim that \mathcal{M} can be made into a model for Δ . Well, \mathcal{M} certainly satisfies $T \cap \Delta \subseteq \text{Th}_{\mathcal{L}}(\mathcal{M})$, and then $\Delta \setminus \text{Th}_{\mathcal{L}}(\mathcal{M})$ is just asserting that \mathcal{M} has some finite number of distinct elements, which is true

More explicitly, let $\lambda \subseteq \kappa$ be a finite subset such that any c_α appearing in a sentence of Δ has $\alpha \in \lambda$. Then choose some element $a_0 \in \mathcal{M}$ and then $|\lambda|$ distinct elements a_α for each $\alpha \in \lambda$. We interpret c_α as a_α for each $\alpha \in \lambda$ and interpret each c_β as a_0 for each $\beta \notin \lambda$. We can see that this new model \mathcal{M}' models Δ , so we are safe.

Anyway, Theorem 1.32 now provides us with a model \mathcal{N}' of T . Notably, \mathcal{N}' can be restricted to an \mathcal{L} -structure by simply forgetting how to interpret the κ new constants, and we see that $\text{Th}_{\mathcal{L}}(\mathcal{N}) \supseteq \text{Th}_{\mathcal{L}}(\mathcal{M})$, so $\mathcal{M} \equiv \mathcal{N}$ follows by Remark 1.64. However, $|\mathcal{N}| \geq \kappa > |\mathcal{M}|$ requires that \mathcal{M} and \mathcal{N} are not isomorphic. ■

Here are some follow-up questions. Fix a language \mathcal{L} .

1. If we have $\mathcal{M} \equiv \mathcal{N}$ and $|\mathcal{M}| = |\mathcal{N}|$, can we construct an example with $\mathcal{M} \not\cong \mathcal{N}$? This is true for some theories $\text{Th}_{\mathcal{L}}(\mathcal{M})$ where this is true but not always. For example, for countable models, this is (roughly speaking) the theory of types.
2. If $\mathcal{M} \equiv \mathcal{N}$, can we find a nonempty index set I and an ultrafilter \mathcal{U} such that $\mathcal{M}^I/\mathcal{U} \cong \mathcal{N}^I/\mathcal{U}$? The converse is certainly true by Theorem 1.58. This forward direction turns out to be yes and is Keisler–Shelah. By the end of the course, we will be able to show this under some assumptions (countable languages, countable structures, and assuming the continuum hypothesis).

1.7 September 14

Today we will prove the Löwenheim–Skolem Theorem.

1.7.1 The Löwenheim–Skolem Theorem

We will want the following definition.

Definition 1.69 (elementary substructure). Fix a language \mathcal{L} and two structures \mathcal{M} and \mathcal{N} . Then we say that \mathcal{M} is an *elementary substructure* of \mathcal{N} , written $\mathcal{M} \leq \mathcal{N}$ if and only if \mathcal{M} is a substructure of \mathcal{N} and $\mathcal{M}_M \equiv \mathcal{N}_M$.

Remark 1.70. It is not enough to have $\mathcal{M} \subseteq \mathcal{N}$ and $\mathcal{M} \equiv \mathcal{N}$. For example, take the language $\mathcal{L} = \{<\}$ and let $\mathcal{M} = (\mathbb{N}, <)$ and $\mathcal{N} = (\mathbb{Z}^+, <)$. Then $\mathcal{M} \subseteq \mathcal{N}$, and $\mathcal{M} \equiv \mathcal{N}$. To see that $\mathcal{M} \equiv \mathcal{N}$ because $\mathcal{M} \cong \mathcal{N}$ (subtracting one is an isomorphism $\mathbb{Z}^+ \rightarrow \mathbb{N}$), which is enough by Proposition 1.65. However, $\mathcal{M} \not\leq \mathcal{N}$: the sentence $\forall x \, 1 \leq x$ is true in \mathcal{M} but not in \mathcal{N} .

Here is the result we are going to show.

Theorem 1.71. Fix a language \mathcal{L} and infinite structure \mathcal{M} . For all subsets $A \subseteq M$, there exists an elementary substructure $\mathcal{N} \leq \mathcal{M}$ containing A with $|\mathcal{N}| = |A| + |\mathcal{L}| + \aleph_0$.

Proof. We essentially do a more careful version of the Henkin construction. Set $T := \text{Th}(\mathcal{M}_A)$. Let \mathcal{L}' and T' be the language and theory extending \mathcal{L} and T (respectively) obtained from the construction in Lemma 1.40 by adding witnessing constants. Quickly, we recall that T' and \mathcal{L}' are constructed inductively as follows.

- Set $T_0 := T$ and $\mathcal{L}_0 := \mathcal{L}$.

- Set \mathcal{L}_{n+1} to be \mathcal{L}_n with a constant c_φ for each \mathcal{L}_n -formula φ with a variable x , and then we add $\exists\varphi(x) \rightarrow \varphi(c_\varphi)$ to T' . The function and relation symbols are the same between \mathcal{L}_n and \mathcal{L}_{n+1} .
- Lastly, \mathcal{L}' is the union of the \mathcal{L}_n s, and T' is the union of the T_n s.

We now expand \mathcal{M} to be a model \mathcal{M}' of T' . One only has to deal with the constants added by \mathcal{L}' . We will do this inductively.

- Set $\mathcal{M}_0 := \mathcal{M}_A$, and we construct \mathcal{M}_n to model T_n .
- Given $\mathcal{M}_n \models T_n$, we construct \mathcal{M}_{n+1} to be an \mathcal{L}_{n+1} -structure as follows. Well, we only need to worry about interpreting the new constants c_φ where φ is an \mathcal{L}_n -formula with free variable x , and we interpret $c_\varphi^{\mathcal{M}_{n+1}}$ as some $a_\varphi \in \mathcal{M}_n$ if $\mathcal{M}_n \models \varphi(a_\varphi)$ if such some a_φ exists, and we set $c_\varphi^{\mathcal{M}_{n+1}}$ to be any element of \mathcal{M}_n if no such a_φ exists.
Then \mathcal{M}_{n+1} certainly satisfies everything in T_n (by inductive hypothesis), and it satisfies every one of the new sentences $\exists x\varphi(x) \rightarrow \varphi(c_\varphi)$ by construction of $c_\varphi^{\mathcal{M}_{n+1}}$, so we conclude $\mathcal{M}_{n+1} \models T_{n+1}$, as needed.
- Lastly, we define \mathcal{M}' to be the union of the \mathcal{M}_n , and we conclude our construction. One can see that $\mathcal{M}' \models T'$ directly from the construction of the previous step because any $\varphi \in T'$ belongs to some T_n for finite n .

To continue the proof, we want the following result to check that we have built an elementary substructure.

Lemma 1.72 (Tarski–Vaught test). Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Call A “realizable” if and only if any \mathcal{L} -formula $\varphi(x_1, \dots, x_n, y)$ and n -tuple $\bar{a} \in A^n$ has $\mathcal{M} \models (\exists y\varphi(\bar{x}, y))(\bar{a})$ if and only if there is some $b \in A$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$. Then A is realizable if and only if there is an elementary substructure $\mathcal{A} \leq \mathcal{M}$ with universe A .

Proof. There is some content here because the assertion $\mathcal{M}_A \equiv \mathcal{A}_A$ does not even make sense without having constructed \mathcal{A} . Anyway, we have two implications to show.

- Suppose that A is the universe of an elementary substructure $\mathcal{A} \leq \mathcal{M}$. We want to show that \mathcal{A} is realizable. Well, let $\varphi(x_1, \dots, x_n, y)$ be an \mathcal{L} -formula, and choose some $\bar{a} \in A^n$. Now, $\mathcal{M} \models (\exists y\varphi(\bar{x}, y))(\bar{a})$ if and only if $\mathcal{M}_A \models \exists y\varphi(\bar{a}, y)$. Now, because $\mathcal{A} \leq \mathcal{M}$, this is equivalent to $\mathcal{A}_A \models \exists y\varphi(\bar{a}, y)$, which is equivalent to having some $b \in A$ such that $\mathcal{A}_A \models \varphi(\bar{a}, b)$, which is equivalent to $\mathcal{M}_A \models \varphi(\bar{a}, b)$, which means there is $b \in A$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$.
- Suppose A is realizable. The main content here is to check that A is the universe of an \mathcal{L} -substructure of \mathcal{M} . We have the following checks.
 - Certainly $A \subseteq M$.
 - For each constant symbol c , we need $c^{\mathcal{M}} \in A$. Well, look at the formula $\varphi(y)$ given by $y = c$. Then $\mathcal{M} \models \exists y\varphi(y)$ by $c^{\mathcal{M}}$, so being realizable grants some $b \in A$ such that $\mathcal{M} \models \varphi(b)$, which means $c^{\mathcal{M}} = b \in A$, as needed.
 - For each n -ary function symbol $f(x_1, \dots, x_n)$ and $\bar{a} \in A$, we need to check $f^{\mathcal{M}}(\bar{a}) \in A$. Well, look at the formula $\varphi(x_1, \dots, x_n, y)$ which is $y = f(x_1, \dots, x_n, y)$. Then $\mathcal{M} \models \exists y\varphi(\bar{a})$, so being realizable promises some $b \in A$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$, which is asserting $f(a_1, \dots, a_n) = b$.

We now need to show $\mathcal{M}_A \equiv \mathcal{A}_A$. We induct to show that an \mathcal{L}_A -sentence ψ has $\mathcal{M}_A \models \psi$ if and only if $\mathcal{A}_A \models \psi$. Let \mathcal{F}' be the set of such \mathcal{L}_A -sentences.

- For atomic formulae, we use Proposition 1.12 so that we don’t have to do any more work.
- The usual arguments tell us that $\varphi, \psi \in \mathcal{F}'$ implies that $\neg\varphi \in \mathcal{F}'$ and $\varphi \wedge \psi \in \mathcal{F}'$. We won’t write this out.

- Lastly, suppose ψ is of the form $\exists y\varphi(y)$. Because $\exists y\varphi$ is an \mathcal{L}_A -sentence, we can write $\varphi(y)$ as $\varphi'(\bar{a}, y)$ where $\varphi'(x_1, \dots, x_n, y)$ is some \mathcal{L} -formula and $\bar{a} \in A^n$.

Now, in one direction, $\mathcal{A}_A \models \psi$ if and only if some $b \in A$ such that $\mathcal{A}_A \models \varphi(a)$, so by induction $\mathcal{M}_A \models \psi(b)$, which implies $\mathcal{M} \models \psi$, as needed.

To go the other direction, we need to pull a witness down from \mathcal{M} to \mathcal{A} , which is harder. Suppose $\mathcal{M}_A \models \psi$. Then $\mathcal{M}_A \models (\exists y\varphi'(x, y))(\bar{a})$, from which being realizable grants $b \in A$ such that $\mathcal{M}_A \models \varphi'(\bar{a}, b)$. This sentence is simpler, so by induction we get $\mathcal{A} \models \varphi'(\bar{a}, b)$, which is equivalent to $\mathcal{A} \models \exists y\varphi(y)$, as needed. ■

Remark 1.73. There is not really anything to do when checking the reverse direction of being realizable: having $b \in A$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$ of course implies that $\mathcal{M} \models (\exists y\varphi(\bar{x}, y))(\bar{a})$ by choosing y to be this $b \in A$. The content is the reverse direction where we pull down the witness from \mathcal{M} to \mathcal{A} .

Now, let the set N be the set of interpretations of constant symbols $c^{\mathcal{M}'}$ for each constant symbol c of \mathcal{L}' . Notably, $A \subseteq \mathcal{L}'$, and $a^{\mathcal{M}'} = a$, so $a \in N$, so $A \subseteq N$. We would like to turn N into an elementary substructure, for which we use Lemma 1.72.

It suffices to check that N is realizable. Let $\varphi(x_1, \dots, x_n, y)$ be an \mathcal{L} -formula and $(a_1, \dots, a_n) \in N^n$. Suppose $\mathcal{M} \models (\exists y\varphi(\bar{x}, y))(\bar{a})$. Then $\mathcal{M}' \models (\exists y\varphi(\bar{x}, y))(\bar{a})$ by choosing the same y , which means $\mathcal{M}' \models \varphi(\bar{a}, y)$, but $\mathcal{M}' \models \exists y\varphi(\bar{a}, y) \rightarrow \varphi(\bar{a}, c)$ for some constant symbol c of \mathcal{L}' . Combining, we get $\mathcal{M}' \models \varphi(\bar{a}, c)$. But then setting $d := c^{\mathcal{M}'}$ (which lives in N !), we achieve $\mathcal{M}' \models \varphi(\bar{a}, d)$.

Thus, N is the universe of some elementary substructure $\mathcal{N} \leq \mathcal{M}$. We saw that N contains A , and we see $|N|$ is at most the size of the constants of \mathcal{L}' , which has size $|\mathcal{L}| + \aleph_0 + |A|$. This completes the proof. ■

One can also go up, which was essentially Proposition 1.68.

Proposition 1.74. Fix an infinite \mathcal{L} -structure \mathcal{M} . For any cardinal $\kappa \geq |M| + |\mathcal{L}|$, there exists an \mathcal{L} -structure \mathcal{N} with cardinality κ and $\mathcal{M} \leq \mathcal{N}$.

Proof. As in Proposition 1.68, let \mathcal{L}' be the language \mathcal{L} where we add constants c_α for each $\alpha \in \kappa$, and then we let T' be

$$\text{Th}(\mathcal{M}_M) \sqcup \{c_\alpha \neq c_\beta : \alpha \neq \beta \text{ for } \alpha, \beta \in \kappa\}.$$

We showed in Proposition 1.68 that T' is finitely satisfiable, so we produce a model \mathcal{N}_0 of T' . Now, let A be the set of interpretations of constants $c^{\mathcal{N}_0}$ for each constant c in \mathcal{L}' . Notably, A contains M , and the map $\kappa \rightarrow A$ given by $\alpha \mapsto c_\alpha^{\mathcal{N}_0}$ is one-to-one, so $|A| \geq \kappa$. On the other hand, $|A|$ has size bounded by the constants of \mathcal{L}' , which has size $\kappa + |M| + |\mathcal{L}|$, which is κ , so $|A|$ has size exactly κ .

Now, by Theorem 1.71, we produce an elementary substructure $\mathcal{N} \leq \mathcal{N}_0$ containing A . Because $\mathcal{M} \subseteq \mathcal{N} \leq \mathcal{N}_0$ and $\mathcal{M} \leq \mathcal{N}_0$ (by construction of \mathcal{N}_0), so we conclude $\mathcal{M} \leq \mathcal{N}$ by chasing our formulae around. ■

1.8 September 19

Here we go.

1.8.1 An Example of the Back-and-Forth Method

For our example, let \mathcal{L} be a language with one binary relation E , which will be considered to be an equivalence relation. Consider the structure \mathcal{M}_0 with universe $(x, y) \in \mathbb{N}^2$ where $x < y$, where $(x, y)E(x', y')$ if and only if $y = y'$.

We claim that there is another countable model with the same theory. For example, we consider \mathcal{M}_ω which is \mathcal{M}_0 with a disjoint copy of $\mathbb{N}^2 \times \{0\}$ where $(x, y, 0)E(x', y', 0)$ if and only if $y = y'$. Let's check that

the theory of \mathcal{M}_0 has the same theory of \mathcal{M}_ω . This essentially follows from compactness (Theorem 1.32) and Theorem 1.71 to the theory T consisting of the elementary diagram of \mathcal{M}_0 plus the sentences

$$\{c_{xy} \neq c_{x'y'} : \text{for } (x, y) \neq (x', y')\} \cup \{c_{xy} E c_{x'y} : x, x', y \in \mathbb{N}\} \cup \{c_{xy} E c_{x'y'} : x, x', y, y' \in \mathbb{N} \text{ where } y \neq y'\},$$

where we have introduced these new constants c_{xy} to an extended language \mathcal{L}' . Namely, Theorem 1.32 permits us to find a countable model of this above theory: to see that the above set of sentences is satisfiable, we note that \mathcal{M}_0 is able to model any finite subset of the above theory is only asking for arbitrary many arbitrarily large equivalence classes, which \mathcal{M}_0 provides.

So we produce a countable model \mathcal{M}' of T . We claim that $\mathcal{M}' \cong \mathcal{M}_\omega$ in the language \mathcal{L} . This will use the back-and-forth method.

Lemma 1.75. Fix everything as above. Then $\mathcal{M}' \cong \mathcal{M}_\omega$, where \mathcal{M}' is considered as an \mathcal{L} -structure.

Proof. We build our isomorphism via approximations $f_i: X_i \rightarrow Y_i$ for $i \in \mathbb{N}$, where $X_i \subseteq \mathcal{M}'$ and $Y_i \subseteq \mathcal{M}_\omega$. We require that $i \leq j$ means $X_i \subseteq X_j$ and $Y_i \subseteq Y_j$ and then $f_j|_{X_i} = f_i$, and we also want f_i to be an isomorphism of \mathcal{L} -structures for $i > 0$. By the end of this process, we will want $\bigcup_{i \in \mathbb{N}} X_i = \mathcal{M}'$ and $\bigcup_{i \in \mathbb{N}} Y_i = \mathcal{M}_\omega$ so that we have a well-defined isomorphism $f: \mathcal{M}' \rightarrow \mathcal{M}_\omega$ at the end. This last bit is going to be a little tricky. For this, we enumerate $\mathcal{M}' = \{a_i\}_{i=0}^\infty$ and $\mathcal{M}_\omega = \{b_i\}_{i=0}^\infty$, and we will ask that each n have $\{a_j : j < n\} \subseteq X_{2n}$ and $\{b_j : j < n\} \subseteq \text{im } f_{2n+1}$.

Alright, let's get started. Take f_0 to be the unique function $X_0 \rightarrow Y_0$ where $X_0 = Y_0 = \emptyset$. One can check that this trivially works for all of our hypotheses. We now induct in two cases.

- Suppose we have $f_{2n}: X_{2n} \rightarrow Y_{2n}$, and we want to produce $f_{2n+1}: X_{2n+1} \rightarrow Y_{2n+1}$. The point is that b_n now needs to appear in the range of f_{2n+1} . We have the following cases.
 - If b_n is already in the range, do nothing. In the following cases, we suppose that b_n is not in the range of f_{2n} already.
 - Suppose that b_n is not equivalent to some element of $\text{im } f_{2n}$. If b_n is in a finite equivalence class, map it to the corresponding unique equivalence class in \mathcal{M}' , which cannot have been chosen so far because f_{2n} is an isomorphism. If b_n lives in an infinite equivalence class, then go find an unused infinite equivalence class in \mathcal{M}' , which is possible because f_{2n} has finite domain currently.
 - Suppose that b_n is equivalent to some element $b' \in \text{im } f_{2n}$. By the nature of f_\bullet being an isomorphism, we are arranging so that the size of the equivalence class of a and $f_\bullet(a)$ are always the same. So the size of the equivalence class of $f_{2n}^{-1}(b')$ must have space (even if finite!) because the element of b_n not being hit so far requires us to have space in the equivalence class of $f_{2n}^{-1}(b')$.
- Going forward the argument is essentially the same just talking in reverse.

Assembling the f_\bullet together produces the desired result. ■

We now conclude by remarking that $\text{Th}_{\mathcal{L}, \mathcal{M}_0}(\mathcal{M}_0) = \text{Th}_{\mathcal{L}, \mathcal{M}_0}(\mathcal{M}_\omega)$, so $\mathcal{M}_0 \leq \mathcal{M}_\omega$.

Remark 1.76. We can now define $\mathcal{M}_n := \mathcal{M}_0 \sqcup \mathbb{N} \times \{0, 1, \dots, n-1\} \times \{0\}$ as a substructure of \mathcal{M}_ω . One can repeat the above argument with \mathcal{M}_0 replaced by \mathcal{M}_n to conclude that $\mathcal{M}_n \leq \mathcal{M}_\omega$ again. We conclude that $\mathcal{M}_0 \equiv \mathcal{M}_n$ for each n . In total, we have produced countably many non-isomorphic models. It turns out that these are all the countable ones.

One might now go back and ask for the number of models of $\text{Th}_{\mathcal{L}, \mathcal{M}_0}(\mathcal{M}_0)$ of cardinality \aleph_1 . It turns out that there are again countably many. The point is that a model \mathcal{M} of cardinality \aleph_1 can be attached the two invariants

$$\begin{aligned} \kappa_0(\mathcal{M}) &:= \#\{[x] \in M/E : [x]_E \text{ has size } \aleph_0\}, \\ \kappa_1(\mathcal{M}) &:= \#\{[x] \in M/E : [x]_E \text{ has size } \aleph_1\}. \end{aligned}$$

One can show that $\mathcal{M}_1 \cong \mathcal{M}_2$ if and only if $\kappa_0(\mathcal{M}_1) = \kappa_0(\mathcal{M}_2)$ and $\kappa_1(\mathcal{M}_1) \cong \kappa_1(\mathcal{M}_2)$ by using some set theory, and then one can produce a model with given invariants κ_0 and κ_1 arbitrarily provided that $\aleph_0 \kappa_0 + \aleph_1 \kappa_1 = \aleph_1$.

1.8.2 Dense Linear Orders Without Endpoints

Let's see another example.

Proposition 1.77. Fix a language \mathcal{L} with a single binary relation $<$. Then $\text{Th}_{\mathcal{L}}(\mathbb{Q}, <)$ is \aleph_0 -categorical.

We should perhaps define \aleph_0 -categorical.

Definition 1.78 (κ -categorical). A theory T of a language \mathcal{L} is κ -categorical if and only if T has exactly one isomorphism class of models of cardinality κ .

In fact, we will show the following.

Proposition 1.79. Fix a language \mathcal{L} with a single binary relation $<$, and let DLO be the following theory, of dense linear orders without endpoints.

- $<$ is a total ordering.
- Dense: $\forall x \forall y (x < y \rightarrow \exists z (x < z \wedge z < y))$.
- Without endpoints: $\forall x \exists y (y < x)$ and $\forall x \exists y (x < y)$.

Then DLO is \aleph_0 -categorical.

Note that \mathbb{Q} models DLO, so Proposition 1.77 will follow. Anyway, let's show Proposition 1.79.

Proof of Proposition 1.79. Let \mathcal{A} and \mathcal{B} be models of DLO. Enumerate $\mathcal{A} = \{a_i\}_{i=0}^{\infty}$ and $\mathcal{B} = \{b_i\}_{i=0}^{\infty}$, and we will work in the same set-up as the back-and-forth argument previously described. Namely, we describe a sequence of compatible isomorphisms $f_i: X_i \rightarrow Y_i$ where X_{2n} contains $\{a_1, \dots, a_{n-1}\}$ and Y_{2n+1} contains $\{b_1, \dots, b_{n-1}\}$. Take f_0 to be the unique function $\emptyset \rightarrow \emptyset$.

- Suppose we have f_{2n-1} , and we want to build f_{2n} . If a_n is already in the domain of f_{2n-1} , do nothing. We have three cases.
 - If $a_n < x$ for all $x \in X_{2n-1}$, use that \mathcal{B} has no endpoints to find $f(a_n)$ less than everyone in Y_{2n-1} .
 - If $a_n > x$ for all $x \in X_{2n-1}$, make a similar argument as the previous case.
 - Otherwise, find $x, y \in X_{2n-1}$ so that $x < a_n < y$, and nothing in X_{2n-1} lives between x and y ; this is possible because $<$ is a total ordering. Then use the density of \mathcal{B} to find some $f(a_n)$ strictly between x and y to complete.
- To extend f_{2n} to f_{2n+1} , repeat the above argument in reverse.

Now, assembling our f_\bullet produces our isomorphism. ■

Remark 1.80. We now might ask how many models DLO has of cardinality \aleph_1 . There are apparently 2^{\aleph_1} many up to isomorphism. Of course, this is an upper bound on the number of models because an ordering is asking for a subset of $\aleph_1 \times \aleph_1$. So the name of the game now is to produce enough models; one cannot really hope to precisely describe all the models.

Example 1.81. It is not too hard to provide two models of DLO of cardinality \aleph_1 which are not isomorphic. We take $\mathcal{M} := \mathbb{R}$ and $\mathcal{N} := \mathbb{R}_1 \sqcup \mathbb{R}_2$. Here, \mathbb{R}_i is a copy of \mathbb{R} where every element of \mathbb{R}_2 is greater than any element of \mathbb{R}_1 , and for any $r \in \mathbb{R}$, we will write r_i for the copy of r in \mathbb{R}_i .

Now, suppose for contradiction there is an isomorphism $\varphi: \mathcal{N} \rightarrow \mathcal{M}$. The point is that \mathcal{N} has a sequence which “goes to infinity” (in \mathbb{R}_1) which still has an upper bound (by \mathbb{R}_2). To rigorize, we have the following steps.

1. Consider the sequence $\varphi(1_1), \varphi(2_1), \dots$ in \mathbb{R} , where $1_1, 2_1, \dots$ are their copies in \mathbb{R}_1 . This is a strictly increasing sequence in \mathbb{R} , and it is bounded above by (say) $\varphi(0_2)$. Thus, noting that the order topology on \mathbb{R} is just the usual topology, our sequence must converge to some $\beta \in \mathbb{R}$. Say $\alpha \in \mathcal{N}$ has $\varphi(\alpha) = \beta$. The point is that this α must lie “between” \mathbb{R}_1 and \mathbb{R}_2 .
2. Indeed, note $\alpha > r_1$ for any $r_1 \in \mathbb{R}_1$ because $\varphi(\alpha) > \varphi(r_1)$. Namely, select any integer $n_1 > r_1$, and we have $\varphi(\alpha) \geq \varphi(n_1) > \varphi(r_1)$.

On the other hand, we claim $\alpha \leq r_2$ for any $r_2 \in \mathbb{R}_2$. Indeed, otherwise we have $\alpha > r_2$ and thus $\varphi(\alpha) > \varphi(r_2)$, so use $\varphi(n_1) \rightarrow \varphi(\alpha)$ as $n_1 \rightarrow \infty$ to find n_1 such that $\varphi(\alpha) > \varphi(n_1) > \varphi(r_2)$. But then $n_1 > r_2$, which contradicts the ordering on \mathcal{N} .

The second step above has produced $\alpha \in \mathcal{N}$ bigger than anything in \mathbb{R}_1 and less than anything in \mathbb{R}_2 , which is a contradiction.

To wrap us up, let’s pick up the following result.

Proposition 1.82. Fix an \mathcal{L} -theory T which is κ -categorical for cardinality κ . If T has only infinite models, then T is complete; i.e., any \mathcal{L} -sentence φ has either $T \models \varphi$ or $T \models \neg\varphi$.

Proof. Let \mathcal{M} be a model of T of cardinality κ . Now, for any sentence φ , if $T \models \varphi$ and $T \models \neg\varphi$, then there is a model \mathcal{M}_+ and \mathcal{M}_- satisfying $T \cup \{\varphi\}$ and $T \cup \{\neg\varphi\}$, respectively. By Theorem 1.71, we may bring \mathcal{M}_+ and \mathcal{M}_- to have cardinality κ , so being κ -categorical requires $\mathcal{M}_+ \cong \mathcal{M}_-$, which is a contradiction because then $\mathcal{M}_+ \equiv \mathcal{M}_-$. ■

Example 1.83. Thus, Proposition 1.79 requires that DLO is complete. As such, the theory DLO must complete to exactly $\text{Th}(\mathbb{Q}, <)$.

THEME 2

ELIMINATING QUANTIFIERS

Freedom is just another word for nothing left to lose.

—Janis Joplin, [Jop10]

2.1 September 21

Today, we will go on to some more nontrivial examples.

2.1.1 Algebraically Closed Fields

Consider the language \mathcal{L} with binary operations $+$ and \cdot , a unary operation $-$, and constants 0 and 1. The theory of fields has the sentences given by the ones in a standard algebra class.

- $\forall x \forall y ((x + y = y + x) \wedge (x \cdot y = y \cdot x)).$
- $\forall x \forall y \forall z (((x + y) + z = x + (y + z)) \wedge ((x \cdot y) \cdot z = x \cdot (y \cdot z))).$
- $\forall x ((x + (-x) = 0) \wedge ((-x) + x = 0)).$
- $\forall x \exists y (x \cdot y = 1).$
- $\forall x \forall y \forall z (x \cdot (y + z) = x \cdot y + x \cdot z).$
- $\forall x ((x + 0 = x) \wedge (x \cdot 1 = x)).$
- $\neg(0 = 1).$

To make this algebraically closed, we want every monic polynomial to have a root. For this, we should go degree-by-degree. For example, for degree d which is a positive integer, we write the sentence φ_d to be

$$\forall a_1 \cdots \forall a_{d-1} \exists x (x^d + a_{d-1}x^{d-1} + \cdots + a_1x + a_0 = 0).$$

Call this theory ACF. Notably, we then have used infinitely many axioms.

As an aside, we note there is no finite set of sentences characterizing algebraically closed fields. Let's show this.

Lemma 2.1. Suppose a satisfiable theory T is finitely axiomatizable: there is a finite set of sentences $\varphi_1, \dots, \varphi_n$ such that $\mathcal{M} \models T$ for a structure \mathcal{M} if and only if $\mathcal{M} \models \varphi_i$ for each φ_i . Then there is a finite subset $T_0 \subseteq T$ such that $\mathcal{M} \models T$ if and only if $\mathcal{M} \models T_0$.

Proof. The reverse direction is clear by just taking T_0 to be our finite set of axioms.

In the other direction, suppose that $\varphi := \varphi_1 \wedge \dots \wedge \varphi_n$ axiomatizes T . We now apply compactness $\Sigma := T \cup \{\neg\varphi\}$. Note Σ is not satisfiable because $\mathcal{M} \models T$ if and only if $\mathcal{M} \models \varphi$. Thus, by Theorem 1.32, we see that Σ cannot be finitely satisfiable. But T is finitely satisfiable, so there is some finite subset of the form $T_0 \cup \{\neg\varphi\}$ which is not satisfiable.

We now check that T_0 does the trick. However, this means that any structure \mathcal{M} such that $\mathcal{M} \models T_0$ requires $\mathcal{M} \models \varphi$, and conversely, $\mathcal{M} \models \varphi$ implies $\mathcal{M} \models T$ implies $\mathcal{M} \models T_0$. Thus, T_0 is the needed subset. ■

Let's apply this lemma to ACF. Let T_0 be some finite subset of ACF, and we show that T_0 is not equivalent to ACF. Add in any of the field axioms necessary, and we know there is some upper bound N such that T_0 is then contained in the field axioms plus $\{\varphi_1, \dots, \varphi_d\}$. To show that T_0 is not equivalent to ACF, we construct a field K/\mathbb{Q} which models T_0 but not ACF. Well, construct K by a tower

$$\mathbb{Q} = K_0 \subseteq K_1 \subseteq \dots,$$

where K_{n+1} consists of all numbers which are roots of polynomials in K_n of degree at most N . Then set $K := \bigcup_{n=0}^{\infty} K_n$, and we see $K \models T_0$.

Well, for a piece of algebra, we note that the polynomial $f_p(x) := x^p - 2 \in \mathbb{Q}[x]$ is irreducible for any prime p . Choosing $p > N$, we then claim that $f_p(x) \in K[x]$ has no root. Indeed, any root would need to live in some $K_{n+1}[x]$, which means that $x^p - 2$ has some root shared with a polynomial of degree at most N whose coefficients live in K_n . However, extracting out the necessary coefficients into a field L , we see that L/\mathbb{Q} has degree coprime to p (it's constructed using roots of polynomials of degrees at most N , and $p > N$ is prime), but then $\mathbb{Q}[x]/(x^p - 2) \subseteq L$ has degree p , so it cannot possibly be a subfield.

Remark 2.2. The same argument shows that one can finitely axiomatize fields of characteristic 0. We produce the theory of characteristic-0 fields by adding in the sentences

$$\underbrace{1 + \dots + 1}_p = 0$$

for each positive prime p . But then no finite subset of these axioms will do because there are fields of arbitrarily large (but still finite) characteristic.

Anyway, here is our theorem.

Theorem 2.3. The completion of ACF are the theories ACF_p where p is a prime or zero, where ACF_p adds in the condition of being characteristic p (via the sentence $1 + \dots + 1 = 0$ for nonzero p and $1 + \dots + 1 \neq 0$ for all lengths when $p = 0$).

In fact, we will show the following stronger result.

Theorem 2.4. Fix p to be prime or zero. Then ACF_p is κ -categorical for any $\kappa > \aleph_0$.

This will be enough to prove Theorem 2.3 by Proposition 1.82 because ACF_p certainly has models of size $\kappa > \aleph_0$ by taking $\overline{k(\kappa)}$ where κ is being used as a transcendence basis. Notably, $\overline{k(\kappa)}$ has size $\kappa + \aleph_0 = \kappa$.

Anyway, let's prove Theorem 2.4.

Proof with algebra. Let k be the smallest field of that characteristic (the finite field when $p > 0$ and \mathbb{Q} when $p = 0$).

Now, suppose we have two fields K_1 and K_2 which satisfy ACF_p of cardinality κ . Now, let $X_i \subseteq K_i$ be a transcendence basis for each i , meaning that X_i is a maximal algebraically independent set of elements. As such, K_i is algebraic over $\mathbb{F}_p(X_i)$. Now, $|k(X_i)| = |X_i| + \aleph_0$, so taking algebraic closure has $\kappa = |K_i| = |k(X_i)| + \aleph_0 = |X_i| + \aleph_0$, so $\kappa = |X_i|$. Thus, $k(X_1) \cong k(X_2)$, so taking algebraic closure enforces $K_1 \cong K_2$ by taking algebraic closure. ■

Corollary 2.5. Let \mathcal{U} be a non-principal ultrafilter on \mathcal{P} . Then we have a field isomorphism

$$\mathbb{C} \cong \prod_{\mathcal{U}} \overline{\mathbb{F}_p}.$$

Proof. By Theorem 1.58, we see that $\prod_{\mathcal{U}} \overline{\mathbb{F}_p}$ is algebraically closed because being algebraically closed field is held in each factor of the ultraproduct. It remains to compute the characteristic. Well, the sentence $1 + \dots + 1 = 0$ for any length p fails to hold in all but finitely many of these factors, so we see that the sentence

$$\underbrace{1 + \dots + 1}_p \neq 0$$

holds in all but finitely many of the factors of our ultrafilter. Thus, the ultraproduct has characteristic 0 by Theorem 1.58 again, and we see that \mathbb{C} has the same cardinality as our ultrafilter, so the result follows by Theorem 2.4. To compute this cardinality, we note that

$$\left| \prod_{\mathcal{U}} \overline{\mathbb{F}_p} \right| \leq \aleph_0^{\aleph_0} = 2^{\aleph_0}.$$

One can then embed this ultraproduct into a tree; one uses Theorem 1.58. More generally, one we will be able to show that $|X_i| \geq 2^i$ for some collection $\{X_i\}_{i \in \mathbb{N}}$ has $\prod_{\mathcal{U}} X_i$ of cardinality 2^{\aleph_0} . ■

Let's improve our proof of Theorem 2.4. We will show the following stronger result.

Theorem 2.6. The theory ACF eliminates quantifiers. In other words, for any formula $\varphi(x_1, \dots, x_n)$, there is a quantifier-free formula $\psi(x_1, \dots, x_n)$ such that $\text{ACF} \models \forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \psi(\bar{x}))$.

Remark 2.7. The theory of Peano arithmetic does not eliminate quantifiers: there are very complicated sets that one can define.

There is a proof in Tarski's RAND paper. We are not going to follow it. We are going to do a back-and-forth argument. To begin, we have the following step.

Proposition 2.8. Fix two algebraically closed fields K_1 and K_2 of cardinality $\kappa > \aleph_0$. Suppose, we have an isomorphism $f: L_1 \rightarrow L_2$ of subfields $L_1 \subseteq K_1$ and $L_2 \subseteq K_2$ where L_1 and L_2 are subfields of cardinality less than κ . Then f extends to an isomorphism $K_1 \rightarrow K_2$.

Proof. We construct this isomorphism using a back-and-forth argument. Treat κ as an ordinal, and enumerate $K_1 = \{a_\alpha : \alpha \in \kappa\}$ and $K_2 = \{b_\alpha : \alpha \in \kappa\}$. We will build a sequence of isomorphisms $g_\alpha: L_\alpha^1 \rightarrow L_\alpha^2$ for each $\alpha \in \kappa$ so that g_β extends g_α whenever $\alpha \leq \beta$. We will also arrange so that $g_0 := f$ and $a_\beta \in L_\alpha^1$ and $b_\beta \in L_\alpha^2$ for each $\beta \in \alpha$; it will also help to have L_α^* always have cardinality less than κ . If we can do this, we simply define $g: K_1 \rightarrow K_2$ by taking the union of all these isomorphisms.

For g_0 , there is nothing to do. If α is a limit ordinal, then take g_α to be the union of the g_β for $\beta < \alpha$. Notably, the domain and codomain are the unions of the domains and codomains; of course, this is still an isomorphism, and it satisfies our necessary property because any $\beta < \alpha$ has a_β and b_β in the domain and

codomain of $g_{\beta+1}$, respectively. Lastly, the domain and codomain is an ascending union of sets of cardinality less than κ , which is typically less than κ .¹

In our last case, take $\alpha := \beta + 1$. Then we need to tell g_β where to send a_β . If a_β is already in the domain, do nothing. Otherwise, there are two cases.

- Suppose that a_β is algebraic over L_β^1 with monic irreducible polynomial $P(x)$. Passing through g_β , we see that $g_\beta(P(x)) \in L_\beta^2[x]$ will fully factor in K_2 , and one of the roots cannot have been hit by g_β because then their pre-images in L_β^1 includes a_β already. So send a_β to a root not hit yet.
- Suppose that a_β is transcendental over L_β^1 . Now, $|\overline{L_\beta^2}| = |L_\beta^2| + \aleph_0 < \kappa$, so there is a transcendental element of K_2 not in L_β^2 . Send a_β to such a transcendental element.

For b_β to go backwards, do the same argument in reverse. ■

Corollary 2.9. Fix algebraically closed fields K_1 and K_2 , and fix tuples $\bar{a} \in K_1^n$ and $\bar{b} \in K_2^n$. Then the following are equivalent.

- The structures (K_1, \bar{a}) and (K_2, \bar{b}) are equivalent in an expanded language.
- $k_1(\bar{a}) = k_2(\bar{b})$ where $k_1 \subseteq K_1$ and $k_2 \subseteq K_2$ are the prime subfields.
- For any quantifier-free formulae θ , we have $K_1 \models \theta(\bar{a})$ if and only if $K_2 \models \theta(\bar{b})$.

2.2 September 26

Today, we will give a structural way to look at quantifier elimination.

2.2.1 A Taste of Types

We split our discussion of quantifier elimination into the following lemmas.

Lemma 2.10. Fix \mathcal{L} -structures \mathcal{A} and \mathcal{B} . Further, fix $\bar{a} \in \mathcal{A}^n$ and $\bar{b} \in \mathcal{B}^n$ with $n \geq 1$. Then the following are equivalent.

- For any quantifier-free \mathcal{L} -formula φ , we have $\mathcal{A} \models \varphi(\bar{a})$ if and only if $\mathcal{B} \models \varphi(\bar{b})$.
- There is an isomorphism of substructures $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{B}' \subseteq \mathcal{B}$ containing \bar{a} and \bar{b} respectively, and the isomorphism sends \bar{a} to \bar{b} .

We will prove this in a moment, but we quickly note that it motivates the following definition.

Definition 2.11 (quantifier-free type). Fix an \mathcal{L} -structure \mathcal{A} and some $\bar{a} \in \mathcal{A}^n$. Then the *quantifier-free type of \bar{a}* , denoted $\text{qftp}^{\mathcal{A}}(\bar{a})$, is the set of quantifier free formulae φ such that $\mathcal{A} \models \varphi(\bar{a})$.

Anyway, here is our proof of Lemma 2.10.

Proof of Lemma 2.10. We have two implications to show.

- We show (b) implies (a). Suppose we have an isomorphism $f: \mathcal{A}' \rightarrow \mathcal{B}'$ as described. Now, suppose $\varphi(\bar{x})$ is a quantifier-free \mathcal{L} -formula with n free variables. Then $\mathcal{A}' \models \varphi(\bar{a})$ if and only if $\mathcal{B}' \models \varphi(\bar{b})$ by the nature of our isomorphism (see Proposition 1.14). Then this comes down to the substructure because φ is quantifier-free by Proposition 1.12.

¹ One needs to do something here in the case that κ is a singular.

- We show (a) implies (b). Define $A' \subseteq A$ to be the set of terms t evaluated on \bar{a} as $t^A(\bar{a})$, and define $B' \subseteq B$ similarly. We do need to check that A' is the universe of an \mathcal{L} -substructure of \mathcal{A} , and the check for B' will be similar. Well, we interpret constants (which are terms) exactly as they were interpreted in \mathcal{A} . We interpret functions exactly as they were interpreted in \mathcal{A} because terms are closed under applying functions. Lastly, relations are defined by intersection with A , which is what is needed to provide a substructure.

We now define $\mathcal{A}' \rightarrow \mathcal{B}'$ by sending the term $t^A(a_1, \dots, a_n)$ to $t^B(b_1, \dots, b_n)$. We have the following checks.

- Well-defined and injective: if s and t are terms with $s^A(\bar{a}) = t^A(\bar{a})$, then this is equivalent to $\mathcal{A} \models (s(\bar{x}) = t(\bar{x}))(\bar{a})$, which is equivalent to $\mathcal{B} \models (s(\bar{x}) = t(\bar{x}))(\bar{b})$ by hypothesis, which at the end is equivalent to $s^B(\bar{b}) = t^B(\bar{b})$.
- Surjective: any element of B' takes the form $t^B(\bar{b})$ for some term t , which is hit by $t^A(\bar{a})$.
- Isomorphism: this has many checks in itself. For any constant symbol c , we see $f(c^{A'}) = c^{B'}$ by viewing c as a term which does not care about the input \bar{a} . Now suppose F is an m -ary function symbol, then

$$\begin{aligned} f(F^{A'}(t_1^{A'}(\bar{a}), \dots, t_m^{A'}(\bar{a}))) &= f(\underbrace{F(t_1(\bar{x}), \dots, t_m(\bar{x}))(\bar{a})}_{\text{some term!}}) \\ &= F^B(t_1^B(\bar{b}), \dots, t_m^B(\bar{b})) \\ &= F^B(f(t_1^A(\bar{a}), \dots, t_m^A(\bar{a}))). \end{aligned}$$

Lastly, let R be an m -ary relation symbol. Then $(t_1^A(\bar{a}), \dots, t_m^A(\bar{a})) \in R^{A'}$ if and only if $\mathcal{A}' \models R(t_1, \dots, t_m)(\bar{a})$ if and only if $\mathcal{A} \models R(t_1, \dots, t_m)(\bar{a})$ by Proposition 1.12, which is now equivalent to $\mathcal{B} \models R(t_1, \dots, t_m)(\bar{b})$ and then equivalent to $\mathcal{B}' \models R(t_1, \dots, t_m)(\bar{b})$. ■

Remark 2.12. The \mathcal{A}' given in the proof above is the smallest substructure of \mathcal{A} containing \bar{a} .

More generally, we might be interested in types.

Definition 2.13 (type). Fix an \mathcal{L} -structure \mathcal{A} . Further, fix an n -tuple $\bar{a} \in A^n$. Then the *type*, denoted $\text{tp}^A(\bar{a})$ is the set of \mathcal{L} -formulae $\varphi(\bar{x})$ such that $\mathcal{A} \models \varphi(\bar{a})$.

Here is the corresponding result.

Lemma 2.14. Fix \mathcal{L} -structures \mathcal{A} and \mathcal{B} , and further fix $a \in A^n$ and $b \in B^n$. Suppose that there are elementary extensions $\mathcal{A}' \geq \mathcal{A}$ and $\mathcal{B}' \geq \mathcal{B}$ with an isomorphism $f: \mathcal{A}' \rightarrow \mathcal{B}'$ sending \bar{a} to \bar{b} . Then $\text{tp}^A(\bar{a}) = \text{tp}^B(\bar{b})$.

Proof. Note that we have elementary expansions $\mathcal{A}_{\bar{a}} \leq \mathcal{A}'_{\bar{a}}$ and $\mathcal{B}_{\bar{b}} \leq \mathcal{B}'_{\bar{b}}$. By hypothesis, the isomorphism $\mathcal{A}' \cong \mathcal{B}'$ sends \bar{a} to \bar{b} , so in fact $\mathcal{A}'_{\bar{a}}$ is isomorphic to $\mathcal{B}'_{\bar{b}}$. Tracking everything through, we see $\mathcal{A} \models \varphi(\bar{a})$ if and only if $\mathcal{A}_{\bar{a}} \models \varphi(\bar{a})$ if and only if $\mathcal{A}'_{\bar{a}} \models \varphi(\bar{a})$ if and only if $\mathcal{B}'_{\bar{b}} \models \varphi(\bar{b})$ if and only if $\mathcal{B}_{\bar{b}} \models \varphi(\bar{b})$ if and only if $\mathcal{B} \models \varphi(\bar{b})$. ■

Remark 2.15. The converse of this result is true, and we will prove it later in this class.

2.2.2 Back to Algebraically Closed Fields

Let's return to our discussion of algebraically closed fields.

Definition 2.16 (eliminates quantifiers). An \mathcal{L} -theory T *eliminates quantifiers* if and only if any formula $\varphi(\bar{x})$ has some quantifier-free formula $\bar{\psi}(\bar{x})$ such that $T \models \forall x(\varphi(\bar{x}) \leftrightarrow \bar{\psi}(\bar{x}))$.

Theorem 2.17. Say that an \mathcal{L} -theory T is “isomorphism-extendable” if and only if it has the following property: for any models $\mathcal{A}, \mathcal{B} \models T$ with fixed n -tuples $a \in A^n$ and $b \in B^n$ equipped with an isomorphism $f: \mathcal{A}' \rightarrow \mathcal{B}'$ of substructures containing \bar{a} and \bar{b} (respectively) which sends \bar{a} to \bar{b} , then any elementary superstructures $\mathcal{A}^* \geq \mathcal{A}$ and $\mathcal{B}^* \geq \mathcal{B}$ have an isomorphism extending f . Then if T is isomorphism-extendable, then T eliminates quantifiers.

Proof. Fix a formula $\varphi(\bar{x})$. Observe that being isomorphism-extendable implies that \bar{a} and \bar{b} having the same quantifier-free type implies that they have the same type by combining Lemmas 2.10 and 2.14.

For technical reasons, we extend the language to \mathcal{L}^* to have some new constants c_1 and c_2 for each of the old constants c . Our functions are the same, and we add in one more unary relation U . The point of introducing \mathcal{L}^* is to be able to talk about two \mathcal{L} -structures of the same type.

Explicitly, given an \mathcal{L}^* -structure where U contains the c_1 s and the complement contains the c_2 s (and these are nonempty), then we can restrict to U and its complement to provide two \mathcal{L} -structures. Conversely, given \mathcal{L} -structures \mathcal{A} and \mathcal{B} , we build an \mathcal{L}^* -structure with universe $A \sqcup B$ as follows: interpret the constants c_1 and c_2 as in \mathcal{A} and \mathcal{B} , respectively. Interpret the values $f(\bar{a})$ and $f(\bar{b})$ for $\bar{a} \in A^\bullet$ and $\bar{b} \in B^\bullet$ as in \mathcal{A} and \mathcal{B} , respectively, and interpret $f(\bar{e})$ for any other \bar{e} however we wish. One does something similar for the relations. Notably, the \mathcal{L}^* -structure, which we call \mathcal{A} , is not exactly the same data as two \mathcal{L} -structures because one has to say what happens on the function and relation symbols when we have not been told by \mathcal{A} and \mathcal{B} alone.

Anyway, let $\varphi(\bar{x})$ be an \mathcal{L} -formula, and we expand \mathcal{L}^* to add in some new constant symbols \bar{a} and \bar{b} . We now relative to build a new theory. The observation is that, using the construction of the previous paragraph, there is a function $\hat{\varphi}^{\mathcal{A}}$ such that $\mathcal{A} \models \varphi(\bar{a})$ if and only if $\mathcal{C} \models \hat{\varphi}^{\mathcal{A}}(\bar{a})$. As such, we adjust T to the theory Σ be an \mathcal{L}^* -theory by adjusting c s to c_1 s and c_2 s in the natural way, and we also add in the sentences $U(a_\bullet)$ and $\neg U(b_\bullet)$. Further, we add in the sentences

$$\{\varphi^{\mathcal{A}}(\bar{a}) \leftrightarrow \psi^{\mathcal{B}}(\bar{b})\}$$

as well as $\hat{\varphi}^{\mathcal{A}}(\bar{a}) \leftrightarrow \varphi^{\mathcal{B}}(\bar{b})$. This theory is inconsistent by the type discussion at the very beginning of this proof: we are being promised that \bar{a} and \bar{b} have the same type, but then they disagree on φ !

Thus, by compactness, there is a finite set Ψ of quantifier-free formulae with the following property for any models $\mathcal{A}, \mathcal{B} \models T$ with $\bar{a} \in A^n$ and $\bar{b} \in B^n$: if $\mathcal{A} \models \psi(\bar{a})$ is equivalent to $\mathcal{B} \models \psi(\bar{b})$ for each $\psi \in \Psi$, then we must have $\mathcal{A} \models \varphi(\bar{a})$ is equivalent to $\mathcal{B} \models \varphi(\bar{b})$. We now construct our quantifier-free formula: for each $X \subseteq \Psi$, we define

$$\theta_X := \bigwedge_{\psi \in X} \psi \wedge \bigwedge_{\psi \in \Psi \setminus X} \neg \psi,$$

and we let G be the set of subsets such that there is a model $\mathcal{A} \models T$ with $\mathcal{A} \models \exists \bar{x}(\varphi(\bar{x}) \wedge \theta_X(\bar{x}))$. Then we set $\eta(\bar{x})$ to be the disjunction over all the θ_X where $X \in G$. Note that $\eta(\bar{x})$ is quantifier-free.

We now claim that $T \models \forall \bar{x}(\eta(\bar{x}) \leftrightarrow \varphi(\bar{x}))$. Suppose $\mathcal{A} \models T$ and we have some $\bar{a} \in \mathcal{A}$ with $\mathcal{A} \models \varphi(\bar{a})$. Then we consider the subset X of Ψ such that $\mathcal{A} \models \psi(\bar{a})$ if and only if $\psi \in X$. Then \mathcal{A} is in fact modelling $\varphi(\bar{a})$ along with the sentences $\psi(\bar{a})$ for each $\psi \in X$ and then $\neg \psi(\bar{a})$ for each $\psi \notin X$. Thus, $\mathcal{A} \models \theta_X(\bar{a}) \wedge \varphi(\bar{a})$, so $X \in G$, and $T \models \forall \bar{x}(\varphi(\bar{x}) \rightarrow \eta(\bar{x}))$ follows.

We now go in the other direction. Suppose $\mathcal{A} \models T$ is a model, and suppose we have $\bar{a} \in A^n$ and $\mathcal{A} \models \eta(\bar{a})$. Then there is some $X \in G$ such that $\mathcal{A} \models \theta_X(\bar{a})$, but being in G promises us a model $\mathcal{B} \models T$ and $\bar{b} \in B^n$ with $\mathcal{B} \models \varphi(\bar{b}) \wedge \theta_X(\bar{b})$. But then any $\psi \in \Psi$ has $\mathcal{A} \models \psi(\bar{a})$ if and only if $\mathcal{B} \models \psi(\bar{b})$ by definition of θ_X , so \mathcal{A} and \mathcal{B} must agree on $\varphi(\bar{b})$. In other words, we conclude $\mathcal{A} \models \varphi(\bar{a})$, and we are done. ■

Corollary 2.18. The theory ACF eliminates quantifiers.

Proof. We show the hypothesis of the above theorem. Given two algebraically closed fields K and L with an isomorphism $f: K' \rightarrow L'$ where $K' \subseteq K$ and $L' \subseteq L$ are algebraically closed subfields, we need an isomorphism $f^*: K^* \rightarrow L^*$ extending f . As long as K and L have the same cardinality, we can simply do this with $K = K^*$ and $L = L^*$. In general, with $|K| \leq |L|$, we might need to use a transcendence basis to expand K and take an algebraic closure, and this is an elementary extension because ACF is κ -categorical. ■

Corollary 2.19. The theory of dense linear order without endpoints eliminates quantifiers.

Proof. Use the theorem. ■

Non-Example 2.20. The theory of an equivalence relation with exactly one equivalence class of size each positive integer does not eliminate quantifiers. To see this, consider the sentence which says that a free variable x is in an equivalence class of size 2.

2.3 September 28

Let's talk about some game.

2.3.1 Ehrenfeucht–Fraïssé Games

For today's lecture, let's discuss Ehrenfeucht–Fraïssé Games. Recall the following definition.

Definition 2.21 (unnested). An atomic \mathcal{L} -formula φ is *unnested* if and only if it takes one of the following forms.

- Equalities: $t_i = t_j$ or $x_i = c$ where the t_\bullet are variables or constants.
- Relations: $R(t_1, \dots, t_n)$ where the t_\bullet are variables or constants.
- Functions: $f(t_1, \dots, t_n) = t_{n+1}$ where the t_\bullet are variables or constants.

For our discussion today, we let U_0 denote the set of finite boolean combinations of unnested atomic formulae, up to provable equivalence (e.g., we don't want to include $\varphi \wedge \varphi$ from φ), and we inductively set U_{n+1} to be finite boolean combinations (again, up to provable equivalence) of formulae of the form $\exists x\psi$ where $\psi \in U_n$ and x is a variable.

Proposition 2.22. Fix a finite language \mathcal{L} . Then for each n and m , there are only finitely many formulae in U_n with the variables x_1, \dots, x_m (up to provable equivalence).

Proof. Fix m , and we induct on n . We start with $n = 0$. For number unnested atomic formulae is finite because the problem is just combinatorics to count sentences of each type. As for the boolean combinations, we note that the boolean algebra generated by a finite set is finite, so there are only finitely many classes up to provable equivalence. Then to go up, we place $\exists x_\bullet$ or not in front of each formula, so there continue to be only finitely many formulae, and the boolean algebra generated continues to be finite, so we are okay. ■

Our observation, now, is that every \mathcal{L} -formula is equivalent to some formula in one of the U_n .

Proposition 2.23. Fix a language \mathcal{L} . Then any \mathcal{L} -formula φ is equivalent to some $\psi \in U_n$ for some n .

Proof. It suffices to check this for atomic formulae; all other formulae follow by adding enough quantifiers and taking boolean combinations. Here are our cases.

- Take sentences of the form $t_1 = t_2$. We now have to induct on the complexity of the terms. If we have an equality of variables $x_i = x_j$ or an equality $x_i = c$ for constant c , there is nothing to say. If we have $c = x_i$, then this is equivalent to the unnested formula $x_i = c$. Lastly, $c = d$ is equivalent to the sentence $\exists x(x = c \wedge x = d)$.

Now if we have something of the type $t_1 = f(s_1, \dots, s_n)$, then by induction, we can achieve any of the formulae $x_{n+1} = t_1$ and $x_i = s_i$ for each i where the x_i are variables. So $t_1 = f(s_1, \dots, s_n)$ is equivalent to

$$\exists x_1 \cdots \exists x_n \left(\bigwedge_{i=1}^n x_i = s_i \wedge x_{n+1} = t_1 \wedge x_{n+1} = f(x_1, \dots, x_n) \right).$$

This induction completes this case.

- For relations, one does essentially the same trick. If we have $R(t_1, \dots, t_n)$, we simply look at the sentences $x_i = t_i$ combined with $R(x_1, \dots, x_n)$, reducing to the previous case. ■

Now let's play a game. Fix a language \mathcal{L} with two \mathcal{L} -structures \mathcal{A} and \mathcal{B} , and we fix a natural number n . The game $EF_n(\mathcal{A}, \mathcal{B})$ of length n is played as follows.

- Player I picks \mathcal{A} or \mathcal{B} and chooses some $a_1 \in A$ or $b_1 \in B$. Then Player II chooses an element $b_1 \in B$ or $a_1 \in A$ from the opposite universe to the one Player I chose.
- Then the above move is repeated until we have two n -tuples (a_1, \dots, a_n) or (b_1, \dots, b_n) .
- Player II wins if, for any unnested atomic formula $\psi(x_1, \dots, x_n)$, we have $\mathcal{A} \models \psi(\bar{a})$ is equivalent to $\mathcal{B} \models \psi(\bar{b})$. Otherwise, Player I wins.

Roughly speaking, Player I wants to make \mathcal{A} and \mathcal{B} look different, and Player II wants them to look similar. We write $\mathcal{A} \equiv^n \mathcal{B}$ to mean that Player II can win the EF_n game.

Example 2.24. Fix the language $\mathcal{L} = \{<\}$, and take \mathcal{A} to be $\omega + \omega^*$, where the ω^* means we have concatenated ω on top of ω^* but in reverse (so that 0^* is the largest element). We then let \mathcal{B} be the set $\{0, 1, 2, \dots, 6\}$ for some natural m , and we play the game. Player I can win the game in four moves, but Player II can win in three moves.

Here is our result.

Proposition 2.25. Fix a finite language \mathcal{L} . For each n and structures \mathcal{A} and \mathcal{B} , Player II has a winning strategy in the $EF_n(\mathcal{A}, \mathcal{B})$ game if and only if $\mathcal{A} \models \psi$ is equivalent to $\mathcal{B} \models \psi$ for all sentences $\psi \in U_n$.

Proof. We prove this by induction on n , but the inductive hypothesis will allow \mathcal{A} and \mathcal{B} to vary. At $n = 0$, we are asking for $\mathcal{A} \models \varphi$ if and only if $\mathcal{B} \models \varphi$ where φ is an unnested atomic formula, so Player II wins if and only if this is satisfied.

For our induction, suppose n , and we get $n + 1$. There are two implications to show.

- In one direction, suppose Player II has a winning strategy. Suppose Player I has picked $a_1 \in A$ (without loss of generality). Then Player II responds with some $b_1 \in B$ according to the winning strategy. Now, the rest of the game is a length- n game in the language \mathcal{L}' expanded by a constant symbol c with the structures \mathcal{A}' and \mathcal{B}' have $c^{\mathcal{A}'} = a_1$ and $c^{\mathcal{B}'} = b_1$. So we are now playing $EF_n(\mathcal{A}', \mathcal{B}')$. So Player II has a winning strategy in $EF_{n+1}(\mathcal{A}, \mathcal{B})$ if and only if, for all $a \in A_1$, there exists $b_1 \in B$ such that Player II has a winning strategy in $EF_n(\mathcal{A}', \mathcal{B}')$. Anyway, by the induction, we get $\mathcal{A}' \equiv^n \mathcal{B}'$ in \mathcal{L}' .

We now show that $\mathcal{A} \equiv^{n+1} \mathcal{B}$. Thus far we are given that $\mathcal{A}' \models \psi$ if and only if $\mathcal{B}' \models \psi$ for any \mathcal{L}' -sentence $\psi \in U_n$. We now do our check. Fix a sentence $\theta \in U_{n+1}$ of the form $\exists x_1 \varphi$ where $\varphi \in U_n$. Then $\mathcal{A} \models \theta$ is equivalent to having some $a_1 \in A$ such that $\mathcal{A} \models \varphi(a_1)$. Let b_1 be the resulting choice of Player II. But now using our hypothesis at the beginning of the paragraph, we achieve $\mathcal{A}' \models \varphi(c)$, so $\mathcal{B}' \models \varphi(c)$, so $\mathcal{B} \models \varphi(b_1)$. The reverse implication is similar.

- Conversely, suppose that $\mathcal{A} \models \psi$ is equivalent to $\mathcal{B} \models \psi$ for all sentences $\psi \in U_n$. We give a winning strategy for Player II. Let's say $a_1 \in A$ is chosen by Player I. Let Ψ be the set of formulae $\psi(x_1) \in U_n$ with at most $(n+1)$ variables such that $\mathcal{A} \models \psi(a_1)$, which is a finite set up to provable equivalence by Proposition 2.22. It is important that Ψ is finite because now

$$\mathcal{A} \models \exists x_1 \bigwedge_{\psi \in \Psi} \psi(x_1).$$

This formula lives in U_{n+1} , so by hypothesis, we get

$$\mathcal{B} \models \exists x_1 \bigwedge_{\psi \in \Psi} \psi(x_1),$$

so we get $b_1 \in B$ satisfying all $\mathcal{B} \models \psi(b_1)$ for $\psi \in \Psi$.

Now build \mathcal{L}' and structures \mathcal{A}' and \mathcal{B}' as before. We claim that $\mathcal{A}' \models \varphi$ if and only if $\mathcal{B}' \models \varphi$ for all \mathcal{L}' -sentences $\varphi \in U_n$. Indeed, simply view φ as an \mathcal{L} -formula $\tilde{\varphi}(x)$ by extracting out the constant c and replacing it with x , and we see $\mathcal{A}' \models \varphi$ is equivalent to $\mathcal{A} \models \tilde{\varphi}(a_1)$, which is indeed equivalent to $\mathcal{B} \models \tilde{\varphi}(b_1)$.

Now by induction, Player II has a winning strategy in the game $EF_n(\mathcal{A}', \mathcal{B}')$, which is equivalent to winning the original game, as discussed in the previous implication. ■

Corollary 2.26. Fix a language \mathcal{L} . Then $\mathcal{A} \equiv \mathcal{B}$ if and only if, for all finite language $\mathcal{L}' \subseteq \mathcal{L}$, we have $\mathcal{A}|_{\mathcal{L}'} \equiv \mathcal{B}|_{\mathcal{L}'}$.

Proof. Play the above game. Note $\mathcal{A} \equiv \mathcal{B}$ if and only if they satisfy the same formulae, which is equivalent to having $\mathcal{A}|_{\mathcal{L}'} \equiv \mathcal{B}|_{\mathcal{L}'}$ for all finite $\mathcal{L}' \subseteq \mathcal{L}$ because any formula will only contain finitely many symbols. Then this is in fact equivalent to satisfying the same \mathcal{L}' -sentences in U_n for all n , which finishes by Proposition 2.25. ■

Remark 2.27. Here is a challenge problem: for which m and n does Player II win the game of length n between the groups \mathbb{Z} and $\mathbb{Z}/m\mathbb{Z}$? There does exist some n such that Player I will always win this game. Approximately speaking, one needs a sentence true in \mathbb{Z} which is false in the $\mathbb{Z}/m\mathbb{Z}$ s.

2.4 October 3

Let's play the game to start off the class.

Example 2.28. We work with ordinals in the language $\mathcal{L} = \{<\}$.

- We play with $\varepsilon_0 = \sup \{\omega, \omega^\omega, \dots\}$ and 2. Then Player II loses after, say, 2 moves: Player I selects anything, Player II selects (say) 0, and then Player I chooses something smaller than what they chose in ε_0 .
- We play with ε_0 and ω_1 . Then Player II can always win. The point is that there is some kind of finite-length back-and-forth argument

2.4.1 Real Closed Fields

Let's discuss real closed fields because they, in some sense, will tell us that Euclidean geometry is decidable (approximately speaking). Our language will be the language $\mathcal{L} = \{+, -, \cdot, <, 0, 1\}$ of ordered rings. The theory of ordered fields OrdFld is axiomatized by writing the axioms for fields, for a total order, and requiring that addition and multiplication respect this ordering. We won't bother writing down the first two lists of axioms, but the third list is given as follows.

- $\forall a \forall b \forall c ((a < b) \rightarrow (a + c < b + c)).$
- $\forall a \forall b \forall c (((a < b) \wedge (c > 0)) \rightarrow (a \cdot c < b \cdot c)).$

So we have a finitely axiomatized our theory OrdFld .

Example 2.29. Any subfield of \mathbb{R} will do is a model.

Example 2.30. We can use compactness to provide a model of \mathbb{R} with an element larger than any other element but the same cardinality.

We will be actually be interested in the theory RCF of real closed fields, which is the theory OrdFld plus the intermediate value theorem for polynomials. This is an infinite list of axioms, approximately saying that, for any model \mathcal{R} with universe R , and polynomial $f \in R[x]$ with inputs $a, b \in R$ such that $f(a) < 0 < f(b)$ has some $c \in R$ such that $f(c) = 0$.

To write this out, we choose a degree of n and write down the sentence

$$\forall a_0 \cdots \forall a_n \forall a \forall b \left(((a < b) \wedge (a_0 a^0 + \cdots + a_n a^n < 0 < a_0 b^0 + \cdots + a_n b^n)) \right. \\ \left. \rightarrow \exists c (a < c < b \wedge a_0 c^0 + \cdots + a_n c^n = 0) \right).$$

We cannot finitely axiomatize these sentences using an argument like Lemma 2.1.

Remark 2.31. Any ordered field $(\mathcal{R}, +, -, 0, 1, <)$ has $(\mathcal{R}, <)$ satisfying DLO. We know that we are a linear order, we have no endpoints because $x + -1 < x < x + 1$ for any $x \in R$, and we are dense because $x < \frac{x+y}{2} < y$ for any $x, y \in R$. Note that checking $x < x + 1$ (for example) requires knowing that $0 < 1$, which is a nontrivial fact on its own (one should use trichotomy and rule out $0 = 1$ by fields and rule out $0 > 1$ because this would imply $-1 > 0$ and then $1 > 0$ by squaring). There are lots of these nontrivial facts (e.g., we also want to know $0 < 1/2 < 1$), but we won't bother to show this.

For ordered fields, there is an order topology, and one can show that various functions like $+$ and \cdot and polynomials are all continuous.

We will define the function $|\cdot| : R \rightarrow R$ given by

$$|x| := \begin{cases} +x & \text{if } x \geq 0, \\ -x & \text{if } x \leq 0. \end{cases}$$

Now, if $x < 0$, then $-x > 0$ by subtraction, so we see that $|x| > 0$ for all $x \neq 0$. The standard casework is also able to prove the triangle inequality $|x + y| \leq |x| + |y|$ by some casework. If both nonpositive or nonnegative, then we have equality, and if they have different signs (say, $x > 0 > y$ without loss of generality and $|x| \geq |y|$), then we are looking at $x + y \leq x - y$, which is true.

For notation, we will also want the function sgn given by

$$\text{sgn}(x) := \begin{cases} +1 & \text{if } x > 0, \\ 0 & \text{if } x = 0, \\ -1 & \text{if } x < 0. \end{cases}$$

Now, one is able to check the following, which tells us that polynomials "go off to infinity."

Proposition 2.32. Fix an ordered field \mathcal{R} and a polynomial $f(x) \in \mathcal{R}[x]$ of positive degree d , written

$$f(x) = \sum_{i=0}^d c_i x^i$$

where $c_d \neq 0$. If

$$x > 1 + \frac{1}{|c_d|} \sum_{i=0}^{d-1} |c_i|,$$

then $\operatorname{sgn} f(x) = \operatorname{sgn} c_d$.

Proof. Boring bounding. Note

$$\operatorname{sgn} f(x) = \operatorname{sgn} c_d \cdot \operatorname{sgn} \left(x^d + \sum_{i=0}^{d-1} \frac{c_i}{c_d} x^i \right),$$

so by scaling down, it is enough to consider the case where $c_d = 1$.

As an aside, we note that any $x \geq 1$ and nonnegative integer n will have $x^n \geq x$, which is true by induction because $x^{n+1} \geq x^n$, where our base case is $x^1 = x \geq 1 = x^0$. With this in mind, we see that x satisfying the desired inequality will have

$$x^d = x \cdot x^{d-1} > \sum_{i=0}^{d-1} |c_i| x^{d-1} \geq \sum_{i=0}^{d-1} |c_i| x^i \geq \sum_{i=0}^{d-1} -c_i x^i,$$

so $f(x) > 0$ follows. ■

Corollary 2.33. If \mathcal{R} is a real closed field and $a \geq 0$, then there exists $b \geq 0$ such that $b^2 = a$.

Proof. If $a = 0$, set $b = 0$. Otherwise, consider the polynomial $f(x) := x^2 - a$. Note $f(0) < 0$, and Proposition 2.32 tells us that $f(1+a) > 0$, so the intermediate value theorem for polynomials tells us that there is some b such that $f(b) = 0$, so $b^2 = a$. ■

Corollary 2.34. If \mathcal{R} is a real closed field, then any polynomial $f(x)$ of odd degree has a root.

Proof. Write

$$f(x) = \sum_{i=0}^d c_i x^i$$

where d is odd and $c_d \neq 0$, and let $N := 2 + \frac{1}{|c_d|} \sum_{i=0}^{d-1} |c_i|$. By Proposition 2.32, we have N such that $\operatorname{sgn} f(N) = \operatorname{sgn} c_d$, and we see similarly that the polynomial $f(-x)$ will now have $\operatorname{sgn} f(-N) = \operatorname{sgn}(-c_d)$. Thus, $f(N)$ and $f(-N)$ have different signs, so the intermediate value theorem for polynomials grants f a root. ■

The above two corollaries turn out to characterize real closed fields.

Remark 2.35. We can now remove the ordering from our real closed fields by declaring that squares are exactly the nonnegative elements. It is in general an interesting question when we can give a field an order; for example, -1 cannot be a sum of squares because $-1 < 0$. This turns out to be good enough to make a field orderable!

2.5 October 5

Here we go.

2.5.1 Quantifier Elimination via Back-and-Forth

Our goal is to show that RCF eliminates quantifiers and is thus complete. Here will be our test.

Proposition 2.36. Fix an \mathcal{L} -theory T . Then T is complete and has quantifier elimination if the following two properties hold.

- (i) There is a “prime” structure: there is an \mathcal{L} -structure \mathcal{A} such that any model $\mathcal{M} \models T$ has an embedding $\mathcal{A} \subseteq \mathcal{M}$.
- (ii) Extension: for any two models \mathcal{M} and \mathcal{N} with an isomorphism $\varphi: \mathcal{M}_0 \rightarrow \mathcal{N}_0$ between substructures $\mathcal{M}_0 \subseteq \mathcal{M}$ and $\mathcal{N}_0 \subseteq \mathcal{N}$, then any chosen element $a \in \mathcal{M}$ has an extension $g: \mathcal{M}' \rightarrow \mathcal{N}'$ extending φ where $a \in \mathcal{M}' \subseteq \mathcal{M}$ and \mathcal{N}' is a substructure of an elementary extension \mathcal{N}^* of \mathcal{N} .

Proof. We will show that (ii) implies that there are elementary extensions $\mathcal{M} \leq \widetilde{\mathcal{M}}$ and $\mathcal{N} \leq \widetilde{\mathcal{N}}$ with an isomorphism $\widetilde{f}: \widetilde{\mathcal{M}} \rightarrow \widetilde{\mathcal{N}}$ extending f . This is a back-and-forth argument, using (ii) to extend our isomorphism one element at a time.

We build a chain of models $\mathcal{M} := \mathcal{M}^0 \leq \mathcal{M}^1 \leq \mathcal{M}^2 \leq \dots$ and $\mathcal{N} := \mathcal{N}^0 \leq \mathcal{N}^1 \leq \mathcal{N}^2 \leq \dots$ and $\widetilde{\mathcal{M}}$ and $\widetilde{\mathcal{N}}$ will be the union of the chains. Roughly speaking, the idea is to construct our models with f_1, f_2, \dots into the following diagram.

$$\begin{array}{ccc}
 \vdots & & \vdots \\
 \uparrow & & \uparrow \\
 \mathcal{M}^2 & \xleftarrow{\widetilde{f}_4} & \mathcal{N}^2 \\
 \uparrow & \nearrow \widetilde{f}_3 & \uparrow \\
 \mathcal{M}^1 & \xleftarrow{\widetilde{f}_2} & \mathcal{N}^1 \\
 \uparrow & \nearrow \widetilde{f}_1 & \uparrow \\
 \mathcal{M}^0 & & \mathcal{N}^0
 \end{array} \tag{2.1}$$

Let's begin by exhibiting \widetilde{f}_1 . Enumerate $\mathcal{M} = \{m_\alpha : \alpha \in \kappa\}$ where $\kappa = |\mathcal{M}|$. Now, we write down our maps.

- (a) Set $g_0 = f$.
- (b) We will have a map $g_\alpha: \mathcal{A}_\alpha \rightarrow \mathcal{N}_\alpha^0$, where $m_\beta \in \mathcal{A}_\alpha$ for any $\beta < \alpha$.
- (c) If $\alpha \leq \beta$ are in κ , then we require $\mathcal{A}_\alpha \subseteq \mathcal{A}_\beta \subseteq \mathcal{M}$ and $\mathcal{N}^0 \leq \mathcal{N}_\alpha^0 \leq \mathcal{N}_\beta^0$.

Then taking the union of the g_α will produce the needed map $\mathcal{M}^0 \rightarrow \mathcal{N}^1$, and reversing the picture produces $\mathcal{N}^1 \rightarrow \mathcal{M}^1$, and we can keep going up the chains.

Anyway, let's construct our g_α . We have already defined g_0 .

- Suppose we have defined $g_\alpha: \mathcal{A}_\alpha \rightarrow \mathcal{N}_\alpha^0$, and we want to get to a successor ordinal $g_{\alpha+1}$. Then (ii) using the single element $m_{\alpha+1} \in \mathcal{M}^0$ on the morphism g_α provides us with an extension $g_{\alpha+1}: \mathcal{A}_{\alpha+1} \rightarrow \mathcal{N}_{\alpha+1}^0$ where $a \in \mathcal{A}_{\alpha+1} \subseteq \mathcal{M}$ and $\mathcal{N}_\alpha^0 \leq \mathcal{N}_{\alpha+1}^0$. So we are done.
- On limit ordinals, we just take a union. If α is a limit ordinal, then we get to suppose that we have defined g_β for all $\beta < \alpha$, and we define

$$\mathcal{A}_\alpha := \bigcup_{\beta < \alpha} \mathcal{A}_\beta \quad \text{and} \quad \mathcal{N}_\alpha^0 := \bigcup_{\beta < \alpha} \mathcal{N}_\beta^0,$$

and we satisfy all the needed hypotheses by how chains work.

Alright, so we have constructed our map $\tilde{f}_1: \mathcal{M}^0 \rightarrow \mathcal{N}^1$ by taking unions of the above g_\bullet s. We can repeat this process to produce the maps \tilde{f}_\bullet and then go up the chain (2.1) to complete the argument. Namely, going up the chain tells us that we get embeddings in both directions whose compositions are the identity, so we do have an isomorphism. Thus, Theorem 2.17 tells us that T eliminates quantifiers.

It remains to check that T is complete, which is where (i) will appear. Fix models \mathcal{M} and \mathcal{N} of T . Now, our prime structure \mathcal{A} embeds into both \mathcal{M} and \mathcal{N} , whose images we will call \mathcal{A}_1 and \mathcal{A}_2 . Now, the above result tells us that we can extend this isomorphism of substructures to an isomorphism of elementary superstructures $\tilde{\mathcal{M}} \rightarrow \tilde{\mathcal{N}}$. Thus, $\mathcal{M} \equiv \tilde{\mathcal{M}} \equiv \tilde{\mathcal{N}} \equiv \mathcal{N}$, so $\text{Th } \mathcal{M} = \text{Th } \mathcal{N}$, which produces completeness. ■

Remark 2.37. Professor Scanlon is lightly considering putting the following weak form of Keisler–Shelah on the exam: if $\mathcal{A} \equiv \mathcal{B}$, then there is a direct limit of ultrapowers of \mathcal{A} and \mathcal{B} which are

Remark 2.38. More generally, the above proof shows that we can complete a theory T which eliminates quantifiers by adding in the diagram of any particular substructure of a model T .

2.5.2 Back to Real Closed Fields

Let's use Proposition 2.36.

Theorem 2.39. The theory RCF eliminates quantifiers and is complete.

Proof. Let's start with the prime structure.

Lemma 2.40. The theory RCF has a prime structure.

Proof. The integers \mathbb{Z} as an ordered integral domain is contained in any ordered field, so it works as our prime substructure. ■

Now for the hard part. Fix real closed fields \mathcal{R}_1 and \mathcal{R}_2 with an isomorphism of substructures $f: \mathcal{A}_1 \rightarrow \mathcal{A}_2$, and choose some $a \in \mathcal{R}_1$. We would like to extend f up to a . Note that there is some content in deciding how to extend \mathcal{A}_1 to a domain of f .

For example, note that \mathcal{A}_1 and \mathcal{A}_2 as substructures of a field must be an integral domain, and so of course we note that f can be extended to $\text{Frac } \mathcal{A}_1 \rightarrow \text{Frac } \mathcal{A}_2$ as a field homomorphism. Additionally, note that this extension also to the fraction field also respects the order: it suffices to note that f will respect positivity, so we note $\text{sgn } f(a) = \text{sgn } a$ for any $a \in \mathcal{A}_1$, so $a/b \in \text{Frac } \mathcal{A}_1$ being positive implies $\text{sgn } a = \text{sgn } b$ and so $\text{sgn } f(a) = \text{sgn } f(b)$ and so $f(a)/f(b)$ is positive. In total, we may assume that \mathcal{A}_1 and \mathcal{A}_2 are ordered fields.

Next up, we may assume that the degree of the field extension $[\mathcal{A}_1(a) : \mathcal{A}_1]$ is minimal among the degrees $[\mathcal{A}_1(a') : \mathcal{A}_1]$ for $a' \in \mathcal{R}_1 \setminus \mathcal{A}_1$ and $[\mathcal{A}_2(b') : \mathcal{A}_2]$ for $b' \in \mathcal{R}_2 \setminus \mathcal{A}_2$. The point is that we can deal with the elements a' and b' one at a time, starting with the smallest possible degree, and this is okay because we can take a countable union, and the total number of elements to deal with are countable over \mathcal{A}_1 , and the number of degrees is also countable.

Now, if a is algebraic over \mathcal{A}_1 , then let p be its minimal monic polynomial over \mathcal{A}_1 ; if α is transcendental, take $p = 0$. Now, define

$$\text{Cut}^-(a/\mathcal{A}_1) = \{\alpha \in \mathcal{A}_1 : \alpha < a\} \quad \text{and} \quad \text{Cut}^+(a/\mathcal{A}_1) = \{\beta \in \mathcal{A}_1 : \alpha < \beta\}.$$

If a is algebraic, then both of these sets are nonempty: Proposition 2.32 grants us a number $N_p \in \mathcal{A}_1$ such that $|x| > N_p$ will have $p(x) \neq 0$ in any ordered field, so $|x| > \alpha$.² Now, we note that we have the chain of

² If a is transcendental, one can in fact have a bigger than anything in \mathcal{A}_1 . For example, compactness provides a model of RCF which is just \mathbb{R} , and then we add in an element bigger than anything in \mathbb{R} .

isomorphisms

$$A_1[a] \cong \frac{A_1[x]}{p(x)} \cong \frac{A_2[x]}{f(p)(x)}.$$

To continue, we need to place a inside A_2 .

Proposition 2.41. Fix an ordered field R . Given a polynomial $F(x) \in R[x]$ and $d \in R$, if $F'(d) > 0$, then there exists $b < d < c$ such that $b < x < d < y < c$ implies $F(x) < F(d) < F(y)$.

Proof. We are basically trying to show that F is locally increasing. Now, we acknowledge that any polynomial $F(x) \in S[x]$ will have

$$F(X + Y) = \sum_{i=0}^{\deg f} \frac{1}{i!} F^{(i)}(X) Y^i.$$

Then

$$F(y) - F(d) = F'(d)(y - d) + \sum_{i=2}^{\deg f} \frac{1}{i!} F^{(i)}(d)(y - d)^i = F'(d)(y - d) \left(1 + \sum_{i=2}^{\deg f} \frac{F^{(i)}(d)}{i! F'(d)} (y - d)^{i-1} \right).$$

Repeating the computation Proposition 2.32, one sees that $|y - d|$ being sufficiently small makes the sign of the bit in parentheses positive, so $\text{sgn}(F(y) - F(d)) = \text{sgn}(y - d)$, and we complete the argument. ■

We will complete the proof next class. ■

2.6 October 10

The exam is in a little over a week. Exercises will be focused on content covered in class (and harder exercises will be chosen from there), but it is possible to be asked about other topics in Marker. Some exercises on the midterm will be taken from exercises assigned to us.

The class began by showing that

2.6.1 Back to Back to Real Closed Fields

For the time being, take a to be algebraic. We claim that there is $\alpha \in \text{Cut}^-(a/A_1)$ and $\beta \in \text{Cut}^+(a/A_1)$ with $\text{sgn}(P(\alpha)) \neq \text{sgn}(P(\beta))$, which is a sign change that we will be able to push over to \mathcal{R}_2 in order to produce a root over there. Well, note $P'(a) \neq 0$ because we are in characteristic 0, so everything is separable. We take $P'(a) > 0$; otherwise simply reverse all signs. Then Proposition 2.41 grants $b < a < c$ with $P(x) < 0 < P(y)$ whenever $b < x < a < c < y$, but technically the argument only gives $b, c \in R_1$, and the same holds for everything between.

It remains to bring these down to A_1 . For this, we use the following lemma.

Lemma 2.42. Fix a real closed field \mathcal{R} . For $Q[x] \in R[x]$ and $\alpha < \beta$ with $Q(\alpha) = Q(\beta) = 0$, there is $\gamma \in [\alpha, \beta]$ such that $Q'(\gamma) = 0$.

Proof. If $Q'(\alpha) = 0$ or $Q'(\beta) = 0$, there is nothing to do. Now, if $Q'(\alpha)$ and $Q'(\beta)$ have different signs, \mathcal{R} being a real closed field grants us our γ .

Lastly, suppose $Q'(\alpha)$ and $Q'(\beta)$ have the same sign. Without loss of generality, make both of them positive. Then there is $\varepsilon > 0$ such that $\alpha < \alpha + \varepsilon < \beta - \varepsilon < \beta$ such that having $\alpha < \gamma < \alpha + \varepsilon$ implies $Q(\gamma) > 0$ and having $\beta - \varepsilon < \gamma < \beta$ implies $Q(\gamma) < 0$. So Q has another root strictly between α and β , so we replace β with this root β' .

Namely, check if $Q'(\beta') \leq 0$, we get our root of Q' ; otherwise, we repeat the process for $[\alpha, \beta']$ to get yet another root β'' . This process must eventually terminate because Q can only have finitely many roots, so we get our needed root of Q . ■

Now, choose $\alpha \in \text{Cut}^-(a/A_1)$ and $\beta \in \text{Cut}^+(a/A_1)$. For concreteness, list the roots of $P(x)$ as $a_1 < a_2 < \dots < a_n$, and suppose $a = a_i$ for some i . Now, set $\lambda := \alpha$ if $a = a_1$ or instead a root of $P'(x)$ between a_{i-1} and $a_i = a$ if $i > 1$. Similarly, set $\mu := \beta$ if $a = a_n$ or instead a root of $P'(x)$ between a_i and a_{i+1} if $i < n$. (These exist by the above lemma.) Notably, λ and μ are at worst roots of polynomials of P' , which has degree less than a , so $\lambda, \mu \in A_1$!

As such, we have $\lambda < a < \mu$ with $\lambda, \mu \in A_1$. Note that $P(\lambda)$ and $P(\mu)$ have different sign: certainly these are not roots, so we have sign in $\{\pm 1\}$, and if they had the same sign, say they are both of sign $P'(a)$, then P being locally strictly monotone at a will produce a root either between λ and a or between a and μ , which contradicts the construction of λ and μ .

The point is that the data (P, λ, μ) uniquely determine a , and these are data we can push through the isomorphism $f: A_1 \rightarrow A_2$. Namely, the sign of $P(\lambda)$ and $P(\mu)$ continue to be different after passing through our isomorphism, so the intermediate value property in \mathcal{R}_2 grants us some $b \in \mathcal{R}_2$ between $f(\lambda)$ and $f(\mu)$. So we get an isomorphism of fields

$$A_1[a] \cong \frac{A_1[x]}{(P(x))} \cong \frac{A_2[x]}{(P(x))} \cong A_2[b].$$

We will later upgrade this to an isomorphism of ordered fields, which will complete the argument in this case.

Before running this check, though, let's take care of the transcendental case. Add a new constant symbol b to our language. We claim that

$$\text{elDiag}(\mathcal{R}_2) \cup \{f(\alpha) < b : \alpha \in \text{Cut}^-(a/A_1)\} \cup \{b < f(\beta) : \beta \in \text{Cut}^+(a/A_1)\} \quad (2.2)$$

is satisfiable. It's enough to check that this is finitely satisfiable. Upon using the linear order in A_1 , it is enough to check that there is $b^{\mathcal{R}_2} \in \mathcal{R}_2$ with $f(\alpha) < b^{\mathcal{R}_2} < f(\beta)$ for some $\alpha \in \text{Cut}^-(a/A_1)$ and $\beta \in \text{Cut}^+(a/A_1)$, for which $\frac{1}{2}(\alpha + \beta)$ will do. Now, let \mathcal{R}_2^* model (2.2); by construction, \mathcal{R}_2 , and we let b denote the interpretation of the corresponding constant, and we get an isomorphism $A_1[a] \cong A_2[b]$. (Note that we can promise b is also transcendental because of yet another compactness argument avoiding the root of any polynomial.) Now choose some $\lambda, \mu \in A_1$ so that $\lambda < a < \mu$, provided they exist.

We now check that our field isomorphism $A_1[a] \cong A_2[b]$ extends to an isomorphism of ordered fields. Well, for any $Q \in A_1[x]$ such that $\deg Q < \deg P$ (take $\deg P = +\infty$ in the transcendental case), we need to check that $\text{sgn } Q(a) = \text{sgn } Q(b)$. Quickly, if $a > A_1$ always, then the sign of $Q(a)$ is the sign of the leading coefficient (we have gone off to infinity), and $f(a) > A_2$ also, so the sign of $Q(b)$ is also the sign of the leading coefficient. The case of $a < A_2$ is similar.

Now, we may recall that we have some extra information λ and μ . Certainly $Q(a) \neq 0$ because $\deg Q < \deg P$. Without loss of generality, we take $Q(a) > 0$. Now, all roots of Q will live in A_1 by our induction, so we let λ_Q denote the maximum of λ and also all the roots y of Q with $y < a$, and we construct μ_Q dually as a minimum greater than a . Now, Q has no roots between λ_Q and μ_Q by construction, so the intermediate value property promises that Q maintains sign over this entire interval, and this sign is the sign of $Q((\lambda_Q + \mu_Q)/2) \in A_1$. The same holds over in A_2 , and we note that this sign will agree with $Q((f(\lambda_Q) + f(\mu_Q))/2) \in A_2$. So the sign of $Q(a)$ is the same as the sign of $Q(f(a))$.

2.6.2 \mathcal{o} -Minimality

In life one might want explicitly eliminate quantifiers, perhaps with few quantifiers and modest complexity. For this, one can use cell decomposition.

Definition 2.43 (*o-minimal*). A theory T in a language \mathcal{L} extending the language of ordered sets is *o-minimal* if and only if the following conditions are satisfied.

1. T restricted to the language of ordered sets is equivalent to DLO
2. Any model $\mathcal{R} \models T$ with an \mathcal{L}_R -formula $\varphi(x)$ has some partition $-\infty = a_0 < a_1 < \dots < a_n = +\infty$ and subsets $I \subseteq \{1, 2, \dots, n-1\}$ and $J \subseteq \{0, \dots, n\}$ such that

$$\mathcal{R} \models \varphi \leftrightarrow \left(\bigvee_{i \in I} x = a_i \vee \bigvee_{j \in J} (a_j < x < a_{j+1}) \right).$$

Remark 2.44. A boolean combination of sets of the form points plus intervals will again then be a boolean combination of sets plus intervals. So if T eliminates quantifiers, we may as well assume that $\varphi(x)$ is quantifier-free and hence atomic for the second check.

2.7 October 12

The exam is in a week. I'm probably going to fail.

2.7.1 More on o-Minimality

Let's check something.

Theorem 2.45. The theory RCF is *o-minimal*.

Proof. We already know that any model restricted to the language of ordered sets is a dense linear order. So we need to check that the definable subsets of a model $\mathcal{R} \models \text{RCF}$ given by a one-variable \mathcal{L}_R -formula $\varphi(x)$ has the partition as needed. By quantifier elimination, we may as well assume that $\varphi(x)$ is quantifier-free, so upon taking boolean combinations, we may as well assume that $\varphi(x)$ is atomic. Well, we note that an atomic formula is equivalent to one of the form $f(x) > 0$ or of the form $g(x) = 0$ where f and g are polynomials; the point is that a general atomic formula is "a term equals or is bigger than some other term."

- In the case $g(x) = 0$, we are looking at either a discrete set of points or all of R , both of which are of the needed form.
- In the case $f(x) > 0$ (where f is nonzero), we note that the intermediate value property has that $f(x) > 0$ is the union of some intervals whose endpoints are roots of $f(x)$. Explicitly, enumerate the roots as $a_1 < a_2 < \dots < a_n$, and we note that $f(x) > 0$ for some x between a_i and a_{i+1} implies that the entire interval will have $f(x) > 0$, and sign changes for f can only take this form.

The above checks complete the proof. ■

We should probably prove the fundamental theorem of *o-minimality*, which is cell decomposition. This requires the notion of a cell.

Definition 2.46 (cell). Fix a model \mathcal{R} of an o -minimal theory T . Then a *cell* is defined as follows.

- A 0-cell is a point.
- A 1-cell in \mathcal{R} is a set of the form (a, b) where $-\infty \leq a < b \leq \infty$.
- From n , an $(n + 1)$ -cell in \mathcal{R}^{n+1} is a set of one of the following forms.
 - We can have

$$\{(x_1, \dots, x_n, y) : (x_1, \dots, x_n) \in X \text{ and } y = f(x_1, \dots, x_n)\}$$
 where $X \subseteq \mathcal{R}^n$ is an n -cell and $f: X \rightarrow \mathcal{R}$ is continuous and definable.
 - We can have $(-\infty, f)_X$ or $(f, g)_X$ or $(g, \infty)_X$ where

$$(f, g)_X := \{(x_1, \dots, x_n, y) : f(\bar{x}) < y < f(\bar{y})\}$$
 where X is an n -cell and $f, g: X \rightarrow \mathcal{R}$ is continuous and definable with $f(\bar{x}) < g(\bar{x})$ always (where $(-\infty, f)_X$ and $(g, \infty)_X$ are defined analogously).
 - Lastly, we can have all of \mathcal{R}^n .

Remark 2.47. An induction shows that n -cells are homeomorphic to open n -balls when \mathcal{R} is $\mathbb{R} \models \text{RCF}$.

We can now define a cell decomposition.

Definition 2.48 (cell decomposition). Fix a model \mathcal{R} of an o -minimal theory T . Then a *cell decomposition* \mathcal{C} of \mathcal{R}^n is a finite set of cells in \mathcal{R}^n such that

$$\mathcal{R}^n = \bigcup_{c \in \mathcal{C}} c.$$

Anyway, here is our theorem.

Theorem 2.49. Fix a model \mathcal{R} of an o -minimal theory T .

- (a) Given a finite collection $X_1, \dots, X_m \subseteq \mathcal{R}^n$ of definable subsets, then there is a cell decomposition \mathcal{C} of \mathcal{R}^n such that each X_i is a union of some of these cells.
- (b) Any definable function $f: \mathcal{R}^n \rightarrow \mathcal{R}$ is piecewise continuous. In other words, there is a cell decomposition \mathcal{C} of \mathcal{R}^n such that f is continuous upon restriction to each cell.

Remark 2.50. The above theory is true even if we only assume that \mathcal{R} is o -minimal, which lets us prove that T is then o -minimal! We will not prove this stronger notion because it would take more time than we want to spend.

We will prove (a) and (b) essentially simultaneously by some kind of awkward induction.

To get us started, we need the following lemma.

Lemma 2.51. Fix a model \mathcal{R} of an o -minimal theory T . Given some \mathcal{L}_R -formula $\varphi(x, y_1, \dots, y_n)$, there is a bound B (depending only on φ) such that

$$\#\partial\{a \in \mathcal{R} : \mathcal{R} \models \varphi(a, \bar{b})\} \leq B$$

for any $\bar{b} \in \mathcal{R}^n$.

Proof. Note that T is o -minimal implies any definable subset X' of a model $\mathcal{R}' \models T$ has ∂X equal to a finite set of points; namely, choose the formula $\varphi(x)$ defining X' , and then use the hypothesis so that X' becomes a set of points plus some intervals, whose boundary is just a finite set of points.

Now, to continue our proof, fix some \bar{b} , and define $X_{\bar{b}} \subseteq \mathcal{R}$ to be the set defined by $\varphi(x, \bar{b})$. We note that $\partial X_{\bar{b}}$ is definable as saying that $y \in \partial X_{\bar{b}}$ if and only if $y \in \overline{X_{\bar{b}}}$ and $y \in \overline{\mathcal{R} \setminus X_{\bar{b}}}$. However, we can describe the closure of a definable set X (defined by $\psi(x)$) by saying that any interval around $y \in \mathcal{R}$ hits X , which can be said as

$$\forall y_- \forall y_+ ((y_- < y < y_+) \rightarrow \exists x ((y_- < x < y_+) \rightarrow \psi(x))).$$

Intersecting, we can define our boundary.

Now, if the lemma were false, then the theory of

$$\text{elDiag } \mathcal{R} \cup \{\# \partial X_{\bar{b}} \geq N\},$$

where \bar{b} have been taken to be some new constants, is finitely satisfiable and hence satisfiable. So compactness provides an elementary extension \mathcal{R}' where $\# \partial X_{\bar{b}}$ is infinite, which contradicts our initial hypothesis. Notably, \mathcal{R}' will still satisfy T because $\mathcal{R} \leq \mathcal{R}'$. ■

Remark 2.52. If we only took \mathcal{R} to be o -minimal instead of the full theory, then the above lemma is actually the hardest part of the proof. Notably, we used that the theory is o -minimal at the end of the proof.

Remark 2.53. This is essentially the typical use of compactness: we know that some value is always finite, so it cannot be arbitrarily large lest compactness enforce infinity.

Alright, let's start proving Theorem 2.49.

Proof of Theorem 2.49 at $n = 1$. We show (a) and (b) separately. For (a), this is essentially the statement of o -minimality. Each of the X_{\bullet} is a finite union of points or intervals whose endpoints live in ∂X_{\bullet} , so we take

$$F := \bigcup_{i=1}^m \partial X_{\bullet}.$$

Write $F = \{a_1, \dots, a_{n-1}\}$ with $a_1 < a_1 < \dots < a_{n-1}$, and add in $a_0 := -\infty$ and $a_n := \infty$. Then our cell decomposition is F plus the intervals (a_i, a_{i+1}) for each i . Then we can write X_{\bullet} as required as points or unions of intervals from points in ∂X_{\bullet} , so we are done.

Now, (b) is harder. Let $f: \mathcal{R} \rightarrow \mathcal{R}$ is definable. We will actually show that f is actually piecewise continuous and either constant or strictly monotone; i.e., there is a cell decomposition \mathcal{C} such that $f|_{\mathcal{C}}$ is continuous and either constant or strictly monotone. The point is that continuity (and monotonicity) can be expressed as a first-order sentence, so this should approximately happen only finitely many times. Anyway, we proceed in steps.

1. We begin by noting that continuity is actually implied by other assumptions. Suppose $f: \mathcal{R} \rightarrow \mathcal{R}$ is definable is piecewise strictly monotone or constant; then we claim that f is piecewise continuous. If f is constant on a cell, then f is of course continuous there, so we just need to worry about being strictly monotone. Also, if the cell is a point, there is nothing to do.

So without loss of generality, let $I := (a, b)$ be an interval on which f is strictly increasing, and we need to show that we can finitely subdivide the interval to make f continuous. Now, the main point is that $f(I)$ is definable, defined by the sentence $\varphi(y) = \exists x f(x) = y$, so o -minimality tells us that it is a finite union of points and intervals. Further, f is strictly increasing and hence injective, so $f(I)$ is infinite, so f has some open intervals in this image.

Now, it is enough to check that f is discontinuous at only finitely many points. Well, if f were discontinuous at infinitely many points, we note that the points of discontinuity is definable and hence a

finite union of points and intervals, so f will actually be discontinuous everywhere in an interval inside I . Re-applying the above logic to this new smaller interval, so f is discontinuous everywhere on I even though $f(I)$ has some intervals in its image.

Now, let $J' \subseteq f(I')$ is an interval, and we want to show that f is continuous. Well, suppose $f(x) \in J'$. Then for any $\varepsilon_1 < f(x) < \varepsilon_2$ where $(\varepsilon_1, \varepsilon_2) \subseteq J'$, we choose $\delta_i = f^{-1}(\varepsilon_i)$ for each i , and monotonicity implies that $\delta_1 < x' < \delta_2$ implies $\varepsilon_1 < f(x') < \varepsilon_2$, as needed.

We will complete the proof next class. ■

2.8 October 17

There is an exam on Thursday. It will be about four or five questions similar to ones on the homework but hopefully of the more reasonable kind (namely, solvable in something like 20 minutes).

2.8.1 The Cell Decomposition Theorem

Let's continue the proof from last class. We continue with our definable function f which we are trying to show is piecewise constant or strictly monotone.

- For notation, let $\lambda(x)$ be a function outputting $+$ if f is locally increasing to the left of x , $-$ if f is locally decreasing to the left of x , 0 if f is locally constant to the left of x , and $*$ otherwise. We define $\mu(x)$ to the needed right versions of these properties. This produces 16 cases for the pair $(\lambda(x), \mu(x))$.

Quickly, we argue that $*$ is in fact never outputted; this follows from o -minimality. By symmetry, we might as well argue this for λ . Define $\rho(y, z)$ to be $<$ or $>$ or $=$ depending on how y and z relate. Now, if $\lambda(x) = *$, then any $\delta > 0$ produces y and z between $x - \delta$ and x such that $\rho(f(y), f(x)) \neq \rho(f(z), f(x))$. This allows us to build ascending sequences $\{y_i\}_{i=0}^\infty$ and $\{z_i\}_{i=0}^\infty$ (always less than x) such that $\rho(f(y_i), f(x)) \neq \rho(f(z_i), f(x))$ always. By the pigeonhole principle, we may reduce to a subsequence so that $\rho(f(y_i), f(x))$ and $\rho(f(z_i), f(x))$ are each constant and not equal. However, the definable set

$$\{y < x : \rho(f(y), f(x)) = \rho(f(y_i), f(x)) \text{ for each } i\}$$

cannot be the union of finitely many intervals because the sequence $\{z_i\}$ puts infinitely many holes in it around x . Explicitly, any interval containing infinitely many of the y_\bullet (which must be possible by the pigeonhole principle) will also contain infinitely many of the z_\bullet .

So we have left to deal with the 9 cases for $(\lambda(x), \mu(x))$.

- Continuing, by o -minimality, we may decompose \mathcal{R} into intervals and points so that λ and μ are both constant on these intervals, essentially using all 9 cases. Let I be such an interval upon which λ and μ are constant. We would like to show that λ and μ are the same on I . Looking locally, we may as well assume that I is a bounded interval.

For example, take $\lambda(x) = 0$ on I ; a similar argument works if $\mu(x) = 0$. We claim that $\mu(x) = 0$ for each $x \in I$. Then we get $\delta < x$ to check constant to the left of x , we may as well assume that $\delta \in I$, and we are promised that $f(y) = f(x)$ for all y between δ and x . Now, choosing any $x > \varepsilon > y$ such that $\rho(f(z), f(y))$ corresponds to $\mu(y)$ for any $y < z$, which is doable because μ is constant in this region. But then $f(z) = f(x)$ is forced because $\delta < z < x$, and $f(x) = f(y)$, so $f(z) = f(y)$, so $\mu(y) = 0$ follows. But μ is constant on I , so μ vanishes everywhere on I .

Next up, suppose λ is $+$ on I but μ is $-$ on I ; in other words, every point is a local maximum!³ The other cases will be analogous. We claim that for all sufficiently large x , there is some $y > x$ such that $f(y) \geq f(x)$. Well, let B be the set of all $x \in I$ such that all $y > x$ have $f(y) < f(x)$. Now, if our claim were false, then B would have infinitely many elements and hence contain an interval. But then looking locally at some point in the interval would require that $\lambda(y) = -$, which is a contradiction.

³ This is weird but not immediately a contradiction: the function $\mathbb{Q} \rightarrow \mathbb{Q}$ defined by $p/q \mapsto 1/q$ (where $\gcd(p, q) = 1$ and $q > 0$) has every point as a local maximum. We will have to use o -minimality.

So because we worked with x large enough, we might as well replace I with an interval upon which the claim was true. We now claim that each $x \in X$ and y sufficiently larger than x will have $f(y) > f(x)$. Well, consider the set

$$B := \{y > x : f(y) < f(x)\}.$$

Because μ is constantly $-$, this set is nonempty; similarly, the complement is nonempty. Further, the set is certainly definable, so we may let z denote the largest boundary point. If $f(x) > f(z)$, then one finds that $f(x) > f(u)$ for some u close but below x , say above $\delta < x$. But then an interval consisting of elements from δ to a little above z does not live in the complement of B , so z is in fact not a boundary point. (Namely, everything after z needs to be less than $f(x)$.)

What?

As such, $f(x) \leq f(z)$. By the previous claim, we produce $w \geq z$ such that $f(z) \leq f(w)$. Further, if there is $\delta < z$ such that each u between δ and z such that $f(u) < f(x)$. But then any v between z and u has $f(v) < f(x)$, so again B contains points beyond z , which is a contradiction.

Similarly, if one has $f(v) \geq f(x)$ for all $v > z$, then because μ is $+$ on I , we see that the set of all v such that $f(v) = f(x)$ must be finite, so there is $t \geq z$ such that $u > t$ implies $f(u) < f(x)$.

What?

We are now ready to define a function $\beta: I \rightarrow I$ sending $x \in I$ to the least element of the set B_x consisting of $y > x$ such that all $z > y$ has $f(z) > f(x)$, which exists by what we've just shown. Quickly, note that there is $\delta < \beta(x)$ such that any w between δ and $\beta(x)$ has $f(w) > f(x)$. Indeed, if no such thing exists, then instead there is some $\delta < \beta(x)$ such that any w between δ and $\beta(x)$ has $f(w) < f(x)$. Choosing any such w will violate the fact that $\beta(x)$ is supposed to be the infimum of B_x .

So we have a property $\theta_{-,+}(v)$ such that we have $\delta_1 < x < \delta_2$ with $\delta_1 < u < v < w < \delta_2$ has $f(u) < f(w)$. We have checked that $\theta_{-,+}(\beta(x))$ by the above argument.

Now, β is definable, so $\beta(I)$ is definable, satisfying $\beta(x) > x$ (and hence infinite), so we can find an interval $J \subseteq \beta(I)$ which is a "cofinal" interval, meaning that any point in I has a larger point living in J . Because J lives in the image of β , we see that, in J , having $\mu(v) = -$ and $\lambda(v) = +$ implies $\theta_{-,+}(v)$. Now, we go ahead and replace J with I because we can.

As a weird trick, we now reverse the ordering and rerun all our arguments. For example, any sufficiently small x has some $y < x$ such that $f(y) \geq f(x)$, and we are able to restrict J to a "coinitial" interval upon which the above statement is true. Continuing, we can show as before that any $x \in J$ and y sufficiently smaller than x has $f(y) > f(x)$, so we are able to define a function α equal to the supremum of all y such that any $z < y$ has $f(z) > f(x)$. As before, we are able to find an interval $K \subseteq \alpha(J)$, and we again get the analogous property $\theta_{+,-}$ everywhere on K . But this is a contradiction because we already have $\theta_{-,+}$.

4. Thus, we have shown that any interval I as defined at the top of the previous step has $\mu = \lambda$ if μ and λ are constant. It remains to show that f is strictly increasing or strictly decreasing or constant on such an interval. The constant case is relatively easy, so without loss of generality, we take $\lambda = \mu = +$. Well, select $x \in I$, and define

$$B_x := \{y > x : f(y) > f(x)\}.$$

Certainly B_x is nonempty because $\mu = +$. We would like to show that B_x contains everything above x . If there is an element of I bigger than x but not in B_x , we may as well as choose some z the minimum of the boundary of B_x . If $f(z) \leq f(x)$, then everything between z and x must have the same value, but this is not okay because we are locally increasing at z . Similarly, if $f(z) > f(x)$, we note that locally increasing at z causes similar problems.

We can now prove (a) of Theorem 2.49, assuming (b). Namely, suppose that definable functions $\mathcal{R}^n \rightarrow \mathcal{R}$ are piecewise continuous, and we prove the cell decomposition theorem in \mathcal{R}^{n+1} . Well, suppose $X \subseteq \mathcal{R}^{n+1} = \mathcal{R}^n \times \mathcal{R}$ is definable. Then for $b \in \mathcal{R}^n$, we define X_b to be the set of $a \in \mathcal{R}$ such that $(b, a) \in X$; note that X_b is a definable subset of \mathcal{R} .

Now, we note that there is an upper bound N (only depending on X) such that each $b \in \mathcal{R}^n$ with $\# \partial X_b \leq N$; this is by some compactness argument. Then we can choose $B_0, \dots, B_N \subseteq \mathcal{R}^n$ such that B_i is the set of b with ∂X_b having i elements. Now, for any $x \in \mathcal{R}$, define a function $g_i: \mathcal{R}^n \rightarrow \mathcal{R}$ as sending b to the i th element of ∂X_b , which is definable and hence piecewise continuous. Namely, one has a cell decomposition

\mathcal{C}' of \mathcal{R}^n such that $g_i|_{\mathcal{C}'}$ is continuous for each i , and we may as well assume that the \mathcal{C}' decomposes the B_\bullet . One can now decompose X using \mathcal{C}' . Explicitly, take $\widehat{\mathcal{C}}$ to be the graphs of the g_i on C for each $C \in \mathcal{C}'$ and also the cells between the g_\bullet s (and also the cells below g_1 and the cell above g_N).

THEME 3

TYPES

I felt profoundly stupid in that moment and he has a PhD in SYNTAX

—Beth Piatote, [Pia]

3.1 October 26

Today we begin discussing types. The final will be a three-day take-home exam during finals week.

3.1.1 Introducing Types

Let's give some examples to motivate types.

Example 3.1. Note that $(\mathbb{N}, 0, s) \leq (\mathbb{N} \sqcup \mathbb{Z}, 0, s)$, where s denotes the successor function (\mathbb{Z} is placed "after" \mathbb{N}). The point is that the theory is the theory of an infinite set with an injective function with no cycles such that only 0 is not in the image of s . This theory eliminates quantifiers, so it is model-complete.

Example 3.2. Note $(\overline{\mathbb{Q}}, 0, 1, +, \times) \leq (\mathbb{C}, 0, 1, +, \times)$ because these are algebraically closed fields. This theory eliminates quantifiers, so it is model-complete.

However, we would still somehow like to tell these structures apart despite being elementarily equivalent. In the case of $\mathbb{N} \subseteq \mathbb{N} \sqcup \mathbb{Z}$, we note that any $a \in \mathbb{N}$ has $a = s^k(0)$ for some k ; equivalently, $\mathbb{N} \sqcup \mathbb{Z}$ has some $a \in \mathbb{N} \sqcup \mathbb{Z}$ such that

$$a \neq s(s(\cdots(s(0))\cdots)).$$

Namely, take anything in the alternate copy of \mathbb{Z} . Similarly, \mathbb{C} has some element $t \in \mathbb{C}$ such that t is not the root of any polynomial with \mathbb{Z} -coefficients. The point is that we want to look more locally at the formulae satisfied by some particular elements of our models.

This motivates the following definition.

Definition 3.3 (type). Fix an \mathcal{L} -structure \mathcal{M} . Given $\bar{a} \in M^n$, we define the *type* $\text{tp}_{\mathcal{L}}^{\mathcal{M}}(\bar{a})$ to be the set of all \mathcal{L} -formulae $\varphi(\bar{x})$ with n free variables such that $\mathcal{M} \models \varphi(\bar{a})$. For a subset $A \subseteq M$, we may abbreviate $\text{tp}_{\mathcal{L}}^{\mathcal{M}}(a)$ to $\text{tp}_A^{\mathcal{M}}(a)$.

The point is that elements of $\mathbb{N} \sqcup \mathbb{Z}$ achieves types which \mathbb{N} does not. Similarly, elements of \mathbb{C} achieves types which $\overline{\mathbb{Q}}$ does not.

However, it is important that we are considering all the formulae at once.

Proposition 3.4. Fix \mathcal{L} -structures $\mathcal{M} \leq \mathcal{N}$, and fix $\bar{b} \in N^n$. For any finite subset $\Delta \subseteq \text{tp}^{\mathcal{N}}(\bar{b})$, there exists $\bar{a} \in M^n$ such that $\Delta \subseteq \text{tp}^{\mathcal{M}}(\bar{a})$.

Proof. Translating, we are asking for

$$\mathcal{M} \models \exists \bar{x} \left(\bigwedge_{\varphi(\bar{x}) \in \Delta} \varphi(\bar{x}) \right).$$

However, the construction of \bar{b} promises

$$\mathcal{N} \models \exists \bar{x} \left(\bigwedge_{\varphi(\bar{x}) \in \Delta} \varphi(\bar{x}) \right),$$

so we are done because $\mathcal{M} \leq \mathcal{N}$. ■

Remark 3.5. The proof above tells us that it is enough for the extension $\mathcal{M} \subseteq \mathcal{N}$ to merely be “existentially closed,” meaning that existential formulae go down.

We can even go the other way.

Proposition 3.6. Fix an \mathcal{L} -structure \mathcal{M} , and let Δ be a set of \mathcal{L} -formulae with at most one free variable x such that any finite subset $\Delta_0 \subseteq \Delta$ has $\mathcal{M} \models \exists x \bigwedge_{\varphi \in \Delta_0} \varphi(x)$. Then there is an elementary superstructure \mathcal{N} of \mathcal{M} such that there is $a \in \mathcal{N}$ with $\Delta \subseteq \text{tp}^{\mathcal{N}}(a)$.

Proof. Add a new constant symbol a to our language. Let Φ denote the set of sentences $\varphi(a)$ for any $\varphi \in \Delta$. As usual, we want to know that $\text{elDiag } \mathcal{M} \cup \Phi$ is satisfiable. Well, by compactness, it is enough to show that $\text{elDiag } \mathcal{M} \cup \Phi_0$ is satisfiable for any finite subset $\Phi_0 \subseteq \Phi$. But \mathcal{M} will do: certainly $\mathcal{M} \models \text{elDiag } \mathcal{M}$, and by hypothesis we have

$$\mathcal{M} \models \exists x \bigwedge_{\varphi \in \Delta_0} \varphi(x)$$

for the subset $\Delta_0 \subseteq \Delta$ corresponding to Φ_0 by replacing a back with x . So we interpret a in \mathcal{M} to be the element promised by the above satisfaction.

Thus, $\text{elDiag } \mathcal{M} \cup \Phi$ is finitely satisfiable and hence satisfiable, so we produce an elementary superstructure \mathcal{N} of \mathcal{M} with $\mathcal{N} \models \Phi$. So $a^{\mathcal{N}}$ is the desired element with $\Delta \subseteq \text{tp}^{\mathcal{N}}(a^{\mathcal{N}})$, as desired. ■

3.1.2 Types with Parameters

Even using types, it is difficult to tell $\mathbb{N} \sqcup \mathbb{Z}$ apart from $\mathbb{N} \sqcup \mathbb{Z} \sqcup \mathbb{Z}$, and it is difficult to tell \mathbb{C} apart from $\overline{\mathbb{C}(t)}$. Namely, the problem is that the formulae in our languages are not using the full power of the models we gave them. For example, $\mathbb{N} \sqcup \mathbb{Z} \sqcup \mathbb{Z}$ has elements which are not reachable from $\mathbb{N} \sqcup \mathbb{Z}$, but one can only say this by using parameters from $\mathbb{N} \sqcup \mathbb{Z}$. So we refine our definition of types.

Definition 3.7 (type). Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Then an n -type is a set P of \mathcal{L}_A -formulae with n free variables such that $P \cup \text{Th}_A(\mathcal{M})$ is satisfiable.

Remark 3.8. For $P \cup \text{Th}_A(\mathcal{M})$ to be satisfiable, compactness tells us that it is enough for it to be finitely satisfiable: namely, it is enough for any finite subset $P_0 \subseteq P$ to have $P_0 \cup \text{Th}_A(\mathcal{M})$ to be satisfiable. For example, it is enough for $\mathcal{M} \models \exists \bar{x} \bigwedge_{\varphi \in P_0} \varphi(\bar{x})$.

We are allowing our n -types to be rather small sets. So we add an adjective to fix this.

Definition 3.9 (complete). Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Then a type P is *complete* if any \mathcal{L} -formula $\varphi(\bar{x})$ with n free variables has either $\varphi(\bar{x}) \in P$ or $\neg\varphi(\bar{x}) \in P$. Otherwise, we say that the type P is *partial*. As notation, we let $S_n^{\mathcal{M}}(A)$ denote the set of all complete n -types.

We would like for our types to be realized by elements of \mathcal{M} , but this need not always be the case (as we have with $\mathbb{N} \leq \mathbb{N} \sqcup \mathbb{Z}$). So we have the following definition.

Definition 3.10 (realizes). Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Given an n -type P , we say that $\bar{a} \in M^n$ *realizes* P if and only if $\mathcal{M} \models \varphi(\bar{a})$ for all $\varphi \in P$. If no such \bar{a} exists for an n -type P , we say that \mathcal{M} *omits* P .

Example 3.11. The set

$$\{x \neq \underbrace{s(s(\cdots(s(0))))}_n : n \in \mathbb{N}\}$$

is a 1-type for $(\mathbb{N}, 0, s)$ (it's satisfiable by the usual compactness argument), but there is no element of \mathbb{N} realizing this type, so this type is omitted. However, this type is realized by elements of \mathbb{Z} in $\mathbb{N} \sqcup \mathbb{Z}$.

We can now immediately generalize Proposition 3.6 to n -types.

Proposition 3.12. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$, and let P be an n -type. Then there is an elementary superstructure \mathcal{N} of \mathcal{M} such that there is $a \in \mathcal{N}$ realizing P .

Proof. Approximately speaking, one can repeat the proof of Proposition 3.6 upon unpacking all the definitions.

As before, it is enough to show that $\text{elDiag } \mathcal{M} \cup P$ is satisfiable, for which it is enough to show that it is finitely satisfiable. Taking conjunctions, we may assume that we are trying to satisfy just two sentences $\varphi(\bar{a}, \bar{b})$ (from $\text{elDiag } \mathcal{M}$) and $\exists \bar{x} \psi(\bar{x}, \bar{a})$ (from P) where $\bar{a} \in A^\bullet$ and $\bar{b} \in \mathcal{M}^\bullet$.

Well, we are given that there is a model \mathcal{N}_0 satisfying $\text{Th}_A(\mathcal{M}) \cup P$. By construction, we are reassured that $\mathcal{N}_0 \models \exists \bar{x} \psi(\bar{x}, \bar{a})$, and we note that

$$\mathcal{N}_0 \models \exists \bar{y} \varphi(\bar{a}, \bar{y})$$

as well because $\exists \bar{y} \varphi(\bar{a}, \bar{y})$ is an \mathcal{L}_A -sentence satisfied by \mathcal{M} . So we interpret the needed constants from \bar{b} as the tuple promised by $\mathcal{N}_0 \models \exists \bar{y} \varphi(\bar{a}, \bar{y})$ to complete the proof. ■

Corollary 3.13. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$, and let P be a subset of \mathcal{L} -formulae with n free variables. Then P is a complete n -type if and only if there is an elementary superstructure \mathcal{N} of \mathcal{M} such that $P = \text{tp}_A^{\mathcal{N}}(\bar{a})$ for some $\bar{a} \in \mathcal{N}^n$.

Proof. Certainly $P = \text{tp}_A^{\mathcal{N}}(\bar{a})$ implies that P is a complete n -type: certainly it is an n -type, and completeness follows because any $\varphi(\bar{x})$ has exactly one of $\mathcal{N} \models \varphi(\bar{a})$ or $\mathcal{N} \models \neg\varphi(\bar{a})$.

Conversely, suppose that P is a complete n -type. Then the previous proposition grants $\mathcal{N} \geq \mathcal{N}$ and $\bar{a} \in \mathcal{N}^n$ such that $P \subseteq \text{tp}_A^{\mathcal{N}}(\bar{a})$. Because P is complete, equality must follow: if $\varphi(\bar{x}) \notin P$, we will have $\neg\varphi(\bar{x}) \in P$, so $\neg\varphi(\bar{x}) \in \text{tp}_A^{\mathcal{N}}(\bar{a})$, so $\varphi(\bar{x}) \notin \text{tp}_A^{\mathcal{N}}(\bar{a})$. ■

3.1.3 Automorphisms

Let's take a moment to discuss automorphisms.

Remark 3.14. Suppose $\sigma: \mathcal{M} \rightarrow \mathcal{M}$ is an \mathcal{L} -automorphism which fixes a subset $A \subseteq M$ pointwise. Then for any $\bar{a} \in M^n$, automorphisms preserving formula satisfaction means that

$$\text{tp}_A^{\mathcal{M}}(\bar{a}) = \text{tp}_A^{\mathcal{M}}(\sigma(\bar{a})).$$

This tells us that automorphisms preserve types.

Example 3.15. However, types are not enough to determine the automorphism orbit of an element of \mathcal{M} . For example, let $\mathcal{N} = (\mathbb{Q}, <)$ and $A := \{1/n : n \geq 1\}$. Now, there is no automorphism switching 0 and 1 while fixing A (being an automorphism must be a homeomorphism for the order topology and thus fix the limit point 0).

However, 0 and 1 have the same type: any \mathcal{L}_A -formula will only use finitely many constants from A , so it is enough to show that $\text{tp}_{A_0}^{\mathcal{M}}(0) = \text{tp}_{A_0}^{\mathcal{M}}(1)$ for any finite subset $A_0 \subseteq A$. But now there is an automorphism switching 0 and 1 while fixing A_0 fixed because there is some positive distance between 0 and A_0 now.

However, our elements are automorphic upon passing to an elementary superstructure.

Proposition 3.16. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Given $\bar{a}, \bar{b} \in M^n$, suppose $\text{tp}_A^{\mathcal{M}}(\bar{a}) = \text{tp}_A^{\mathcal{M}}(\bar{b})$. Then there is an elementary extension $\mathcal{N} \geq \mathcal{M}$ and an automorphism $\sigma: \mathcal{N} \rightarrow \mathcal{N}$ fixing A pointwise and swapping $\sigma(\bar{a}) = \bar{b}$.

Note that Remark 3.14 provides the converse.

3.2 October 31

Today we discuss partial elementary embeddings.

3.2.1 Partial Elementary Embeddings

Last class we stated Proposition 3.16, which we will show today. The main character of the proof will be the following definition.

Definition 3.17 (partial elementary map). Fix \mathcal{L} -structures \mathcal{M} and \mathcal{N} with a subset $A \subseteq M$. Then a map $f: A \rightarrow N$ is a *partial elementary map* if and only if it preserves types: for all \mathcal{L} -formulae $\varphi(\bar{x})$ and $\bar{a} \in A$, we have $\mathcal{M} \models \varphi(\bar{a})$ if and only if $\mathcal{N} \models \varphi(f(\bar{a}))$.

The point is that we want to extend such maps to full elementary embeddings.

Example 3.18. Such extensions are not possible in general. For example, use the elementary substructure $(\mathbb{N}, s) \leq (\mathbb{N} \sqcup \mathbb{Z}, s)$. However, there is a partial elementary map from $\mathbb{N} \subseteq \mathbb{N} \sqcup \mathbb{Z}$ back to all of \mathbb{N} , which cannot be extended to a full elementary embedding simply because there is nowhere for \mathbb{Z} to go!

Somehow the above problem is “set-theoretic” in that (\mathbb{N}, s) is too small to be an elementary superstructure of $(\mathbb{N} \sqcup \mathbb{Z}, s)$. So perhaps we should only hope to have an extension of a partial elementary map after taking an elementary superstructure of \mathcal{N} . In an attempt to do this inductively, we pick up the following lemma.

Lemma 3.19. Fix \mathcal{L} -structures \mathcal{M} and \mathcal{N} with a subset $A \subseteq M$. Suppose $f: A \rightarrow N$ is a partial elementary map. For any $b \in M$, there is an elementary extension \mathcal{N}' of \mathcal{N} and a partial elementary map $g: (A \cup \{b\}) \rightarrow N'$ extending f .

Proof. For convenience, identify A with its image in \mathcal{N} via f . Also, we may assume that $b \notin A$, for otherwise we can take $f = g$. Choose a new constant symbol c , which is where b is going to go. Now, let T be the theory

$$\text{elDiag } \mathcal{N} \cup \{\varphi(c) : \varphi \in \text{tp}_A^{\mathcal{M}}(b)\}.$$

Here, φ is an \mathcal{L}_A -formula. This will complete the proof: namely, let \mathcal{N}' be a model, which we see is an elementary extension of \mathcal{N} , and we define g extending f by defining $g(b) := c^{\mathcal{N}'}$. And by construction we have $\mathcal{M} \models \varphi(\bar{a})$ if and only if $\mathcal{N} \models \varphi(\bar{a})$ for any $\bar{a} \in (A \cup \{b\})$.

We now check that T is satisfiable by compactness: after taking conjunctions, we may reduce the right-hand side to a single formula $\varphi(c)$ where φ is an \mathcal{L}_A -formula, and we note that $\mathcal{M} \models \exists x \varphi(x)$ by hypothesis on b and so $\mathcal{N} \models \exists x \varphi(x)$ because f is a partial elementary map. So \mathcal{N} is the required model by interpreting c to witness $\mathcal{N} \models \exists x \varphi(x)$. ■

And now here is our transfinite induction.

Lemma 3.20. Fix \mathcal{L} -structures \mathcal{M} and \mathcal{N} with a subset $A \subseteq M$. Suppose $f: A \rightarrow N$ is a partial elementary map. Then there is an elementary extension \mathcal{N}' of \mathcal{N} and an elementary embedding $g: \mathcal{M} \rightarrow \mathcal{N}'$ extending f .

Proof. Find a cardinal κ so that we can enumerate $\mathcal{M} = \{a_\alpha : \alpha \in \kappa\}$. Now, define $A_0 := A$ and $\mathcal{N}_0 := \mathcal{N}$ and $f_0 := f$, and we will define a sequence of maps $f_\alpha: A_\alpha \rightarrow \mathcal{N}_\alpha$ for $\alpha \leq \kappa$ by transfinite recursion, arranged so that the following hold for each $\alpha \leq \kappa$.

- $A_\alpha = A \cup \{a_\beta : \beta \in \alpha\}$.
- $\mathcal{N} \leq \mathcal{N}_\alpha$
- f_α is a partial elementary embedding.

These are satisfied by construction at $\alpha = 0$. Well, by the induction, there are two checks we have to do.

- Suppose $\alpha = \beta + 1$ is a successor ordinal. Then Lemma 3.19 allows us to extend f_β up to a partial elementary embedding $f_\alpha: A_\alpha \rightarrow \mathcal{N}_\alpha$ where $\mathcal{N}_\beta \leq \mathcal{N}_\alpha$. Because $\mathcal{N} \leq \mathcal{N}_\beta$ also, we see $\mathcal{N} \leq \mathcal{N}_\alpha$.
- Suppose α is a limit ordinal. Then we note A_α is the union of the A_β for $\beta \in \alpha$, and so we define \mathcal{N}_α as the union of the \mathcal{N}_β for $\beta \in \alpha$. Because we have an ascending chain $\{\mathcal{N}_\beta\}_{\beta \in \alpha}$, it follows that $\mathcal{N}_0 \leq \mathcal{N}_\alpha$. Lastly, we define f_α by extending all the f_β , and f_α is a partial elementary embedding because such a thing can be checked on the level of points of A_β for each β .

So at all stages of our recursion, we know how to keep going. This completes the transfinite induction. ■


3.2.2 Back to Automorphisms

We are now ready to show Proposition 3.16.

Proposition 3.16. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Given $\bar{a}, \bar{b} \in M^n$, suppose $\text{tp}_A^{\mathcal{M}}(\bar{a}) = \text{tp}_A^{\mathcal{M}}(\bar{b})$. Then there is an elementary extension $\mathcal{N} \geq \mathcal{M}$ and an automorphism $\sigma: \mathcal{N} \rightarrow \mathcal{N}$ fixing A pointwise and swapping $\sigma(\bar{a}) = \bar{b}$.

Proof. To begin, note that the function $f: (A \cup \{\bar{a}\}) \rightarrow M$ defined by $f|_A = \text{id}_A$ and $f(\bar{a}) = \bar{b}$ is a partial elementary embedding. This simply holds because \bar{a} and \bar{b} satisfy all the same \mathcal{L}_A -formulae. So Lemma 3.20 produces an elementary extension \mathcal{N}_1 of \mathcal{M} with an elementary extension $f_1: \mathcal{M} \rightarrow \mathcal{N}_1$ extending f .

To continue, $f_1^{-1}: f_1(M) \rightarrow M$ is a partial elementary embedding because f_1 is an elementary embedding, so using Lemma 3.20 produces a full elementary extension \mathcal{M}_1 of \mathcal{M} and an elementary extension $g_1: \mathcal{N}_1 \rightarrow \mathcal{M}_1$ extending f_1^{-1} . Repeating this step, we produce an elementary extension \mathcal{N}_2 of \mathcal{N}_1 with an elementary embedding $f_2: \mathcal{M}_1 \rightarrow \mathcal{N}_2$ which extends g_1^{-1} . Iterating this process, we build the following diagram.



Now, at the end of it all, let \mathcal{N} be the union of all the \mathcal{N}_\bullet , which is also the union of all the \mathcal{M}_\bullet (for suitable notion of union). All these vertical arrows are elementary embeddings, so \mathcal{N} is an elementary extension of \mathcal{M} . Lastly, we realize that the map $\sigma: \mathcal{N} \rightarrow \mathcal{N}$ given by the union of all the f_\bullet s sends $\bar{a} \mapsto \bar{b}$ because we are extending f , and σ will be invertible with inverse given by the union of all the g_\bullet s. ■

Remark 3.21. A careful examination of the above proof reveals that we have actually proven the following: suppose that we have \mathcal{L} -structures \mathcal{M} and \mathcal{N} and a subset $A \subseteq M$ such that there is a partial elementary embedding $f: A \rightarrow N$. Then there are elementary superstructures \mathcal{M}' of \mathcal{M} and \mathcal{N}' of \mathcal{N} with an isomorphism $\sigma: \mathcal{M}' \rightarrow \mathcal{N}'$ extending f . Indeed, the proof of this result is the above proof minus the first two sentences.

Next class we will put a topology on our types.

3.3 November 2

Today we discuss Stone spaces. I will not record any topological background.

3.3.1 The Stone Topology

As usual, we will let \mathcal{M} be an \mathcal{L} -structure, and we let $A \subseteq M$ be a subset. Recall we defined $S_n^{\mathcal{M}}(A)$ to be the set of complete n -types with parameters from A .

Definition 3.22 (Stone topology). Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. For each \mathcal{L}_A -formula $\varphi(\bar{x})$ with n free variables, we define

$$[\varphi] := \{P \in S_n^{\mathcal{M}}(A) : \varphi \in P\}.$$

The *Stone topology* is the topology generated by the $[\varphi]$ as a sub-basis.

Remark 3.23. Note that we are able to bring in some semantics: for a complete type P , observe $\varphi \wedge \psi \in P$ if and only if $\varphi \in P$ and $\psi \in P$, where both directions can be argued by contradiction. For example, if $\varphi \in P$ and $\psi \in P$ but $\varphi \wedge \psi \notin P$, then $\neg(\varphi \wedge \psi) \in P$, but this is now impossible to satisfy along with φ and ψ . These arguments tell us that

$$[\varphi] \cap [\psi] = [\varphi \wedge \psi]$$

Similarly, one can see that $S_n^{\mathcal{M}}(A) \setminus [\varphi] = [\neg\varphi]$ because $\neg\varphi \in P$ if and only if $\varphi \notin P$ by completeness. Combining, we get

$$[\varphi] \cup [\psi] = S_n^{\mathcal{M}}(A) \setminus ([\neg\varphi] \cap [\neg\psi]) = [\neg(\neg\varphi \wedge \neg\psi)] = [\varphi \vee \psi],$$

which of course we could also have proven directly similarly to the argument with \wedge .

In fact, we have a basis.

Lemma 3.24. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. For a given nonnegative integer n , the sets $[\varphi]$ form a basis of a topology on $S_n^{\mathcal{M}}(A)$

Proof. It is enough to show that the intersection of any two basic sets $[\varphi]$ and $[\psi]$ can be written as the union of basic open sets. But this is automatic from Remark 3.23. ■

Remark 3.25. Thus, if \mathcal{L} is finite or even countable, we have provided a countable basis for the topology on $S_n^{\mathcal{M}}(A)$.

So the open sets of the stone topology on $S_n^{\mathcal{M}}(A)$ are unions of the basic open sets $[\varphi]$ for various \mathcal{L} -formulae φ with n free variables. For example, this allows us to explicitly describe convergence: for a net $\{p_\alpha\}_{\alpha \in \Lambda}$, we have $p_\alpha \rightarrow q$ if and only if any basic open set $[\varphi]$ containing q (i.e., $\varphi \in q$), there is some $\lambda \in \Lambda$ such that $\alpha \geq \lambda$ implies $p_\alpha \in [\varphi]$ (i.e., $\varphi \in p_\alpha$).

Let's discuss the topology on $S_n^{\mathcal{M}}(A)$.

Proposition 3.26. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Then $S_n^{\mathcal{M}}(A)$ is totally disconnected when given the Stone topology.

Proof. We show that $S_n^{\mathcal{M}}(A)$ is totally disconnected. In other words, we have to show that singletons are the largest connected sets. So for any set $S \subseteq S_n^{\mathcal{M}}(A)$ with more than one point, we want to show that S is not connected. Well, we are given two points $P, Q \in S$; they are distinct, so find $\varphi \in P$ with $\varphi \notin Q$. But then $S \subseteq [\varphi] \cup [\neg\varphi]$ even though $[\varphi] \cap [\neg\varphi]$ and $P \in S \cap [\varphi]$ and $Q \in S \cap [\neg\varphi]$. So $[\varphi]$ and $[\neg\varphi]$ disconnect S . ■

Remark 3.27. The argument above in fact shows that $S_n^{\mathcal{M}}(A)$ is Hausdorff as well: we have placed any two distinct complete types P and Q into disjoint open subsets $[\varphi]$ and $[\neg\varphi]$ where $\varphi \in P \setminus Q$.

For something a little more interesting, let's use compactness.

Theorem 3.28. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Then $S_n^{\mathcal{M}}(A)$ is compact and totally disconnected when given the Stone topology.

Proof. We show that $S_n^{\mathcal{M}}(A)$ is compact.

1. We translate covers into semantics. Fix a subset of \mathcal{L}_A -formulae Φ . The main claim is that $\{[\varphi] : \varphi \in \Phi\}$ covers $S_n^{\mathcal{M}}(A)$ if and only if $\{\neg\varphi : \varphi \in \Phi\} \cup \text{Th}_A \mathcal{M}$ is not satisfiable.

In one direction, suppose that $\{\neg\varphi : \varphi \in \Phi\} \cup \text{Th}_A \mathcal{M}$ is satisfiable by a structure \mathcal{N} and tuple \bar{a} . Then we let P be the type $\text{tp}_A^{\mathcal{N}}(\bar{a})$. By construction, P is complete, and $\mathcal{N} \models \neg\varphi(\bar{a})$ for each $\varphi \in \Phi$, so we conclude that P is a complete type not covered by one of the $[\varphi]$ for $\varphi \in \Phi$.

In the other direction, suppose that $\{[\varphi] : \varphi \in \Phi\}$ fails to cover $S_n^{\mathcal{M}}(A)$. So choose P which does not live in any $[\varphi]$ for $\varphi \in \Phi$, implying that $\neg\varphi \in P$ for each $\varphi \in \Phi$. Now, Corollary 3.13 grants us some elementary superstructure \mathcal{N} of \mathcal{M} and some $\bar{a} \in N^n$ so that $P = \text{tp}_A^{\mathcal{N}}(\bar{a})$. Thus, by construction, $\mathcal{N} \models \neg\varphi(\bar{a})$ for each $\varphi \in \Phi$ and $\mathcal{N} \models \text{Th}_A \mathcal{M}$ because we have an elementary superstructure, so we are done.

2. Now, Suppose that we have an open cover \mathcal{U} of $S_n^{\mathcal{M}}(A)$ which we would like to reduce to a finite subcover. By writing each open set in \mathcal{U} as a union of basic open subsets, we may assume that \mathcal{U} has only basic open subsets, which we enumerate as $[\varphi]$ for various $\varphi \in \Phi$. We would like to extract a finite subcover. The previous step implies that

$$\{\neg\varphi : \varphi \in \Phi\}$$

fails to be satisfiable, so by compactness, a finite subset fails to be satisfiable, so we have some finite $\Phi_0 \subseteq \Phi$ such that

$$\{\neg\varphi : \varphi \in \Phi_0\}$$

fails to be satisfiable, so by the previous step once again, we see that $\{[\varphi] : \varphi \in \Phi_0\}$ is the needed finite subcover. ■

Remark 3.29. Notably, the main input to the above proof was the compactness theorem! In some sense, this is where the compactness theorem gets its name.

There are a bunch of other functoriality checks one can do with continuous maps.

3.3.2 Isolated Types

Topology motivates the following definition.

Definition 3.30 (isolated). Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Then a type $P \in S_n^{\mathcal{M}}(A)$ is *isolated* if and only if P is isolated in the Stone topology. In other words, there exists an open subset around P only containing P .

Let's get a better understanding of this term.

Proposition 3.31. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Let $P \in S_n^{\mathcal{M}}(A)$ be a type. Then the following are equivalent.

- (a) P is isolated.
- (b) $\{P\} = [\varphi]$ for some \mathcal{L}_A -formula φ .
- (c) There is an \mathcal{L}_A -formula $\varphi \in P$ such that, for any other \mathcal{L}_A -formula ψ , we have $\psi \in P$ if and only if $\text{Th}_A(\mathcal{M}) \models (\varphi \rightarrow \psi)$.

Note that the completeness of $\text{Th}_A(\mathcal{M})$ makes $\text{Th}_A(\mathcal{M}) \models (\varphi \rightarrow \psi)$ equivalent to $(\varphi \rightarrow \psi)$ being in $\text{Th}_A(\mathcal{M})$, which is equivalent to $\mathcal{M} \models (\varphi \rightarrow \psi)$.

Proof. Note that (b) implies (a) by the definition of being isolated. For (a) implies (b), we note that $\{P\}$ is an open set by definition, so we can find some φ such that $P \in [\varphi]$ and $[\varphi] \subseteq \{P\}$ by using our basis, so of course $\{P\} = [\varphi]$ follows.

So the interesting part is showing that (c) is equivalent to the other two. The main claim is that $[\varphi] \subseteq [\psi]$ if and only if $\text{Th}_A(\mathcal{M}) \models (\varphi \rightarrow \psi)$. We show the implications separately.

- In one direction, if $\text{Th}_A(\mathcal{M}) \models (\varphi \rightarrow \psi)$, then $\mathcal{M} \models (\varphi \rightarrow \psi)$, then if $P \in [\varphi]$, we have $\varphi \in P$, so $\psi \in P$ by completeness, so $P \in [\psi]$.
- In the other direction, if \mathcal{M} fails to satisfy $(\varphi \rightarrow \psi)$, then there is some $\bar{a} \in M$ such that $\mathcal{M} \models \varphi(\bar{a}) \wedge \neg\psi(\bar{a})$. Thus, $\text{tp}_A^{\mathcal{M}}(\bar{a}) \in [\varphi] \setminus [\psi]$.

We now show (b) implies (c): we have $\{P\} = [\varphi]$, so $\psi \in P$ if and only if $P \subseteq [\psi]$ if and only if $[\varphi] \subseteq [\psi]$, which by the claim is equivalent to $\mathcal{M} \models (\varphi \rightarrow \psi)$. Lastly, we show (c) implies (b): given our special φ , we want to show that $\{P\} = [\varphi]$. Well, certainly $P \in [\varphi]$. Conversely, if $Q \neq P$ for some complete n -type Q , pick up $\psi \in P \setminus Q$, but then $\mathcal{M} \models (\varphi \rightarrow \psi)$, so by the claim, $[\varphi] \subseteq [\psi]$, but then $Q \notin [\psi]$, so $Q \notin [\varphi]$. ■

Remark 3.32. The point is that isolated types are determined by a single formula. Note that the formula φ yielding P is unique up to equivalence by (c) because then $\text{Th}_A(\mathcal{M}) \models (\varphi \leftrightarrow \psi)$ if $\{P\} = [\varphi] = [\psi]$.

Example 3.33. Let $\mathcal{M} = (\mathbb{R}, 0, 1, +, \times, \leq)$. Then $P = \text{tp}_{\emptyset}^{\mathcal{M}}(A)$ is isolated given by formula $x = 0$.

More generally, we have the following.

Proposition 3.34. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. Suppose that $\bar{b} \in M$ is definable over A . Then $\text{tp}_A^{\mathcal{M}}(\bar{b})$ is an isolated type.

Proof. Well, suppose $\varphi(\bar{x})$ defines \bar{b} , and we claim that $\text{tp}_A^{\mathcal{M}}(\bar{b}) = [\varphi]$, for which we use Proposition 3.31. Well, we see that $\psi(\bar{x}) \in \text{tp}_A^{\mathcal{M}}(\bar{b})$ if and only if $\mathcal{M} \models \psi(\bar{b})$ if and only if $\mathcal{M} \models \forall \bar{x}(\varphi(\bar{x}) \rightarrow \psi(\bar{x}))$, which is equivalent to $\mathcal{M} \models (\varphi \rightarrow \psi)$. To finish, we note that this is equivalent to $\varphi \rightarrow \psi$ living in $\text{Th}_A(\mathcal{M})$, which is equivalent to $\text{Th}_A(\mathcal{M}) \models (\varphi \rightarrow \psi)$ by completeness. ■

In fact, we have the following partial converse.

Proposition 3.35. Fix an \mathcal{L} -structure \mathcal{M} and a subset $A \subseteq M$. If P is an isolated type, then $P = \text{tp}_A^{\mathcal{M}}(\bar{a})$ for some \bar{a} .

Proof. Suppose that $\{P\} = [\varphi]$. But now $\text{Th}_A(\mathcal{M}) \cup \{\exists \bar{x} \varphi(\bar{x})\}$ is satisfiable because this is the same as $\text{Th}_A(\mathcal{M}) \cup P$. So $\exists \bar{x} \varphi(\bar{x}) \in \text{Th}_A(\mathcal{M})$ by completeness, so $\mathcal{M} \models \varphi(\bar{a})$ for some $\bar{a} \in M$. To complete the argument, we note that $\psi \in P$ if and only if $\text{Th}_A(\mathcal{M}) \models (\varphi(\bar{a}) \rightarrow \psi(\bar{a}))$, so $\mathcal{M} \models \psi(\bar{a})$, so $\psi \in \text{tp}_A^{\mathcal{M}}(\bar{a})$. So $P \subseteq \text{tp}_A^{\mathcal{M}}(\bar{a})$, and equality follows by completeness. ■

Example 3.36. One can use the above proposition to show that there are types which aren't isolated in $\mathcal{M} = (\mathbb{R}, 0, 1, +, \times, \leq)$. For example, take the type given by any transcendental.

3.4 November 7

We continue.

3.4.1 Types of Theories

Let's move from types of models \mathcal{M} to types of theories T .

Definition 3.37 (type). Fix an \mathcal{L} -theory T . Then an n -type is a set P of \mathcal{L} -formulae with n free variables such that $P \cup T$ is satisfiable. A type is *complete* if $\varphi \in P$ or $\neg\varphi \in P$ for each \mathcal{L} -formula φ with n free variables. We let $S_n(T)$ denote the set of complete n -types.

Example 3.38. For an \mathcal{L} -structure \mathcal{M} , we have

$$S_n(\text{Th } \mathcal{M}) = S_n^{\mathcal{M}}(\emptyset).$$

The main content here is that we are allowing T to not be complete. Note that $S_n(T)$ has a topology given by the basic open sets

$$[\varphi] := \{P \in S_n(T) : \varphi \in P\}.$$

The checks on this topology are the same as when T is complete; namely, we have defined basis as in Lemma 3.24, and it is totally disconnected as in Proposition 3.26, and it is compact as in Theorem 3.28.

We are also able to provide definitions motivated by this topology.

Definition 3.39 (isolated). Fix an \mathcal{L} -theory T . Then a complete n -type P is *isolated* if and only if there is an \mathcal{L} -formula φ with n free variables such that $\{P\} = [\varphi]$.

Remark 3.40. The argument of Proposition 3.31 generalizes immediately to show that the following are equivalent for a complete n -type P and \mathcal{L} -formula φ with n free variables.

- (a) P is isolated with $\{P\} = [\varphi]$.
- (b) There is an \mathcal{L} -formula $\varphi \in P$ such that, for any other \mathcal{L} -formula ψ , we have $\psi \in P$ if and only if $T \models (\varphi \rightarrow \psi)$.

As before, the main input is to show that $[\varphi] \subseteq [\psi]$ if and only if $T \models (\varphi \rightarrow \psi)$, and the proof of this is quite similar.

So we take the above remark as providing our definition of isolated types.

Definition 3.41 (isolated). Fix an \mathcal{L} -theory T . Then an n -type P is *isolated* if and only if there is an \mathcal{L} -formula φ such that $T \cup \varphi$ is satisfiable and the following holds: for any other \mathcal{L} -formula ψ , we have $\psi \in P$ if and only if $T \models (\varphi \rightarrow \psi)$.

We can also use topology to define a notion of density.

Definition 3.42 (dense). Fix an \mathcal{L} -theory T . Then a set X of complete n -types is *dense* if and only if X intersects each nonempty basic open set of $S_n(T)$. In other words, for each \mathcal{L} -formula φ with n free variables such that $T \cup \{\varphi\}$ is satisfiable, there is a complete n -type $P \in X$ such that $\varphi \in P$.

Example 3.43. The set $S_n(T) \subseteq S_n(T)$ is dense.

Example 3.44. Suppose $T = \text{Th } \mathcal{M}$ for some \mathcal{L} -structure \mathcal{M} . Then the set

$$X := \{P \in S_n^{\mathcal{M}}(\emptyset) : P \text{ is realized in } \mathcal{M}\}$$

is dense in $S_n(T)$. Indeed, suppose φ is an \mathcal{L} -formula with n free variables such that $T \cup \{\varphi\}$ is satisfiable. But $T = \text{Th } \mathcal{M}$ is complete, so $\mathcal{M} \models \exists \bar{x} \varphi(\bar{x})$, so we may find $\bar{a} \in M$ such that $\mathcal{M} \models \varphi(\bar{a})$. Thus, $\text{tp}^{\mathcal{M}}(\bar{a}) \in X$ contains φ , as needed.

Remark 3.45. Let T be an \mathcal{L} -theory. If P is a complete isolated n -type, and $X \subseteq S_n(T)$ is dense, then we claim $P \in X$. Indeed, write $\{P\} = [\varphi]$; then $[\varphi] \cap X \neq \emptyset$ needs $P \in X$, as desired.

Professor Montalban recommends reading types of discrete linear orders and of algebraically closed fields to understand what is going on. Here is a taste for the sort of thing one can show.

Proposition 3.46. Fix a discrete linear order $(\mathcal{M}, <)$, and let $A \subseteq M$ be a subset. Then the types in $S_1^{\mathcal{M}}(A)$ which are not realized in A correspond to a cut (L, U) of A . (Here, $L \cup U = A$ and L is closed downwards and U is closed upwards.)

Sketch. The point is that every formula in a complete n -type is equivalent to a quantifier-free formula, which amounts to requiring some list of satisfiable inequalities. These lists of inequalities amount to a cut. ■

Proposition 3.47. Fix a discrete linear order $(\mathcal{M}, <)$, and let $A \subseteq M$ be a subset. Then a complete 1-type $P \in S_1^{\mathcal{M}}(A)$ not realized in A corresponding to the cut (L, U) of A fails to be isolated if and only if L fails to have a maximum or U fails to have a minimum.

Notably, if $L = \emptyset$ or $U = \emptyset$, then P remains isolated.

Sketch. The point here is that we need to be determined by a single inequality. Being “above L and below U ” being encoded into a single formula requires that L or U contain their supremum or infimum (respectively). ■

3.4.2 Type Omitting

Here is our theorem.

Theorem 3.48 (Type omitting). Fix a countable language \mathcal{L} , and let T be an \mathcal{L} -theory. Further, let P be an n -type which is not isolated. Then there is a countable model $\mathcal{M} \models T$ which omits P .

The “non-isolated” hypothesis on P is necessary: for example, if $T = \text{Th}(\mathbb{N}, 0, s)$, then the type of 0 is always realized, which is notably an isolated type. More generally, isolated types are always realized by Proposition 3.35. Theorem 3.48 above is the converse.

Proof of Theorem 3.48. We do a Henkin construction. Namely, we use an argument like Lemma 1.43 to expand our language to \mathcal{L}^* by adding in new constant symbols \mathcal{C} to our language, and then we extend T to an \mathcal{L}^* -theory T^* to be complete (and satisfiable) and have witnesses. We will also arrange our construction so that each tuple $(c_1, \dots, c_n) \in \mathcal{C}^n$ has some $\varphi \in P$ such that $\neg\varphi(c_1, \dots, c_n)$ is in T^* . Then the construction of Lemma 1.44 produces the needed model \mathcal{M} whose universe is \mathcal{C} modded out by some equivalence relation dictated by T^* . Namely, having $\neg\varphi(c_1, \dots, c_n)$ in T implies that P is omitted because the universe of \mathcal{M} arises exactly from \mathcal{C} .

We will construct T^* to be $T \cup \{\theta_0, \theta_1, \dots\}$ by adding one sentence at a time; by compactness, the satisfiability of T^* follows from the satisfiable at each finite step. (Technically, we will eventually have $T^* \models \varphi$ or $T^* \models \neg\varphi$ for each φ at the end of the construction.) We will also require that $T \models (\theta_{n+1} \rightarrow \theta_n)$ for each n , for psychological reasons. For convenience, we will also need the following enumerations.

- Let $\{\varphi_n : n \in \omega\}$ be an enumeration of all \mathcal{L}^* -sentences.
- Let $\{\psi_n(x) : n \in \omega\}$ be an enumeration of \mathcal{L}^* -formulae with one free variable.
- Enumerate the n -tuples \bar{c} of constants in \mathcal{C} .

We now proceed in steps.

0. On step $s = 0$, we let θ_0 be the sentence $\forall x(x = x)$, which is always true.
1. On steps which are $s+1 = 3i+1$, we deal with completeness. Here, let θ_{s+1} be either $\theta_s \wedge \varphi_i$ or $\theta_s \wedge \neg \varphi_i$, one of which we know is going to be satisfiable with T .
2. On steps which are $s+2 = 3i+2$, we deal with witnesses. Here, we choose a constant $c \in \mathcal{C}$ not in θ_s , and we let θ_{s+1} be the sentence

$$\theta_s \wedge (\exists x \psi_i(x) \rightarrow \psi_i(c)),$$

which continues to be satisfiable by interpreting c to be the needed witness to $\exists x \psi_i(x)$ (if it exists) in a model of θ_s .

3. On steps which are $s+3 = 3i+3$, we deal with omitting P . Let \bar{c} be the i th n -tuple of constants in \mathcal{C} . We would like to find some $\varphi(\bar{x})$ in P such that $T \cup \{\theta_s \wedge \neg \varphi(\bar{c})\}$ is satisfiable.

The point is to contradict the fact that P is not isolated. Write θ_s as $\theta(\bar{d}, \bar{c})$ where \bar{d} are the constants which appear in θ_s despite not appearing in \bar{c} . By the satisfiability of $T \cup \{\theta_s\}$, we see that $T \cup \{\exists \bar{y} \exists \bar{x} \theta(\bar{y}, \bar{x})\}$. But because P is not isolated, there is $\varphi \in P$ such that T does not prove $\forall \bar{y} \exists \bar{x} (\theta(\bar{y}, \bar{x}) \rightarrow \varphi(\bar{x}))$. Thus, there is a model \mathcal{M} of T such that all $\bar{a} \in M$ have some $\bar{b} \in M$ with

$$\mathcal{M} \models (\theta(\bar{b}, \bar{a}) \wedge \neg \varphi(\bar{a})).$$

Interpreting constants in (\bar{d}, \bar{c}) as in (\bar{b}, \bar{a}) , we have shown that $T \cup \{\theta_s \wedge \neg \varphi(\bar{c})\}$ is satisfiable by \mathcal{M} , as required. ■

Remark 3.49. The above proof can also show that we can omit countably many non-isolated types $\{P_n\}_{n \in \mathbb{N}}$ simply by modifying the third step to yield the sentence $\neg \varphi_m(\bar{c})$ where $\varphi_m \in P_n$; the point here is to use the countability of $\mathbb{N} \times \mathbb{N}$.

3.5 November 9

We continue.

Example 3.50. Let $\mathbb{N} = (\mathbb{N}, 0, 1, +, \times, \leq)$, and choose a countable proper elementary extension \mathcal{M} of \mathbb{N} , so for example, \mathcal{M} is still a model of PA. Now, we build a proper elementary extension \mathcal{N} of \mathcal{M} . For example, for each $m \in M$, we can try to omit the type p_m defined by

$$\{(v < m)\} \cup \{(v \neq h) : h < m\}.$$

We could then build a model $\mathcal{N} \models \text{elDiag } \mathcal{M}$ omitting all the types p_m (which are not isolated because they are not realized in \mathcal{M}).

3.5.1 Prime and Atomic Models

For the next few weeks, we examine prime, atomic, and saturated models.

Definition 3.51 (prime). Fix an \mathcal{L} -theory T . Then a model \mathcal{M} of T is a *prime model* if and only if $\mathcal{N} \models T$ implies that there is an elementary embedding $\mathcal{M} \rightarrow \mathcal{N}$.

We would like to know when these exist and that they are unique. These require proof.

Example 3.52. In the theory ACF_p of algebraically closed fields of characteristic p , then $\overline{\mathbb{F}_p}$ is a prime model: it will embed into any other algebraically closed field of characteristic p , and these are elementary extensions because ACF_p eliminates quantifiers and hence is model-complete.

Example 3.53. Similarly, the theory DLO has \mathbb{Q} as a prime model for essentially the same reason.

Remark 3.54. A theory T needs to be complete to have prime models. Namely, suppose \mathcal{M} is a prime model. Then $\mathcal{N} \models T$ implies that $\mathcal{M} \leq \mathcal{N}$, so $\text{Th } \mathcal{M} = \text{Th } \mathcal{N}$. Thus, $\mathcal{M} \models \varphi$ if and only if $\mathcal{N} \models \varphi$ for all models \mathcal{N} of T , which is then equivalent to $T \models \varphi$, so completeness of T follows.

Remark 3.55. Suppose our language \mathcal{L} is countable. Given a prime model \mathcal{M} of T , then $\text{tp}^{\mathcal{M}}(\bar{a})$ must be isolated. Indeed, if not, then there is a model \mathcal{M}' omitting $\text{tp}^{\mathcal{M}}(\bar{a})$, but then the promised elementary extension $\mathcal{M} \leq \mathcal{M}'$ requires there to be an element in \mathcal{M}' with the same type! Note that this implies that the isolated types are dense in $S_n(T)$, which need not be the case in general.

And now for atomic models.

Definition 3.56 (atomic). Fix an \mathcal{L} -theory T . Then a model \mathcal{M} of T is an *atomic model* if and only if $\text{tp}^{\mathcal{M}}(\bar{a})$ is isolated for all $\bar{a} \in M$.

Actually, these are the same.

Proposition 3.57. Let \mathcal{L} be a countable language, and let T be a complete \mathcal{L} -theory with an infinite model. Then a model \mathcal{M} is prime if and only if it is atomic and countable.

Proof. Remark 3.55 tells us that prime models are atomic, and Theorem 1.71 tells us that T does have countable models, so any prime model must be countable in order to embed into such models.

Thus, the difficulty will come from the converse direction. Suppose \mathcal{M} is atomic and countable, and let $\mathcal{N} \models T$. We want to show that there is an elementary embedding $\mathcal{M} \leq \mathcal{N}$. Well, enumerate \mathcal{M} as $\{m_i : i \in \omega\}$. We will create our elementary embedding inductively: namely, we want to define a sequence $n_0, n_1, \dots \in N$ such that

$$\text{tp}^{\mathcal{M}}(m_0, \dots, m_k) = \text{tp}^{\mathcal{N}}(n_0, \dots, n_k).$$

This will imply our elementary embedding: for any $\varphi(\bar{x})$, we see that $\mathcal{M} \models \varphi(m_0, \dots, m_k)$ if and only if $\mathcal{N} \models \varphi(n_0, \dots, n_k)$. (Even if a formula φ does not use every single variable in $\{m_0, \dots, m_k\}$, we might as well include them anyway.) Let's do our induction.

- At step 0, we may find n_0 because the isolated type $\text{tp}^{\mathcal{M}}(m_0)$ is realized in \mathcal{N} by Proposition 3.35.
- At step $k + 1$, we want to find n_{k+1} such that

$$\text{tp}^{\mathcal{M}}(m_0, \dots, m_k, m_{k+1}) = \text{tp}^{\mathcal{N}}(n_0, \dots, n_k, n_{k+1}).$$

The issue here is that we want to find n_{k+1} without adjusting n_0, \dots, n_k . To get around this, we find $\varphi(x_0, \dots, x_k)$ isolate $\text{tp}^{\mathcal{M}}(m_0, \dots, m_k)$. We want n_{k+1} such that $\mathcal{N} \models \varphi(n_0, \dots, n_k)$, but we know that

$$\mathcal{M} \models \exists x_{k+1} \varphi(m_0, \dots, m_k, x_{k+1}),$$

so the sentence $\exists x_{k+1} \varphi(x_0, \dots, x_k, x_{k+1})$ belongs to $\text{tp}^{\mathcal{M}}(m_0, \dots, x_k)$, so the inductive hypothesis implies that it belongs to $\text{tp}^{\mathcal{N}}(n_0, \dots, n_k)$ too, so

$$\mathcal{N} \models \exists x_{k+1} \varphi(n_0, \dots, n_k, x_{k+1}),$$

which provides us with the needed n_{k+1} . ■

If \mathcal{M} and \mathcal{N} are both countable and atomic, one can turn the above argument into a genuine back-and-forth, allowing us to conclude that $\mathcal{M} \cong \mathcal{N}$.

Proposition 3.58. Let \mathcal{L} be a countable language, and let T be a complete \mathcal{L} -theory with an infinite model. Then any two countable and atomic models \mathcal{M} and \mathcal{N} are isomorphic.

Proof. Turn Proposition 3.57 into a back-and-forth argument. Essentially, enumerate both \mathcal{M} and \mathcal{N} and then alternate steps going back and forth to make sure we produce a bijection. ■

Theorem 3.59. Let \mathcal{L} be a countable language, and let T be a complete \mathcal{L} -theory with an infinite model. Then the following are equivalent.

- (a) T has a prime model.
- (b) T has an atomic model.
- (c) The isolated types in $S_n(T)$ are dense for all n .

Proof. Note that (a) implies (b) by Remark 3.55. For (b) implies (a), we note that any atomic model \mathcal{M} by Theorem 1.71 has some $\mathcal{M}_0 \leq \mathcal{M}$ which is countable, but then all types will remain isolated, so Proposition 3.57 completes. Next, (a) implies (c) by Remark 3.55.

Thus, the difficulty comes from showing that (c) implies (b). For the proof, say that a sentence $\varphi(\bar{x})$ is *isolating* if and only if $T \models \exists \bar{x} \varphi(\bar{x})$ and any $\psi(\bar{x})$ has either $T \models \forall \bar{x} (\varphi(\bar{x}) \rightarrow \psi(\bar{x}))$ or $T \models \forall \bar{x} (\varphi(\bar{x}) \rightarrow \neg \psi(\bar{x}))$. Namely, $\varphi(\bar{x})$ implies a complete set of formulae.

Now, consider the set P_n of formulae $\neg \varphi(\bar{x})$ where this is an isolating formula $\varphi(\bar{x})$ with n free variables. This is countably many types, so we would like to use Remark 3.49 to provide a model $\mathcal{M} \models T$ which omits all the types P_n . For this, we must check that P_n is not an isolated (partial) n -type. (If they are not consistent with T , they are omitted automatically, so we don't have to worry about that.)

Well, suppose for the sake of contradiction that P_n is isolated by $\psi(\bar{x})$. This means that $T \models \exists \psi(\bar{x})$ and $T \models (\psi(\bar{x}) \rightarrow \varphi(\bar{x}))$ for each $\varphi(\bar{x})$ in P . But by the hypothesis (c), we may find an isolated n -type Q containing $\psi(\bar{x})$, meaning that it is isolated by the formula $\theta(\bar{x})$, so $T \models \forall \bar{x} (\theta(\bar{x}) \rightarrow \psi(\bar{x}))$. This is contradiction because $\neg \theta(\bar{x})$ lives in P_n , so $T \models \forall \bar{x} (\psi(\bar{x}) \rightarrow \neg \theta(\bar{x}))$, meaning we have shown θ implies $\neg \theta$.

We now complete the proof. Remark 3.49 now grants us a model \mathcal{M} omitting all the types P_n . Thus, each $\bar{a} \in M^n$ cannot realize P_n , means that $\mathcal{M} \models \theta(\bar{a})$ for some isolating formula $\theta(\bar{x})$, but then $\text{tp}^{\mathcal{M}}(\bar{a})$ must be the type isolated by $\theta(\bar{x})$. So \mathcal{M} is an atomic model. ■

Remark 3.60. Now combining with Proposition 3.58 assures us that these prime models are in fact unique.

3.6 November 14

Office hours will be on Tuesdays at 11:30AM.

Example 3.61. Let's build a theory whose isolated types are not dense. Our language \mathcal{L} will contain countably many unary relations $\{P_0, P_1, \dots\}$, and let T be the theory consisting of sentences of the form

$$\exists x \left(\bigwedge_{i \in S} P_i(x) \wedge \bigwedge_{i \notin S} \neg P_i(x) \right),$$

where $S \subseteq \mathbb{N}$ is any finite subset. On the other hand, for each subset $S \subseteq \mathbb{N}$, there is a type P_S consisting of the sentences $P_i(x)$ for $i \in S$ and $\neg P_i(x)$ for $i \notin S$; note $P_S \cup T$ is consistent by compactness. Further, P_S is a complete type, which one can see by showing that T eliminates quantifiers by the usual syntactic arguments; from here, $P_S(x)$ will imply any formula or its negation because one can replace formula by a quantifier-free one. Now, we note that all the 1-types take this form again by the quantifier elimination.

However, T has no isolated types. Indeed, suppose that a type P_S is contained in $[\varphi]$; we will argue that $[\varphi]$ has another 1-type. We may assume that φ is quantifier-free, and we may assume that φ has only conjunctions. The problem is that φ only mentions finitely many of the P_i , so we can find multiple complete 1-types P_S and P_Q living in $[\varphi]$ which complete φ .

3.6.1 Homogeneous Models

Another useful kind of model for our discussion is homogeneous models.

Definition 3.62 (κ -homogeneous). Fix an \mathcal{L} -structure \mathcal{M} . Then \mathcal{M} is κ -homogeneous if and only if the following holds: for any subset $A, B \subseteq M$ of cardinality less than κ equipped with a partial elementary embedding $f: A \rightarrow M$, and given an element $a \in M$, then there is $b \in M$ and some partial elementary embedding $f^*: A \cup \{a\} \rightarrow M$ extending f and sending $f^*: a \mapsto b$. We then say that \mathcal{M} is homogeneous if and only if \mathcal{M} is $|M|$ -homogeneous.

Intuitively, homogeneity allows us to extend partial elementary embeddings from subsets one element at a time. By an inductive argument, one achieves the following.

Proposition 3.63. Fix a homogeneous \mathcal{L} -structure \mathcal{M} . Given subsets $A, B \subseteq M$ of strictly smaller cardinality than M , any partial elementary embedding $f: A \rightarrow M$ extends to an automorphism of \mathcal{M} .

Proof. We do transfinite induction, applying a back-and-forth argument. Enumerate the elements of M by $\{m_\alpha : \alpha \in \kappa\}$, where $\kappa = |M|$. We now build a sequence of partial elementary embeddings $f_\alpha : M_\alpha \rightarrow M$ of partial elementary embeddings satisfying the following.

- $f_0 = f$.
- f_β extends f_α whenever $\beta \geq \alpha$.
- $|\text{im } f_\alpha| \leq |A| + 2\alpha$.
- m_α is in the domain and range of $f_{\alpha+1}$.

Taking the union of the f_α will complete the proof. Indeed, the union of partial elementary maps is a partial elementary map, but the union now contains all of M in the domain and codomain.

We quickly deal with limit stages in our induction first. Namely, if α is a limit ordinal, we define f_α as the union of all the previous f_β s. This is a union of partial elementary embeddings f_β , so f_α is a partial elementary embedding too. The $\alpha + 1$ check has no content, and the extensions are satisfied by construction. Lastly, the size of the image of the $\text{im } f_\alpha$ s is the supremum of all the $\text{im } f_\beta$ s, which is upper-bounded by the supremum of all the $|A| + 2\beta$, which is $|A| + 2\alpha$.

We now must argue the successor stage. Suppose we are given f_α , and we must construct g_α . To add $m_{\alpha+1}$ to the domain, we use the homogeneity of \mathcal{M} . On the other hand, applying the same argument to

the inverse $g_\alpha^{-1}: \text{im } g_\alpha \rightarrow \text{dom } g_\alpha$ allows us to extend g_α^{-1} to the new element $m_{\alpha+1}$ in the image, which is exactly the $f_{\alpha+1}$ we needed. Notably, we have only added two elements in total, so the inequality on $\text{im } f_{\alpha+1}$ is still satisfied. ■

Non-Example 3.64. Consider $A := \mathbb{Q}^{2\mathbb{Z}}$ as a subspace of $M := \mathbb{Q}^{\mathbb{Z}}$. But then there is a partial elementary embedding $A \cong M$, and it cannot be extended to an automorphism because it is already surjective! Note that \mathcal{M} is in fact homogeneous: one may assume that A is a full subspace, and then one extends to a single extra point in \mathcal{M} arbitrarily as long as the single extra point is away from A .

Here's an example.

Lemma 3.65. Fix a countable language \mathcal{L} and an \mathcal{L} -structure \mathcal{M} . If \mathcal{M} is atomic, then \mathcal{M} is \aleph_0 -homogeneous.

Proof. Fix a finite subset $A \subseteq M$ and a partial elementary embedding $f: A \rightarrow M$. Given some $c \in M$, we must extend f to $A \cup \{c\}$. Namely, because A is already finite, it will be enough to must find some $d \in M$ such that

$$\text{tp}^{\mathcal{M}}(a_1, \dots, a_n, c) = \text{tp}^{\mathcal{M}}(f(a_1), \dots, f(a_n), d),$$

where $A = \{a_1, \dots, a_n\}$. Because \mathcal{M} is atomic (!), we note that $\text{tp}^{\mathcal{M}}(\bar{a}, c)$, we can find some $\theta(\bar{x}, y)$ isolating this type. Now, $\exists y \theta(\bar{x}, y)$ lives in $\text{tp}^{\mathcal{M}}(\bar{a})$ and hence in $\text{tp}^{\mathcal{M}}(f(\bar{a}))$, so we can find some d such that $\mathcal{M} \models \theta(f(\bar{a}), d)$. This d is the one required because θ isolated the type, so $\mathcal{M} \models \theta(f(\bar{a}), d)$ requires that

$$\text{tp}^{\mathcal{M}}(a_1, \dots, a_n, c) = [\theta] \subseteq \text{tp}^{\mathcal{M}}(f(a_1), \dots, f(a_n), d),$$

so the completeness of these types enforces equality. ■

We can even go in the other direction.

Theorem 3.66. Fix a countable language \mathcal{L} , and let \mathcal{M} and \mathcal{N} models of a complete \mathcal{L} -theory T which are countable homogeneous models realizing the same types in $S_n(T)$ for all n . Then $\mathcal{M} \cong \mathcal{N}$.

Proof. These models are countable, so we will do a back-and-forth argument. Enumerate \mathcal{M} by $\{m_i : i \in \mathbb{N}\}$ and \mathcal{N} by $\{n_i : i \in \mathbb{N}\}$. We will build finite partial elementary maps $f_k: M_k \rightarrow N_k$ where $\text{dom } f_k$ contains m_i for $i < k$ and $\text{im } f_k$ contains n_i for $i < k$. At $i = 0$, we simply take the empty function for f_0 .

Now, suppose we are given f_k , and we want to build f_{k+1} . We will discuss how to add m_k to the domain of f_k ; taking the inverse will allow us to add n_k to the image of f_k , so we will omit writing out the argument. Anyway, fully enumerate the domain of f_k by \bar{a} and the image of f_k by \bar{b} . We would like to add in m_k , so we set $P := \text{tp}^{\mathcal{M}}(\bar{a}, m_k)$.

At this point, we would like to use the homogeneity of \mathcal{N} . Well, \mathcal{M} realizes P , so \mathcal{N} must realize P too, so we can find some (\bar{c}, d) realizing P . But then $\text{tp}^{\mathcal{N}}(\bar{c}) = \text{tp}^{\mathcal{M}}(\bar{a}) = \text{tp}^{\mathcal{N}}(\bar{b})$. So we may define a partial elementary embedding by sending $\bar{c} \mapsto \bar{b}$, which by homogeneity extends to a map $(\bar{c}, d) \mapsto (\bar{b}, n')$ for some $n' \in \mathcal{N}$. This n' is the needed element where m_k should go because $\text{tp}^{\mathcal{M}}(\bar{a}, m_k) = \text{tp}^{\mathcal{N}}(\bar{c}, d) = \text{tp}^{\mathcal{N}}(\bar{b}, n')$ by construction. ■

3.6.2 Saturated Models

At the end of class, we are now ready to define saturated models.

Definition 3.67 (κ -saturated). An \mathcal{L} -structure \mathcal{M} is κ -saturated if and only if any $A \subseteq M$ of cardinality less than κ has all types $P \in S_n^{\mathcal{M}}(A)$ realized in \mathcal{M} . Then \mathcal{M} is saturated if and only if it is $|M|$ -saturated.

Example 3.68. Consider the theory DLO of dense linear orders.

- The model \mathbb{Q} is \aleph_0 -saturated. The point is that, by quantifier elimination, any finite set $A \subseteq \mathbb{Q}$ has only the types saying that
- The model \mathbb{Q} is not \aleph_1 -saturated because, for example, there is a 1-type saying that the given element is bigger than every integer, which is not realized. In fact, no countable model \mathcal{M} is \aleph_1 -saturated because one can build a type saying that the given element is not equal to each individual element of \mathcal{M} .

3.7 November 16

Today we continue with saturated models.

3.7.1 Countable Saturated Models

Whenever necessary, we will continue with \mathcal{L} being a countable language and T being a complete \mathcal{L} -theory with infinite models.

Example 3.69. Let's describe a model of DLO which is \aleph_1 -saturated. Let \mathcal{M} consist of functions $f: \omega_1 \rightarrow \mathbb{Q}$ with countable support, ordered lexicographically: namely, $f < g$ if and only if the least i with $f(i) \neq g(i)$ has $f(i) < g(i)$. This is a dense linear order (between any two functions f and g , they differ at some least point i , and so define a function between the two by sending $i \mapsto \frac{1}{2}(f(i) + g(i))$).

To see that \mathcal{M} is \aleph_1 -saturated, one makes some argument using the countable support. Now, we note that it has cardinality 2^{\aleph_0} : there are certainly at least 2^{\aleph_0} functions, and the number of functions can be upper-bounded by the number of countable "graphs" living in $\omega_1 \times \mathbb{Q}$, which simply has size $\aleph_1^{\aleph_0} \leq 2^{\aleph_0 \times \aleph_0} = 2^{\aleph_0}$. So under the continuum hypothesis, \mathcal{M} achieves a cardinality of \aleph_1 , so \mathcal{M} is indeed saturated.

Example 3.70. The algebraically closed field \mathbb{C} is \aleph_1 -saturated. The point is that quantifier elimination means that we only need to worry about types which either say that an element is a root of some polynomial or avoids being a root of some set of polynomials with parameters in a countable set. But \mathbb{C} is big enough that it has enough transcendental elements, so we are okay.

Here are some other ways of thinking about saturated types.

Proposition 3.71. Fix an \mathcal{L} -structure \mathcal{M} and an infinite cardinal κ . Then the following are equivalent.

- \mathcal{M} is κ -saturated.
- For all $A \subseteq M$ of cardinality less than κ , all (possibly partial) types P with parameters in A are realized in \mathcal{M} .
- For all $A \subseteq M$ of cardinality less than κ , any type $P \in S_1^{\mathcal{M}}(A)$ is realized in \mathcal{M} .

Proof. Of course, (b) implies (a) with no content because complete types are types. To show that (a) implies (b), note that a type can be extended to a complete type, so having (a) is enough to realize the needed partial type.

Similarly, we note that (a) implies (c) again with no content. For the converse, suppose we have a complete n -type $P(\bar{x}) \in S_n^{\mathcal{M}}(A)$ which we want to realize. Then we will add in one variable at a time. Namely, start with $P_1(x_1) \in S_1^{\mathcal{M}}(A)$ be the type of formulae in P only mentioning x_1 , and we see that this is a complete 1-type, which is then realized in \mathcal{M} by some a_1 .

Now, because κ is infinite, we may consider the complete type $P_2(a_1, x_2) \in S_1^{\mathcal{M}}(A \cup \{a_1\})$ of formulae in P only mentioning x_1 and x_2 but with x_1 replaced with a_1 . Notably, $P_2(a_1, x_2)$ is indeed consistent with \mathcal{M} : for each $\varphi(x_1, x_2) \in P_2(x_1, x_2)$, we see that $\exists x_2 \varphi(x_1, x_2)$ would need to live in P_1 , so $\mathcal{M} \models \exists x_2 \varphi(a_1, x_2)$, so $\varphi(a_1, x_2)$ is consistent. So we may find a_2 so that (a_1, a_2) realizes P_2 . From here, we can continue the argument inductively up to P_3 and so on. ■

Remark 3.72. We only needed κ to be infinite in order to show that (c) implies (a).

Let's go ahead and start to construct saturated models.

Example 3.73. Consider the theory ACF_0 of algebraically closed fields of characteristic 0. Then an extension F/\mathbb{Q} with countable transcendence degree will be countable and so in particular countable and saturated. The point is that we are only going to consider parameter sets A which are finite, and then F realizes all types to consider, which by quantifier elimination either is asking to be the root of some polynomial in A or to be transcendental with respect to parameters in the finite set A .

Theorem 3.74. Fix a countable language \mathcal{L} and a complete \mathcal{L} -theory T . Then T has a countable saturated model if and only if $S_n(T)$ is countable.

Proof. The forward direction is easier: if T has a countable saturated model \mathcal{M} , then each distinct complete n -type P must correspond to a distinct n -tuple in \mathcal{M} (note distinct complete types have distinct sentences, so they cannot be satisfied by the same n -tuple!), of which there are only countably many, so there can only be countably many types.

We now show the reverse direction. Suppose we have a countable model $\mathcal{M} \models T$. Then we claim that there is a countable model $\mathcal{M}' \geq \mathcal{M}$ realizing any countable set of types. Indeed, add in (countably many) new constants \bar{c}^P corresponding to each of the countably many types $P(\bar{x})$, and everything is finitely satisfiable by \mathcal{M} because T is complete. So we can find a model \mathcal{M}' satisfying everything including $\text{elDiag } \mathcal{M}$.

Iterating this construction grants an ascending sequence

$$\mathcal{M}_0 \leq \mathcal{M}_1 \leq \mathcal{M}_2 \leq \dots,$$

where \mathcal{M}_{i+1} is constructed from \mathcal{M}_i by asking for all the types in $S_1^{\mathcal{M}_i}(A)$ to be realized for each finite subset $A \subseteq \mathcal{M}_i$. The key point is that $S_1^{\mathcal{M}_i}(A)$ is still countable: for each 1-type $P(x, \bar{a}) \in S_1^{\mathcal{M}_i}(A)$ (where \bar{a} is an enumeration of A), we can turn this into an $(n+1)$ -type $P(x, \bar{y}) \in S_{n+1}(T)$, of which there are only countably many. This function is countable-to-one, so we conclude that there are only countably many types in $S_1^{\mathcal{M}_i}(A)$.

To complete the proof, we let \mathcal{M} be the union of the \mathcal{M}_i s, which is an elementary extension of all the \mathcal{M}_i , so \mathcal{M} will realize all the types from all the \mathcal{M}_i . Now \mathcal{M} is countable, so we still only need to worry about finite parameter sets $A \subseteq \mathcal{M}$, which means that we might as well put the parameter set inside some fixed level \mathcal{M}_i , so we know the types will be realized in \mathcal{M}_{i+1} . ■

Remark 3.75. In fact, the above proof has shown that any countable model \mathcal{M}_0 can be embedded (in an elementary way) into a saturated countable model (provided a single saturated countable model exists).

3.7.2 Universal Models

The above proof motivates the following definition.

Definition 3.76 (κ -universal). Fix a complete \mathcal{L} -theory T . A model \mathcal{M} of T is κ -universal if and only if any model \mathcal{N} of T of cardinality less than κ has some elementary embedding $\mathcal{N} \hookrightarrow \mathcal{M}$. Then we say \mathcal{M} is *universal* if only if it is $|M|^+$ universal.

Example 3.77. Consider the theory T of infinitely many equivalence classes with infinite size. Then any model \mathcal{M} with \aleph_1 many classes of size \aleph_1 will be universal by embedding in the necessary number of equivalence classes. Notably, this example informs us that universal models are not isomorphic despite having some kind of universal property.

So in the above language, we now know that countable and saturated implies universal. More generally, we have the following.

Theorem 3.78. Fix a complete \mathcal{L} -theory T and a model \mathcal{M} of T . Then the following are equivalent.

- (a) \mathcal{M} is κ -saturated.
- (b) \mathcal{M} is κ^+ -universal and κ -homogeneous.

We will show this next class.

3.8 November 21

Today we will finish our discussion of saturated models.

3.8.1 Miscellaneous Saturated Models

We begin with the following lemmas.

Lemma 3.79. Fix a complete \mathcal{L} -theory T . If \mathcal{M} is κ -saturated, then \mathcal{M} is κ -homogeneous.

Proof. Fix some subset $A \subseteq M$ of cardinality less than κ and a partial elementary embedding $f: A \rightarrow M$ which we would like to extend to an element $a \in M$.

To do this, set $p(x) := \text{tp}_A^{\mathcal{M}}(a)$. We now push the parameters to B , defining $q(x)$ to be the set of formulae defined by taking any constants from A in any formula $\varphi(x) \in p(x)$ and pushing them to B . We would like to show that $q(x)$ is actually a type (with parameters in B), and \mathcal{M} being κ -saturated implies that it will be realized by some $b \in M$, which will complete the argument by extending f to send $a \mapsto b$.

So it remains to show that $q(x)$ is consistent with $\text{Th}_B \mathcal{M}$. Well, it suffices to show that $q(x) \cup \text{Th}_B \mathcal{M}$ is satisfiable, for which it is enough to show that it is finitely satisfiable, so it is enough to show that any single formula $\varphi(x, \bar{b}) \in q(x)$ is consistent with $\text{Th}_B \mathcal{M}$. Well, this is $\varphi(x, f(\bar{a}))$ for some $\bar{a} \in A$, and f being a partial elementary embedding then implies that $\varphi(x, f(\bar{a})) \in \text{Th}_B \mathcal{M}$ is equivalent to $\varphi(x, \bar{a}) \in \text{Th}_A \mathcal{M}$, which is true by construction of $\varphi(x, \bar{b}) \in q(x)$. ■

Lemma 3.80. If \mathcal{M} and \mathcal{N} are saturated models of the same cardinality, then $\mathcal{M} \cong \mathcal{N}$.

Proof. This is a back-and-forth argument proceeding via transfinite induction. Let κ be the cardinality of M and N , and enumerate $M = \{m_\alpha : \alpha \in \kappa\}$ and $N = \{n_\alpha : \alpha \in \kappa\}$. Then one proceeds as in Proposition 3.57. ■

Remark 3.81. The above proof actually shows that any partial elementary embedding $f_0: A \rightarrow \mathcal{N}$ with $|A| < \kappa$, then we can extend f_0 to an isomorphism. Indeed, just replace the 0th step of the transfinite induction with f_0 .

While we're here, we also show Theorem 3.78.

Theorem 3.78. Fix a complete \mathcal{L} -theory T and a model \mathcal{M} of T . Then the following are equivalent.

- (a) \mathcal{M} is κ -saturated.
- (b) \mathcal{M} is κ^+ -universal and κ -homogeneous.

Proof. We already show that κ -saturated implies κ -homogeneous. To show that κ -saturated implies κ^+ -universal, argue as in Lemma 3.80 but only use embeddings in one direction (i.e., do the forth but not the back).

We now show the other direction; choose \mathcal{M} as in (b). We would like to show that \mathcal{M} realizes any type $P(x) \in S_1^{\mathcal{M}}(A)$ for some $A \subseteq M$ of cardinality less than κ . By Proposition 3.12, we know that $p(x)$ will be realized in some elementary extension \mathcal{N} of \mathcal{M} , and Theorem 1.71 allows us to reduce the size of \mathcal{N} back down to κ , so κ^+ -universality grants an elementary embedding $\varphi: \mathcal{N} \hookrightarrow \mathcal{M}$. However, this elementary embedding need not fix A , so we use the homogeneity of \mathcal{M} to move $\varphi(A^{\mathcal{N}})$ back to $A \subseteq M$, and extending this up to automorphism moves the element realizing $p(x)$ in \mathcal{N} to the needed element of \mathcal{M} . ■

3.8.2 Construction of Saturated Models

We now have the following step to construct saturated models.

Lemma 3.82. Fix an infinite cardinal κ and a language \mathcal{L} with $|\mathcal{L}| < \kappa$ and a complete \mathcal{L} -theory T . For any model \mathcal{M} , there is a κ^+ -saturated model \mathcal{N} which is an elementary extension of \mathcal{M} and with $|\mathcal{N}| \leq |M|^\kappa$.

Proof. As an intermediate step, we claim that there is an elementary extension $\mathcal{M}' \geq \mathcal{M}$ which realizes any type $p(x) \in S_1^{\mathcal{M}}(A)$ where $A \subseteq M$ is a subset of cardinality $|A| \leq \kappa$. We will do this by compactness: let T_0 be the theory of $\text{elDiag } \mathcal{M}$ plus the sentences in all the types $p(c_p)$ we need to satisfy, where c_p is a new constant we added. We want to show that T_0 is satisfiable, for which we note that it is enough to check that only finitely many sentences in finally many of these types is satisfiable with $\text{elDiag } \mathcal{M}$, for which we use Proposition 3.12.

We take a moment to recognize that there are $|M|^\kappa$ subsets A available, and for each subset A , there are $2^{|A|+|\mathcal{L}|+\aleph_0} = 2^\kappa$ possible subsets of sentences (and hence possible types), so the number of added constants in the above construction was at most

$$|M|^\kappa \cdot 2^\kappa = |M|^\kappa.$$

So our \mathcal{M}' may have size at most $|M|^\kappa$.

We now define an elementary sequence using the above steps

$$\underbrace{\mathcal{M}_0}_{\mathcal{M}} \leq \mathcal{M}_1 \leq \mathcal{M}_2 \leq \cdots \leq \mathcal{M}_\alpha \leq \cdots$$

for each $\alpha \in \kappa^+$, where at successor stages we use the above claim, and at limit stages we take unions (which remains an elementary extension due to the chain). Then we define \mathcal{N}_{κ^+} as the union of all these chains.

We take a moment to compute the size of these models. Looking at our stages, we claim that $|N_\alpha| \leq |M|^\kappa$ by induction: at limit stages, we are taking the union of at most κ -many sets of cardinality at most $|M|^\kappa$, which is okay; at successor steps, we note that

$$|N_{\alpha+1}|^\kappa \leq |N_\alpha|^\kappa \leq |M|^{\kappa \times \kappa} = |M|^\kappa$$

by the induction.

It remains to check that \mathcal{N}_{κ^+} is κ^+ -saturated. Well, for any subset $A \subseteq \mathcal{N}_{\kappa^+}$ of size at most κ , there is some $\alpha < \kappa^+$ such that $A \subseteq \mathcal{N}_\alpha$,¹ but then $\mathcal{N}_{\alpha+1}$ realizes any type with parameters in A . ■

¹ Namely, finding a bijection of A to a subset $\lambda \subseteq \kappa^+$, of cardinality at most κ , we note that the supremum of λ is an ordinal which is the union of all elements in A , which has cardinality at most $\kappa \times \kappa = \kappa$, so we must have $\lambda < \kappa^+$.

Theorem 3.83. Fix an infinite cardinal κ and a language \mathcal{L} with $|\mathcal{L}| < \kappa$ and a complete \mathcal{L} -theory T .

- If $\kappa^+ = 2^\kappa$, then there is a saturated model of cardinality κ^+ .
- If λ is infinite, and each $\tau < \lambda$ has $2^\tau \leq \lambda$, then there is a saturated model of cardinality λ .

Proof. Note $2^\kappa = \kappa^\kappa$, so one can iterate the construction of Lemma 3.82 to produce the needed model. Namely, in the first case, there is nothing to do because being κ^+ -saturated implies just being saturated by the size condition. For the second case, one builds a sequence $\{\mathcal{M}_\alpha : \alpha \leq \lambda\}$ where \mathcal{M}_α is \aleph_α^+ -saturated but $|\mathcal{M}_\alpha| \leq \lambda$ and then take the union. ■

Remark 3.84. Some cardinal arithmetic shows that cardinals satisfying the second case exist. For example, one can find λ with $\aleph_\lambda = \lambda$ (take the limit of $\aleph_0, \aleph_{\aleph_0}, \aleph_{\aleph_{\aleph_0}}, \dots$), where the result is true.

3.9 November 28

The final exam will be released Wednesday the 13th of December morning (at 5AM) and due on Friday the 15th of December at 10PM. It will be released and submitted by email. It is expected to be the length of a regular exam, open-note.

3.9.1 Primer on Indiscernibles

We are going to do a little from chapter 5, which is on indiscernibles. Let T be an \mathcal{L} -theory, where \mathcal{L} is a countable language. Recall that if T has an infinite model, compactness implies that T has an infinite model of cardinality κ for each infinite cardinal κ . However, if we require that T omit some type, then we may no longer have arbitrarily large models.

Example 3.85. Take the language \mathcal{L} consisting of countably many constants c_i for each $i \in \mathbb{N}$, where T is the theory that $c_i \neq c_j$ whenever $i \neq j$. Then omitting the type $p(x)$ given by $x \neq c_i$ for each $i \in \mathbb{N}$ requires that any model of T is now countable!

Example 3.86. Take the language \mathcal{L} consisting of countably many constants c_i for each $i \in \mathbb{N}$, and we add in a unary relation ω and a binary relation E . The point is that we will require models to look like they have some ω and also subsets of ω . With this in mind, we require our theory T to have the following.

- Constants are distinct: $c_i \neq c_j$ for each $i \neq j$.
- Constants are in ω : $\omega(c_i)$ for each i .
- Subsets are outside ω : for all x and y with xEy , we have $\omega(x)$ and $\neg\omega(y)$.
- Subsets satisfy extensionality: given x and y with $\neg\omega(x) \wedge \neg\omega(y)$, if $wEx \leftrightarrow wEy$ for each w , then $x = y$.

We now ask to omit the type $p(x)$ given by requiring $\omega(x)$ and $x \neq c_i$. This means that any model \mathcal{M} of T omitting $p(x)$ has that $\mathcal{M} \models \omega(a)$ implies $a = c_i^{\mathcal{M}}$ for any $a \in M$. Now every element outside ω in M can be read off as living in $\mathcal{P}(M) \subseteq \mathcal{P}(\omega)$.

One can also iterate the above example, adding in a relation for subsets of subsets of ω and so on. Namely, there is a language \mathcal{L} and theory T such that omitting a given type $p(x)$ has only models of size at most \beth_α for $\alpha < \omega$. With some effort, we can extend this to any countable α . However, we will show that if T has a model of size \beth_{ω_1} omitting a given type $p(x)$, then T has a model of any larger size.

As another motivation, take T to be a theory in a countable language. We would like models with “lots” of automorphisms.

Example 3.87. Consider $T = \text{Th}(\mathbb{N}, 0, 1, +, \cdot, <)$. There are only countably many types, so find a model \mathcal{M} of T of cardinality \aleph_1 ; then we can surely find two elements of the same type, so the two elements have an automorphism between them in an elementary extension by Proposition 3.16.

We might want to enforce having countable models, which the above example does not require.

3.9.2 Introducing Indiscernibles

Anyway, let’s provide our definition.

Definition 3.88. Fix a countable language \mathcal{L} and \mathcal{L} -theory T . Further, fix a model \mathcal{M} and a linear ordering (I, \leq) . Then $\{x_i\}_{i \in I} \subseteq \mathcal{M}$ is a *sequence of indiscernibles* if and only if the x_i are distinct, and any ordered sequences $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ will have

$$\mathcal{M} \models (\varphi(x_{i_1}, \dots, x_{i_n}) \leftrightarrow \varphi(x_{j_1}, \dots, x_{j_n}))$$

for any \mathcal{L} -formula φ with n free variables.

Example 3.89. For any dense linear order, any subset will produce a sequence of indiscernibles, using the ordering provided by the linear ordering. Indeed, one can write down an automorphism sending any finite subset to another finite subset.

Example 3.90. Give $\mathbb{Z} \times \mathbb{Z}$ the lexicographical ordering. Then $\{(x, 0) : x \in \mathbb{Z}\}$ is a sequence of indiscernibles: there is no first-order way to tell these elements apart because they are already “infinitely” apart.

Remark 3.91. Given a set of indiscernibles $\{x_i : i \in I\}$ as above, one can define its type as

$$\text{tp}(I) := \{\varphi(v_1, \dots, v_n) : \mathcal{M} \models \varphi(x_{i_1}, \dots, x_{i_n}) \text{ for any } i_1 < \dots < i_n\}.$$

Note the choice of x_{i_1}, \dots, x_{i_n} does not matter by definition of the indiscernibles.

Let’s try to find indiscernibles.

Theorem 3.92. Let T be an \mathcal{L} -theory with infinite models. Fix a linear ordering (I, \leq) . Then there is a model $\mathcal{M} \models T$ with a sequence $\{x_i : i \in I\}$ of indiscernibles.

This result requires Ramsey’s theorem. For notation, let $[X]^n$ denote the sets of X with n elements.

Theorem 3.93 (Ramsey). Any k -coloring $c: [\mathbb{N}]^n \rightarrow \{0, 1, \dots, k\}$ has an infinite monochromatic set $H \subseteq \mathbb{N}$. Namely, there is ℓ such that any $S \in [H]^n$ will have $c(S) = \ell$.

For example, $n = 0$ has nothing to show, and $n = 1$ is the Pigeonhole principle. If we have time, we will prove Theorem 3.93 next class. Anyway, here is our proof of Theorem 3.92.

Proof of Theorem 3.92 using Theorem 3.93. We use compactness. Add to our language new constants c_i for each $i \in I$, and let T' be the theory T adding in the requirements that $c_i \neq c_j$ for each $i \neq j$ and also the sequences

$$\varphi(c_{i_1}, \dots, c_{i_n}) \leftrightarrow \varphi(c_{j_1}, \dots, c_{j_n})$$

for each pair of ordered sequences $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ in I . We need to show that T' is satisfiable, so it is enough by compactness to show that T' is finitely satisfiable. Well, any finite subset $\Delta \subseteq T'$ will only mention finitely many formulae $\varphi_1, \dots, \varphi_k$ and finitely many constants c_1, \dots, c_n .

We now show that Δ is satisfiable. Well, pick up some model \mathcal{M} , and we may as well assume that \mathcal{M} is countable.² We would now like to assign the finitely many constants c_\bullet to distinct places to satisfy Δ . For this, we use Theorem 3.93. Namely, order M as $M = \{a_i : i \in \mathbb{N}\}$, and we define the coloring

$$c: [M]^n \rightarrow 2^k,$$

where $c(\{a_1, \dots, a_n\})$ is the subset of $\{1, 2, \dots, k\}$ determined by having ℓ if and only if $\mathcal{M} \models \varphi_\ell(a_1, a_2, \dots)$ and not having ℓ otherwise. (The arity of φ_ℓ matters here, which is why we have written a_1, a_2, \dots .) Now, Theorem 3.93 produces an infinite subset $H \subseteq M$ such that c is constant on $[H]^n$. Thus, we may send the constants c_1, \dots, c_n wherever we please in H (ordered properly). ■

Remark 3.94. The above proof in fact promises a model \mathcal{M} of cardinality $|\mathcal{L}| + |I| + \aleph_0$.

Let's think a little about using indiscernibles to produce automorphisms. This requires an adjective to our theory.

Definition 3.95 (Skolem functions). Fix an \mathcal{L} -theory T . Then T has *built-in Skolem functions* if and only if any \mathcal{L} -formula $\varphi(\bar{x}, \bar{y})$ has a term $t(\bar{x})$ such that

$$T \models \forall \bar{x} (\exists \bar{y} \varphi(\bar{x}, \bar{y}) \rightarrow \varphi(\bar{x}, t(\bar{x}))).$$

We call $t(\bar{x})$ a *Skolem function*.

Remark 3.96. A theory T with built-in Skolem functions is model-complete: for any substructures $\mathcal{M} \subseteq \mathcal{N}$, we can use the Tarski–Vaught test to show $\mathcal{M} \leq \mathcal{N}$. Namely, looking at Lemma 1.72, choosing some formula $\varphi(\bar{x}, y)$ and $\bar{a} \in M^n$ with $\mathcal{N} \models \exists y \varphi(\bar{a}, y)$, then the existential is witnessed by a Skolem function, meaning that we can find it in \mathcal{M} .

Example 3.97. Peano arithmetic has “definable” Skolem functions by finding the least element satisfying some formula. Adding in function symbols produces an extension of the theory with built-in Skolem functions.

3.10 November 30

It's the last class for this course.

3.10.1 Skolem Functions

We note that we can always add in Skolem functions to a theory.

Proposition 3.98 (Skolemization). Fix a language \mathcal{L} and theory T . Then we can extend the language and theory to \mathcal{L}^* and T^* , respectively, so that T^* has built-in Skolem functions.

² Even if \mathcal{L} is countable, Δ only uses a finite subset of this language, so we may restrict T to this language when concerned with the satisfiability of Δ .

Proof. Set $\mathcal{L}_0 := \mathcal{L}$ and $T_0 := T$. Then for each i , we can make \mathcal{L}_{i+1} from \mathcal{L}_i by adding a function symbol f_φ for each formula $\varphi(\bar{x}, y)$ and then make T_{i+1} from T_i by adding the sentences

$$\forall \bar{x} (\exists y \varphi(\bar{x}, y) \rightarrow \varphi(\bar{x}, f(\bar{x}))).$$

Taking the union over all the \mathcal{L}_i and T_i will produce the needed \mathcal{L}^* and T^* . ■

Remark 3.99. A theory T with Skolem functions has quantifier elimination: any formula $\varphi(\bar{x}, y)$ with Skolem function f_φ has

$$T \models \forall \bar{x} (\exists y \varphi(\bar{x}, y) \leftrightarrow \varphi(\bar{x}, f_\varphi(\bar{x}))),$$

effectively eliminating the outer existential from $\exists y \varphi(\bar{x}, y)$. So we can eliminate quantifiers, by hand, one at a time. For example, it follows that T is model-complete, recovering Remark 3.96.

Anyway, we are permitted the following definitions.

Definition 3.100 (Skolemization). Fix an \mathcal{L} -theory T . Then a *Skolemization* of T is an extension T^* of T to a new language $\mathcal{L}^* \supseteq \mathcal{L}$ which has built-in Skolem functions.

To continue, we pick up the following definition.

Definition 3.101. Fix an \mathcal{L} -theory T with built-in Skolem functions. Given a subset X of a model $\mathcal{M} \models T$, we let $H(X)$ be the \mathcal{L} -substructure generated by X .

Remark 3.102. Explicitly, the universe of $H(X)$ is the collection of terms whose inputs come from X .

Remark 3.103. By Remark 3.96, we note that $H(X) \leq \mathcal{M}$ in the above situation.

3.10.2 Building Automorphisms

We now get the following theorem on automorphisms of indiscernibles.

Theorem 3.104. Fix an \mathcal{L} -theory T with built-in Skolem functions. Further, fix a model \mathcal{M} and a sequence of indiscernibles $\{x_i\}_{i \in I}$, with $X := \{x_i : i \in I\}$. Given any order-preserving automorphism $\sigma : I \rightarrow I$, there is a unique automorphism $\tau : H(X) \rightarrow H(X)$ such that $\tau(x_i) = x_{\sigma(i)}$ for each $i \in I$.

Proof. The uniqueness is easier, and only needs us to preserve the relevant structure (namely, not bijectivity): any element of $H(X)$ is of the form $t(\bar{x})$ where t is a term and $\bar{x} \in X$, but then any automorphism $\tau : H(X) \rightarrow H(X)$ with $\tau(x_i) = x_{\sigma(i)}$ must have

$$\tau(t(\bar{x})) = t(\tau(\bar{x})),$$

thus defining τ .

For existence, we need to show that the τ defined above is well-defined and an automorphism. Here are our checks.

- Suppose we have equal terms $t(x_{i_1}, \dots, x_{i_k}) = s(x_{j_1}, \dots, x_{j_\ell})$, and we need to show that

$$t(x_{\sigma i_1}, \dots, x_{\sigma i_k}) \stackrel{?}{=} s(x_{\sigma j_1}, \dots, x_{\sigma j_\ell}).$$

Well, this follows from X being a sequence of indiscernibles: we see that the sentence

$$t(y_{i_1}, \dots, y_{i_k}) = s(y_{j_1}, \dots, y_{j_\ell})$$

is going to be in the type of $\{x_{i_1}, \dots, x_{i_k}, x_{j_1}, \dots, x_{j_\ell}\}$ (suitably ordered and so on), so because applying σ (which is order-preserving!) cannot change the type, we see that the sentence still holds upon applying σ , achieving the needed equality.

- We show that τ preserves the structure. Essentially the same argument goes through. For example, $\mathcal{M} \models R(t_1(\bar{x}), \dots, t_n(\bar{x}))$ is equivalent to $\mathcal{M} \models R(t_1(\sigma\bar{x}), \dots, t_n(\sigma\bar{x}))$ because we are asking for $R(t_1(\bar{x}), \dots, t_n(\bar{x}))$ to be in the type of \bar{x} , which does not change because X is a sequence of indiscernibles, and σ is order-preserving.
- We show that τ is an isomorphism. Well, we can define some automorphism τ' which sends x_i to $x_{\sigma^{-1}i}$ using the argument above, but then τ and τ' are inverse maps of structures: $\tau \circ \tau' = \text{id}$ and $\tau' \circ \tau = \text{id}$ are morphisms sending $x_i \mapsto x_i$ and so must be the identity by uniqueness, so we do indeed have an automorphism. ■

As a corollary, we get to build models that have lots of automorphisms!

Corollary 3.105. Fix an \mathcal{L} -theory T with \mathcal{L} countable, where T has infinite models. For any $\kappa \geq \aleph_0$, there is a model $\mathcal{M} \models T$ of cardinality κ and 2^κ automorphisms.

Proof. We may as well extend T to be a complete theory which has built-in Skolem functions by taking the Skolemization. (Note that the language remains countable due to the construction in Proposition 3.98.) Now, choose a linear ordering (I, \leq) of cardinality κ , and then Theorem 3.92 promises the existence of a model $\mathcal{M} \models T$ of cardinality κ with a sequence of indiscernibles $\{x_i : i \in I\}$. We now may replace \mathcal{M} with $H(\{x_i : i \in I\})$, which we note still has cardinality I .

Now, Theorem 3.104 promises that $\mathcal{M} = H(\{x_i : i \in I\})$ will have at least as many automorphisms as order-preserving maps $I \rightarrow I$, so it remains to choose I so that I has 2^κ order-preserving maps $I \rightarrow I$. Well, choose $I := \mathbb{Z} \times \kappa$ to be κ many copies of \mathbb{Z} , named \mathbb{Z}_α for $\alpha \in \mathbb{Z}$. Then for any subset $S \subseteq \kappa$, we can build an order-preserving map σ_S by applying a $+1$ shift to \mathbb{Z}_α for each $\alpha \in S$ and do nothing to \mathbb{Z}_α for $\alpha \notin S$. ■

Here is another application.

Lemma 3.106. Fix an \mathcal{L} -theory T , and let $X = \{x_i : i \in I\}$ be a sequence of indiscernibles. For any linear ordering J , there is a model \mathcal{N} with a sequence of indiscernibles $Y = \{y_j : j \in J\}$ such that

$$\text{tp}^{\mathcal{M}}(X) = \text{tp}^{\mathcal{N}}(Y).$$

Further, if T has built-in Skolem functions and \mathcal{M} omits any type $p(x)$, then we may require that \mathcal{N} omits the same type $p(x)$.

Proof. We use compactness. For the first claim, build \mathcal{N} by compactness, as in Theorem 3.92. Namely, add in constants $\{c_j : j \in J\}$ and force \mathcal{N} to satisfy $\text{elDiag } \mathcal{M}$ along with the sentences $\varphi(c_{j_1}, \dots, c_{j_n})$ when $j_1 < \dots < j_n$ and φ lives in $\text{tp}^{\mathcal{M}}(X)$. This is finitely satisfiable by \mathcal{M} because any finite segment of J looks just like a finite segment of I .

To get the second claim, move everything up to a Skolemization first and then replace the given model \mathcal{N} with $H(Y)$. Indeed, the type of any sequence of elements in \mathcal{N} only uses finitely many elements of \mathcal{N} , so we can move the used subsequence of y_\bullet s back to x_\bullet s to show that the types in \mathcal{N} are a subset of the types in \mathcal{M} . ■

Theorem 3.107. Fix a countable language \mathcal{L} and an \mathcal{L} -theory T with built-in Skolem functions. Further, suppose that any $\alpha \in \omega_1$ has a model $\mathcal{M} \models T$ of size at least \beth_α which omits a given type $p(x) \in S_1(T)$. Then there is a model $\mathcal{N} \models T$ still omitting $p(x)$ with an infinite sequence of indiscernibles.

Proof. Intuitively, we will pick up some large subset of \mathcal{M} and slowly make it smaller and smaller in order to avoid the type and all satisfy the same formulae (to be a sequence of indiscernibles). The process of making it smaller but maintaining an infinite size is some kind of coloring problem with every formula colored by which formulae they satisfy. So we will want the following purely combinatorial result.

Theorem 3.108. Fix a set B of size $\beth_n(\kappa)^+$ and a coloring $c: [B]^{n+1} \rightarrow \kappa$. Then there is a monochromatic subset $A \subseteq B$ of cardinality κ^+ .

Here, we recall that \beth_n means repeating \mathcal{P} a total of n times.

Let's provide a few more details. Add in constants $\{c_i : i \in \mathbb{N}\}$ to our language. To begin our compactness argument, we add to our theory $\Sigma \supseteq T$ the requirements that each the c_i are distinct and that we are making the $\{c_i : i \in \mathbb{N}\}$ into a sequence of indiscernibles. This is still satisfiable, so we will try to complete Σ in a way that has this continue to be satisfiable. Now, the main idea is to add in the constraints that each term $t(\bar{v})$ has $\varphi(y) \in p(y)$ such that $\neg\varphi(t(\bar{c}))$.

To do this construction, we will build Σ inductively via a sequence $\Sigma_0 \subseteq \Sigma_1 \subseteq \dots$. At each step of the construction $\alpha \in \omega_1$, we ensure that we have a model \mathcal{N}_α omitting $p(x)$ but satisfying Σ_s , and we ensure that \mathcal{N}_α has a sequence of \beth_α indiscernibles. ■

Remark 3.109. The above theorem answers many of the questions of Section 3.9.1 on trying to construct models with lots of automorphisms but still omitting a type.

BIBLIOGRAPHY

- [Mar02] D. Marker. *Model Theory : An Introduction*. Graduate Texts in Mathematics. Springer New York, 2002. ISBN: 9780387987606. URL: <https://books.google.com/books?id=QieAHk--GCcC>.
- [Jop10] Janis Joplin. *Me and Bobby McGee*. 2010. URL: https://www.youtube.com/watch?v=OTHRg_iSWzM.
- [Doe14] Anthony Doerr. *All the Light We Cannot See*. Scribner, 2014.
- [Shu16] Neal Shusterman. *Scythe*. Arc of a Scythe. Simon & Schuster, 2016.
- [Pia] Beth Piatote. *ku'nu*.

LIST OF DEFINITIONS

- atomic, 63
- atomic formula, 6
- cell, 46
- cell decomposition, 46
- complete, 53
- definable, 8
- dense, 60
- diagram, 21
- elementarily equivalent, 21
- elementary equivalence, 8
- elementary substructure, 24
- eliminates quantifiers, 35
- embedding, 5
- expansion, 19
- filter, 16
- finitely satisfiable, 9
- formula, 6
- free variable, 6
- homomorphism, 5
- isolated, 58, 60, 60
- isomorphism, 5
- κ -categorical, 28
- κ -homogeneous, 65
- κ -saturated, 66
- κ -universal, 68
- language, 5
- o -minimal, 45
- partial elementary map, 54
- prime, 63
- quantifier-free type, 33
- realizes, 53
- satisfiable, 9
- sentence, 6
- Skolem functions, 73
- Skolemization, 74
- Stone topology, 56
- structure, 5
- term, 6
- theory, 7, 8, 21
- truth, 7
- type, 34, 51, 52, 60
- ultrafilter, 17
- ultraproduct, 19
- unnested, 36
- witness, 11