Model theory notes.

Kevin Buzzard

April 26, 2012

1 Introduction

Ambrus Pal is (Jan-Mar 2008) giving some lectures on model theory. I don't know the basic definitions and facts though! So here are some notes containing the basic facts and possibly more.

2 Definitions.

A *language* $\mathcal{L}$ is:

(i) a set $\mathcal{F}$ of "function symbols", and a positive integer n_f for each $f \in \mathcal{F}$

(ii) a set $\mathcal{R}$ of "relation symnols", and a positive integer n_R for each $R \in \mathcal{R}$, and

(iii) a set $\mathcal{C}$ of "constant symbols".

For example the "language of rings" might be the language with $\mathcal{F} = \{+, -, .\}$ with $n_+ = n_- = n_. = 2$, $\mathcal{R}$ empty, and $\mathcal{C} = \{0, 1\}$. Note that there are other ways to do it! Later on, when we get to theories (that is, axioms) we will see that we can either make minus ($-$) part of the language, as we just did above, or we can bung in an axiom stating that everything has an additive inverse, and hence we can "deduce" the minus function from the axioms. Sometimes innocuous changes like this really make a difference (for example, although I haven't explained any of what this means yet, we'll see later that something called "quantifier elimination" might be false in one theory and true in a completely equivalent theory with "more language and less axioms"). But for sanity's sake we have to choose once and for all what we are talking about, and Ambrus in his course choses the following: the language of rings will have $\mathcal{F} = \{+, -, ., 0, 1\}$, the language of groups will have $\mathcal{F} = \{1, .\}$ (1 the identity constant, . the multiplication function with $n_. = 2$, but no inverse function: we will assert existence of inverses as part of the axioms) and the language of orderings will just be one relation $\mathcal{F} = \{<\}$.

An *$\mathcal{L}$-structure*, typically denoted, $\mathcal{M}$ is the following things:

(i) a non-empty set M, the *underlying set* of the structure,

(ii) a function $f^{\mathcal{M}} : M^{n_f} \to M$ for each $f \in \mathcal{F}$,

(iii) a relation, that is, a subset $R^{\mathcal{M}} \subseteq M^{n_R}$, for each relation-symbol R, and

(iv) an element $c^{\mathcal{M}} \in M$ for each constant-symbol c.

Examples: in the language of groups sometimes people write 1_G for the identity element of G; here we would say that $\mathcal{G}$ was the group, G was the underlying set, and write $1^{\mathcal{G}}$ for the identity element.

Note: if $\mathcal{L}$ is the language of rings then an $\mathcal{L}$-structure is not necessarily a ring! We have no axioms yet. An $\mathcal{L}$-structure is just a set S with two distinguished elements 0 and 1 and three completely random functions $+, -, * : S^2 \to S$ that can be anything.

An *$\mathcal{L}$-embedding* $\eta : \mathcal{M} \to \mathcal{N}$ between two $\mathcal{L}$-structures is an injection $M \to N$ with commutes with the function-symbols and constant-symbols in the obvious way, but there's a "catch": it must also have the property that $(m_1, m_2, \ldots, m_{n_R}) \in R^{\mathcal{M}}$ if *and only if* $(\eta(m_1), \eta(m_2), \ldots, \eta(m_{n_R})) \in R^{\mathcal{N}}$. Initially this surprised me, but perhaps if for every relation you also adjoin its "opposite" then you can see that this makes sense. Something that also surprised me was that the definition only allowed injections—but if you regard "equals" as a relation then the if and only if part of the definition forces the map to be an injection!

A bijective $\mathcal{L}$-embedding is called a $\mathcal{L}$-isomorphism.

If $M \subseteq N$ and η is the inclusion and η happens to be an $\mathcal{L}$-embedding, we say $\mathcal{M}$ is a *substructure* of $\mathcal{N}$ and $\mathcal{N}$ is an *extension* of $\mathcal{M}$.

A *word* in $\mathcal{L}$ is just a finite string of symbols built using the elements of $\mathcal{F} \cup \mathcal{R} \cup \mathcal{C}$, countably infinitely many variable symbols $v_1, v_2, \ldots$, the symbols $=$, $\vee$, $\wedge$, $\neg$, $\forall$, $\exists$, (and), and a comma. A *term* is, vaguely, a word that can be interpreted as an element of an $\mathcal{L}$-structure, possibly after substitution of variables. More formally, the set of *terms* of a language is the smallest set containing all the constant-symbols, all the variables, and such that if $t_1, t_2, \ldots, t_{n_f}$ are terms then $f(t_1, t_2, \ldots, t_{n_f})$ is a term for all function-symbols f. If a term t mentions only the variables $v_1, v_2, \ldots, v_d$ then given a structure $\mathcal{M}$ and elements $m_1, m_2, \ldots, m_d$ in that structure, there's an obvious way of evaluating the term at these elements (formally one makes a recursive definition, recursing on length of term); we call the resulting function $t^{\mathcal{M}}$.

An *atomic formula* is a word which is either of the form

(i) $t_1 = t_2$, with t_1 and t_2 terms, or

(ii) $R(t_1, t_2, \ldots, t_{n_R})$ for R a relation-symbols and the t_i terms.

The set of *formulae* for a structure $\mathcal{L}$ is the smallest set of words containing the atomic formulae and such that

(i) if ϕ is a formula then so is $\neg\phi$

(ii) If ϕ and ψ are formulae then so is $\phi \wedge \psi$ and $\phi \vee \psi$

(iii) If ϕ is a formula then so is $\forall v_i \phi$ and $\exists v_i \phi$.

The set of formulae can be built recursively, and hence one can do "induction on length of formula".

A variable v occurs *freely* in a formula if it's not inside an $\exists v$ or $\forall v$ quantifier; otherwise it's *bound*. A crappy example of a formula in the language of rings is $v = 0 \vee \exists v\ v.v = 1$ and this is a bit rubbish because v occurs both as a free and a bound variable; Marker's book restricts to formulae in which no variable occurs in both a free and a bound way in any subformula.

Note that we're quantifying over elements of the structure, not subsets. This is first order logic!

A *sentence* is a formula with no free variables.

If ϕ is an $\mathcal{L}$-formula with free variables in the set $(v_1, v_2, \ldots, v_d)$, if $\mathcal{M}$ is an $\mathcal{L}$-structure, and if $a_1, a_2, \ldots, a_d \in \mathcal{M}$ then by induction on length of formula one can give a truth-value to $\phi(a_1, a_2, \ldots, a_d)$. If $\phi(\vec{a})$ turns out to be true in $\mathcal{M}$ we write $\mathcal{M} \models \phi(\vec{a})$, pronounced "$\phi(\vec{a})$ is true in $\mathcal{M}$" or "$\mathcal{M}$ satisfies $\phi(a)$". If ϕ is a sentence then we just write $\mathcal{M} \models \phi$. Note that for every sentence and for every $\mathcal{L}$-structure $\mathcal{M}$, either $\mathcal{M} \models \phi$ or $\mathcal{M} \models \neg\phi$.

A *theory*, or an $\mathcal{L}$*-theory*, is just a set of sentences in the language $\mathcal{L}$. We say that an $\mathcal{L}$-structure $\mathcal{M}$ is a *model* for the theory T if $\mathcal{M} \models \phi$ for all $\phi \in T$, and we write $\mathcal{M} \models T$.

Finally we can do mathematics: for example ring theory is the set of all structure for the language of rings which are models for the theory which consists of the axioms for a ring. And so on.

The *full theory* of an $\mathcal{L}$-structure $\mathcal{M}$ is the set of all $\mathcal{L}$-sentences ϕ such that $\mathcal{M} \models \phi$. This set is called $\mathrm{Th}(\mathcal{M})$ and note that for every ϕ, precisely one of ϕ and $\neg\phi$ will be in $\mathrm{Th}(\mathcal{M})$.

Two $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$ are *elementarily equivalent*, denoted $\mathcal{M} \equiv \mathcal{N}$, if they have the same theory, that is, if the $\mathcal{L}$-sentences that are true for $\mathcal{M}$ coincide with the $\mathcal{L}$-sentences that are true for $\mathcal{N}$. This is a *much* much weaker notion than that of isomorphism. For example I suspect that in the theory of fields, $\overline{\mathbf{Q}}$ and $\mathbf{C}$ will be elementarily equivalent. On the other hand an easy formula induction shows that isomorphic $\mathcal{L}$-structures are elementarily equivalent.

Examples of theories: if $\mathcal{L}$ is the language of orderings then examples of theories are complete orderings (I call these "total orderings"), dense orderings without endpoints, discrete orderings with a bottom but no top, and so on. If $\mathcal{L}$ is the language of groups then examples of theories are the theory of abelian groups, or the theory *DAG* of non-zero torsion-free divisible abelian groups (note that the natural thing to do here is to use infinitely many axioms: $\forall x(x+x+x+x+\cdots+x) = 0 \to x = 0$ with n terms in the sum, for $n = 1, 2, 3, \ldots$. Note that there doesn't seem to be a "theory of torsion abelian groups"![1] Finally, if $\mathcal{L}$ is the language of rings, then examples of theories are the theory of rings, of commutative rings, of integral domains, of fields, the theory ACF of algebraically closed fields (again you need infinitely many axioms, one for each "degree", the nth one saying that any polynomial of degree n has a root), the theory ACF_0 of algebraically closed fields of characteristic 0, or the theory ACF_p of algebraically closed fields of characteristic p (with p a prime). It's an exercise to check that all these can be axiomatised. Note however that if $\mathcal{L}$ is the empty language and one attempts to develop the theory of topological spaces, one is in trouble, because you can't write out the axioms for a topological space using these (first-order) formulae that we have just defined: you need to quantify over subsets, which we're not allowed to do. I don't know whether this observation is "profound" or not though; I'm not really sure what to make of it, to be honest!

3 Filters.

A filter on a set I is a subset F of the power set of I with the empty set not in F, and if $s, t \in F$ then $s \cap t \in F$, and if s is in F and $s \subseteq u$ then $u \in F$. An ultrafilter $\mathcal{U}$ is a maximal filter; this has the property that for any subset J of I, exactly one of J and $I \backslash J$ is in $\mathcal{U}$.

4 The compactness theorem.

If $\mathcal{L}$ is a language and T is an $\mathcal{L}$-theory, say that T is *satisfiable* if there's an $\mathcal{L}$-structure $\mathcal{M}$ such that $\mathcal{M} \models T$.

Theorem 1 (Compactness theorem.). *T is satisfiable iff all finite subsets of T are satisfiable.*

Before I start the proof I need to explain a fundamental construction. If I is an index set and $\{X_i : i \in I\}$ is a collection of sets indexed by I, and $\mathcal{U}$ is an ultrafilter on I, then define an equivalence relation on $\prod_i X_i$ by $(a_i)\tilde{(}b_i)$ iff $\{i : a_i = b_i\} \in \mathcal{U}$. This is an equivalence relation (easy) and the set of equivalence classes is denoted $(\prod_i X_i)/\mathcal{U}$.

The fundamental construction in the proof of the compactness theorem is the following. If $\mathcal{L}$ is a language and $\{\mathcal{M}_i : i \in I\}$ is a collection of $\mathcal{L}$-structures, and $\mathcal{U}$ is an ultrafilter on I, then there's an ultraproduct $\prod_i \mathcal{M}_i/\mathcal{U}$ whose underlying set is $\prod_i M_i/\mathcal{U}$, and where the functions, relations and constant-symbols are defined in the obvious way: for example, if $n = n_f$ and $(a^1, a^2, \ldots, a^n) \in \prod_i M_i/\mathcal{U}$ then choose representatives $(a_i^r)_{i \in I}$ for a^r, $1 \leq r \leq n$, and define $f^{\prod_i \mathcal{M}_i/\mathcal{U}}(a^1, a^2, \ldots, a^n)$ to be $(f^{\mathcal{M}_i}(a_i^1, a_i^2, \ldots, a_i^n))_{i \in I}$, and check that this is well-defined. Relations: say the relation is true on $(a^1, \ldots, a^n)$ iff the set of i for which it's true on the ith chunk, is in $\mathcal{U}$.

Let $\mathcal{M}$ be $\prod_i \mathcal{M}_i/\mathcal{U}$. Say ϕ is a formula with n free variables. Say $a^1 = (a_i^1)_{i \in I}$ and so on are elements of $\prod_i M_i$. The basic fact is:

Lemma 2 (Los Lemma). *$\mathcal{M} \models \phi(a^1, a^2, \ldots, a^n)$ iff $\{i : i \in I, \mathcal{M}_i \models \phi(a_i^1, a_i^2, \ldots, a_i^n)\} \in \mathcal{U}$.*

The proof of the Los Lemma is easy formula induction. Usually one only uses the Los Lemma for sentences but formula induction doesn't work for sentences! The only time one uses the ultra-ness of the filter is when proving the induction step for $\neg$.

The *proof of the compactness theorem* now goes as follows: Say all finite subsets have a model. Let I denote the set of finite subsets of T. For each $i \in I$ let $\mathcal{M}_i$ denote a model for the elements of i. If $\phi \in T$ let $\hat{\phi}$ denote the subset of I consisting of i with $\phi \in i$. Let $\mathcal{F}$ denote the subset of the power set of I consisting of the $\hat{\phi}$ for $\phi \in T$. One checks without too much trouble that $\mathcal{F}$ extends to an ultrafilter $\mathcal{U}$ on I. Now $\prod_i \mathcal{M}_i/\mathcal{U}$ is a model for all of T by the Los Lemma.

[1] Indeed, jumping ahrad a little, if one considers the ultraproduct of the finite cyclic groups C_1, C_2, C_3,..., then this will satisfy the statement "$n.1 \neq 0$" for every n, so the property of being torsion isn't preserved under ultraproducts and hence isn't axiomatisable in this way.

5 Applications of the compactness theorem.

By a *graph* I mean a set V and a symmetric binary relation E on V. The theorem (apparently due to Erdos and de Bruijn) is that a graph has a k-colouring (k some positive integer) iff all finite subsets have a k-colouring. The proof goes like this. Fix $k \geq 1$. let $\mathcal{L}$ denote the language with no functions, a constant c_v for each element v of V, and $k+1$ relations: one binary one called E and k unary ones called R_0, R_1,…,R_{k-1}. Let T be the following set of axioms: for each $v, w \in V$ with vEw in the graph, bung in an axiom c_vEc_w. Now bung in an axiom saying that for all x, exactly one of $R_0(x), R_1(x), \ldots, R_{k-1}(x)$ is true, and bung in another one saying that for all x and y, if xEy and $R_i(x)$ is true, then $R_i(y)$ is false. A finite subset of these axioms only mentions finitely many vertices of the graph, so it has a model. So there's a model for all the axioms and this is a colouring; for each v just associate the colour i where i is the unique i such that $R_i(c_v)$.

A shocking application is:

Theorem 3. *(super-weak going up) If T is an $\mathcal{L}$-theory and T has an infinite model then T has models of arbitrarily large cardinality.*

Note that one needs the assumption: for example if $\mathcal{L}$ is the empty language then there's an axiom which says "any model satisfying me has 3 elements".

Proof. For a cardinal κ just bung in a constant-symbol for every element of κ, call the resulting language $\mathcal{L}'$, and consider the $\mathcal{L}'$-theory which is the full theory of $\mathcal{M}$ (our infinite model) and the extra axioms $c_\alpha \neq c_\beta$ for all $\alpha \neq \beta \in \kappa$. □

Note: There are much better going up theorems. This one is "super-weak" because (a) given a model it produces a new one but the new one might not be elementarily equivalent to the first one, and (b) we don't get precise bounds on cardinalities. The best going up theorem (later on) is a much stronger statement about T having models which are elementarily equivalent to a given model and which have exactly some cardinality.

6 Elementary embeddings, elementary equivalence, and going up and going down.

Say that an embedding $\eta : \mathcal{M} \to \mathcal{N}$ is an *elementary embedding* if for every formula ϕ with n free variables, and for every $\vec{a} \in M^n$, $\mathcal{M} \models \phi(a)$ iff $\mathcal{N} \models \phi(\eta(a))$. In words, every formula in N that only mentions stuff in M will be true in N if and only if it's true in M. Notation: $\mathcal{M} \prec \mathcal{N}$.

Non-example: $\mathbf{Z}$ injects into $\mathbf{Q}$ in the theory of rings, but it's not an elementary embedding because $\exists x : x + x = 1$ is true in $\mathbf{Q}$ and makes sense within $\mathbf{Z}$ but isn't true in $\mathbf{Z}$. Similarly $\neg(\exists x : x + x = 1)$ is false in $\mathbf{Q}$ and makes sense within $\mathbf{Z}$ but is true in $\mathbf{Z}$.

Of course if $\mathcal{M} \prec \mathcal{N}$ then $\mathcal{M} \equiv \mathcal{N}$; this is just a special case: $\mathcal{M} \equiv \mathcal{N}$ just means that a sentence (that is, a formula with no free variables at all) is true in $\mathcal{M}$ iff it's true in $\mathcal{N}$.

Here's a terrifying example of two structures $\mathcal{M} \subseteq \mathcal{N}$ with $\mathcal{M} \equiv \mathcal{N}$ (indeed $\mathcal{M} \cong \mathcal{N}$!) but $\mathcal{M} \not\prec \mathcal{N}$. Take the language of orders and let $\mathcal{M}$ be $\mathbf{Z}_{\geq 1}$ considered as a subset of $\mathbf{Z}_{\geq 0}$. If $\phi(x)$ is the sentence $(\exists y)(y < x)$ then $\phi(1)$ is true in $\mathbf{Z}_{\geq 0}$ but not in $\mathbf{Z}_{\geq 1}$. Another example: $2\mathbf{Z} \subset \mathbf{Z}$ in the language of groups. This is related to the notion of "model completeness", which is mentioned later on in these notes.

Notation: If $\mathcal{M} \prec \mathcal{N}$ we say $\mathcal{M}$ is an *elementary substructure* of $\mathcal{N}$ and that $\mathcal{N}$ is an *elementary extension* of $\mathcal{M}$.

Here's a neat trick: enlarging our language a bit. Say $\mathcal{M}$ is a $\mathcal{L}$-structure; let $\mathcal{L}_\mathcal{M}$ denote the language $\mathcal{L}$ plus an extra constant c_m for each $m \in M$. Of course, $\mathcal{M}$ is, in a natural way, a $\mathcal{L}_\mathcal{M}$-structure! Let $\mathrm{Diag}(\mathcal{M})$ denote the set of all sentences ϕ of $\mathcal{L}_\mathcal{M}$ such that $\mathcal{M} \models \phi$ as a $\mathcal{L}_\mathcal{M}$-structure. This set $\mathrm{Diag}(\mathcal{M})$ is called the *elementary diagram* of $\mathcal{M}$, or just the *diagram* of $\mathcal{M}$. I suspect that the point of this notion is:

Lemma 4. *Let $\mathcal{N}$ be an $\mathcal{L}_\mathcal{M}$-structure such that $\mathcal{N} \models \mathrm{Diag}(\mathcal{M})$. Then there's an elementary $\mathcal{L}$-embedding $\mathcal{M} \to \mathcal{N}$.*

Proof. The obvious thing works: send $m \in M$ to $(c_m)^\mathcal{N}$. Everything is either immediate or almost immediate. An example of an almost-immediate thing: the map is injective because if $m, r \in M$ are distinct then $\mathcal{M} \models_{\mathcal{L}_\mathcal{M}} c_m \neq c_r$ and hence $c_m \neq c_r \in \mathrm{Diag}(\mathcal{M})$, so it's true in $\mathcal{N}$. Another one: if R is a relation and $\vec{m} = (m_1, m_2, \ldots)$ a vector of the appropriate length then $R^\mathcal{M}(\vec{m})$ is true iff $\mathcal{M} \models R(c_{m_1}, c_{m_2}, \ldots)$ (which is a sentence in $\mathcal{L}_\mathcal{M}$) and this is true if *and only if* $R(c_{m_1}, c_{m_2}, \ldots) \in \mathrm{Diag}(\mathcal{M})$, because $\mathrm{Diag}(\mathcal{M})$ is a *full theory*: either a sentence is in, or its negation is. Note also that you need formula induction to verify that for a formula ϕ of $\mathcal{L}$ some free variables and a vector $\vec{m} \in M^n$, and with ϕ' the corresponding *sentence* in $\mathcal{L}_\mathcal{M}$, we have $\mathcal{N} \models_\mathcal{L} \phi(\vec{m})$ iff $\mathcal{N} \models_{\mathcal{L}_\mathcal{M}} \phi'$. □

As a consequence of this powerful trick ("beefing up languages") we deduce

Theorem 5. *(weak going up) If $\mathcal{M}$ is an $\mathcal{L}$-structure such that M is infinite, then for any cardinal κ there's an $\mathcal{L}$-structure $\mathcal{N}$ with $|N| \geq \kappa$ and an elementary embedding $\mathcal{M} \to \mathcal{N}$.*

Proof. Apply the super-weak going up result to $\mathrm{Diag}(\mathcal{M})$ in $\mathcal{L}_\mathcal{M}$. □

It's "weak" because it doesn't get κ on the nose.
Now let's go the other way.

Lemma 6. *(weak going down) If $\mathcal{N}$ is an $\mathcal{L}$-structure and $X \subseteq N$ is a non-empty subset, and $\kappa = \max\{|X|, |\mathcal{L}| + \aleph_0\}$ then there's a substructure $\mathcal{M}$ of $\mathcal{N}$ with $X \subseteq M$ and $|M| \leq \kappa$.*

Proof. Easy: build N recursively starting with N_0, the constants and X, and then just throw in at stage $n+1$ the images of the functions applied to stage n. The proof that the resulting gadget isn't too big is elementary cardinal arithmetic. □

Note that if $|X| \geq |\mathcal{L} + \aleph_0$ then this lemma guarantees that $|M| = \kappa$, because $X \subseteq M$. However the lemma is still "weak" because it doesn't say that $\mathcal{M}$ is elementarily equivalent to $\mathcal{N}$.

Now let me explain the Tarski-Vaught criterion. First note that if $\mathcal{M}$ is a sub-$\mathcal{L}$-structure of $\mathcal{N}$ then for any *atomic* formula with free variables $v_1, v_2, \ldots, v_n$, and for any $\vec{m} \in M^n$ we have $\mathcal{M} \models \phi(\vec{a})$ iff $\mathcal{N} \models \phi(\vec{a})$; an atomic formula is just a relation between terms so this is basically trivial.

Now here's a criterion for a substructure to be an elementary substructure.

Lemma 7 (Tarski-Vaught criterion). *If $\mathcal{M}$ is a substructure of $\mathcal{N}$ then $\mathcal{M}$ is an elementary substructure of $\mathcal{N}$ if and only if the following is true:*

For all formulae ϕ with free variables $u_1, u_2, \ldots, u_n, v$ and for all $\vec{a} \in M^n$, if ($\exists c \in N$ with $\mathcal{N} \models \phi(\vec{a}, c)$)), then ($\exists b \in M$ with $\mathcal{N} \models \phi(\vec{a}, b)$).

Note that the TV criterion may as well be an iff, because if there is $b \in M$ with $\mathcal{N} \models \phi(\vec{a}, b)$ then, setting $c = b$, we deduce that there is $c \in N$ with $\mathcal{N} \models \phi(\vec{a}, c)$.

Note also that in both cases we have $\mathcal{N}$ modelling things. We don't seem to talk about $\mathcal{M}$ modelling anything!

In words: to check a substructure is an elementary substructure all you have to do is to check the "exists" part; you have to check that whenever there is c in N with $\phi(c)$ true, you can find b in M with $\phi(b)$ true. This is exactly the problem with $\mathbf{Z} \subset \mathbf{Q}$; $(\exists x)(x + x = 1)$ has a solution in $\mathbf{Q}$ but not one in $\mathbf{Z}$.

Proof. If $\mathcal{M}$ is an elementary substructure of $\mathcal{N}$ then the TV criterion is trivially true (apply the definition of "elementary substructure" to the statement $\exists x \phi(\vec{a})$). The other way is formula induction, the point being that $\wedge$, $\vee$ and $\neg$ are easy, $\forall$ follows from $\exists$, so it suffices to deal with $\exists$. Here's the argument in full. We have a formula ϕ with $n+1$ free variables $u_1, u_2, \ldots, u_n, x$. By formula induction (i.e., the inductive hypothesis) we know that for all $\vec{a} \in M^n$ and $b \in M$ we have $\mathcal{M} \models \phi(\vec{a}, b)$ iff $\mathcal{N} \models \phi(\vec{a}, b)$. Now assume the TV criterion, and choose $\vec{a} \in M^n$. We want to show that $\mathcal{M} \models \exists x \phi(\vec{a}, x)$ iff $\mathcal{N} \models \exists x \phi(\vec{a}, x)$. Here's how we do it. By definition, $\mathcal{M} \models \exists x \phi(\vec{a}, x)$ iff there is $b \in M$ such that $\mathcal{M} \models \phi(\vec{a}, b)$. By the inductive hypothesis this is true iff there is $b \in M$ such that $\mathcal{N} \models \phi(\vec{a}, b)$ (note that this latter iff is not vacuously true: $\phi(\vec{a}, b)$ doesn't have any free variables in but it might have lots of quantifiers, so changing universe is a big deal; we really are using the inductive hypothesis here). By the TV criterion this is true iff there exists $c \in N$ with $\mathcal{N} \models \phi(\vec{a}, c)$. And this is true iff $\mathcal{N} \models \exists x \phi(\vec{a}, x)$. And that's what we wanted! □

Before we prove the best possible going-down we need to introduce a way of making the TV criterion "concrete", that is, we need to be able to get from c to b. If ϕ is a formula in a language $\mathcal{L}$ and ϕ has free variables $u_1, u_2, \ldots, u_n, v$, and if $\mathcal{M}$ is an $\mathcal{L}$-structure, then we say that $f : M^n \to M$ is a *Skolem function for* $\exists v \phi$ if for all $\vec{a} \in M^n$, if $\mathcal{M} \models \exists v \phi(\vec{a}, v)$ then $\mathcal{M} \models \phi(\vec{a}, f(\vec{a}))$. Of course by the axiom of choice, such things always exist. So now we use this trick (due to Skolem) to prove

Theorem 8. *(Lowenheim-Skolem going down) If $\mathcal{N}$ is an $\mathcal{L}$-structure and $X \subseteq N$ is a subset such that $|X| = \kappa \geq |\mathcal{L}| + \aleph_0$ then there exists $\mathcal{M} \prec \mathcal{N}$ with $X \subseteq M$ and $|\mathcal{M}| = \kappa$.*

This is the proper "going down" theorem. We get an elementary embedding and a given cardinality (assumed sufficiently large).

Proof. Let $\mathcal{L}'$ denote the language $\mathcal{L}$ plus a new function-symbol $f_{\exists v \phi}$ for all $\mathcal{L}$-formulae ϕ and for all free variables v of ϕ. If ϕ has $r + 1$ free variables in all then set $n_{f_{\exists v \phi}} = r$. Note that $|\mathcal{L}'| \leq \kappa$ because there are only countably many variable symbols! Now if $\mathcal{N}$ is an $\mathcal{L}$-structure then use Skolem functions to make $\mathcal{N}$ an $\mathcal{L}'$-structure. By "weak going down" there's a sub-$\mathcal{L}'$-structure $\mathcal{M}$ with $X \subseteq M$ and $|M| = \kappa$. Skolem's observation is that, as $\mathcal{L}$-structures, we have $\mathcal{M} \prec \mathcal{N}$. To verify this it suffices to use the TV criterion. Say ϕ is a formula with variables u_1, u_2,…,u_n,v, and $\vec{a} \in M^n$. Then $\exists c \in N : \mathcal{N} \models \phi(\vec{a}, c))$ implies that $\mathcal{N} \models \phi(\vec{a}, f^{\mathcal{N}}_{\exists v \phi}(\vec{a}))$. But $b := f^{\mathcal{N}}_{\exists v \phi}(\vec{a}) \in M$! So $\mathcal{N} \models \phi(\vec{a}, b)$ and this is precisely TV. □

Corollary 9 (going up theorem.)**.** *If T is an $\mathcal{L}$-theory and if T has an infinite model then for all $\kappa \geq |\mathcal{L}| + \aleph_0$, T has a model of size κ.*

Proof. By weak going up, T has a model $\mathcal{N}$ with $|N| \geq \kappa$. Take a subset X of N of size κ. By going down there's an $\mathcal{L}$-structure $\mathcal{M}$ and an elementary embedding $\mathcal{M} \to \mathcal{N}$ with $|M| = \kappa$. Now $M \prec N$ implies $M \equiv N$, and hence $M \models T$. □

7 Completeness and catagorical-ness.

An $\mathcal{L}$-theory T is *complete* if any two models of T are elementarily equivalent! Non-example: the theory of fields of characteristic zero, because the sentence $\exists x : x^2 = 2$ is true in some and not true in others. Trivial example: a full theory, for example $\mathrm{Th}(\mathcal{M})$ for some structure $\mathcal{M}$ (recall that the full theory is just a list of *every* sentence which is true in $\mathcal{M}$).

A sentence ϕ of $\mathcal{L}$ is a *consequence* of a theory T if, for all models $\mathcal{M}$ of T we have $\mathcal{M} \models \phi$. I have a moral objection to this definition: a consequence of a bunch of sentences should be something which one can prove from the sentences in some formal proof-theory type of way. But Ambrus says that he doesn't want to do any proof theory, and anyway, by the completeness theorem, which he doesn't want to do because he doesn't want to do any proof theory, one can formalise my definition and then prove that it's equivalent to his definition.

Notation: if ϕ is a consequence of T then we write $T \models \phi$. So "$\models$" is transitive: if $\mathcal{M} \models T$ and $T \models \phi$ then $\mathcal{M} \models \phi$.

A useful remark is:

Lemma 10. *An $\mathcal{L}$-theory T is complete iff for all sentences ϕ, either $T \models \phi$ or $T \models \neg\phi$.*

Proof. If T has no models at all then $T \models \phi$ for all sentences ϕ and the lemma is true in this case. So let's assume T has models. If for all sentences ϕ either $T \models \phi$ or $T \not\models \phi$, then for any two models $\mathcal{M}$ and $\mathcal{N}$ and any sentence ϕ, if $\mathcal{M} \models \phi$ then, because $\mathcal{M} \models T$, we can't have $T \models \neg\phi$, so we must have $T \models \phi$, and this means that $\mathcal{N} \models \phi$. Hence T is complete. Conversely, if T is complete and has a model $\mathcal{M}$ then for a sentence ϕ, either $\mathcal{M} \models \phi$, in which case all models for T satisfy ϕ, and hence $T \models \phi$, or $\mathcal{M} \models \neg\phi$ in which case the same argument shows that $T \models \neg\phi$. □

If κ is a cardinal and $\mathcal{L}$ is a language, we say that an $\mathcal{L}$-theory T is *κ-categorical* if any two models of T of cardinality κ are *isomorphic* to each other! We'll see examples of this later. In words this says that "there's at most one model of size κ". I guess vector spaces will give examples (use a standard basis argument), and so will algebraically closed fields of a fixed characteristic (use a transcendence degree argument), for κ large enough. What I'm saying is that it's not hard to think of interesting categorical theories. There is also a really stupid class of examples: if a language contains uncountably many constants and the theory contains sentences demanding that all these constants are distinct, then clearly the theory is $\aleph_0$-categorical, because there are no countable models at all, and in general being κ-categorical is not going to be of much use at all if $|\mathcal{L}| > \kappa$. On the other hand, here's an amazing consequence of going up and going down, which shows how strong κ-categoricality is if $\kappa \geq |\mathcal{L}|$ (which is true in all the sensible examples above):

Proposition 11 (Vaught's test)**.** *If T is a satisfiable theory which is κ-categorical for some infinite $\kappa \geq |\mathcal{L}|$, and such that every model of T is infinite, then T is complete!*

This is kind of amazing because it seemed to me *a priori* to be hard to think of complete theories.

Proof. Take two models $\mathcal{M}$ and $\mathcal{N}$ of T. We need to show that $\mathcal{M} \equiv \mathcal{N}$. By assumption, M and N are infinite. So by going up and going down we can find $\mathcal{M}'$ with $\mathcal{M}' \equiv \mathcal{M}$ and $|\mathcal{M}'| = \kappa$, and we can also find $\mathcal{N}'$ with $\mathcal{N}' \equiv \mathcal{N}$ and $|\mathcal{N}'| = \kappa$. By κ-catagoricality we have $\mathcal{M}' \cong \mathcal{N}'$ and hence certainly $\mathcal{M}' \equiv \mathcal{N}'$. So we're home. □

8 Back and forth.

This is a technique used to construct isomorphisms between models and hence verify κ-categoricalness.

Let ACF denote the theory of algebraically closed fields (in the language of rings). If p is a prime then let ACF_p denote the theory of alg closed fields of characteristic p, that is, add in the extra axiom $1 + 1 + 1 + 1 + \cdots + 1 = 0$. Let ACF_0 denote the theory of alg closed fields of char 0, that is, add in infinitely many extra axioms $\neg(1 + 1 + 1 + \cdots + 1 = 0)$, one for each prime.

Theorem 12. *The following theories are complete.*
(i) Dense linear orderings without endpoints
(ii) DAG (Torsion-free divisible abelian groups of size greater than 1).
(iii) ACF_p *and* ACF_0.

Proof. Use Vaught's test.

(i) First we need to check that every model is infinite. This is true because given x in a dense linear ordering without endpoints, x isn't an endpoint, so $x < x_1 < x_2 < \cdots$.

So now it suffices to check that the theory is $\aleph_0$-categorical. We use "back and forth". Take two countable models $\mathcal{M}$ and $\mathcal{N}$. List the elements of M (say $m_1, m_2, \ldots$) and N. We recursively define subsets $X_n \subseteq M$ and $Y_n \subseteq N$ and order-preserving bijections $f_n : X_n \to Y_n$ such that if $i \leq n$ then $x_i \in X_n$ and $X_i \subseteq X_n$ and such that the f_i are compatible. This will be enough, because the union of the f_n will be the bijection we seek (easy check that it works).

Here's the construction. Set $X_1 = \{m_1\}$ and $Y_1 = \{n_1\}$ and let f_1 be the bijection. For the induction step we have X_i and Y_i; let X_i' be $X_i \cup \{m_{i+1}\}$ (this might not be X_{i+1}) and extend f_i to an injection $f_i' : X_i' \to Y$ by the axioms: if $m_{i+1} \in X_i$ then we do nothing; if it's new and in the middle of X_i then use denseness of Y, and if it's new and at an end then use no-endpoint-ness of Y.

Now let Y_i' be the image of f_i' and let Y_{i+1} be $Y_i' \cup \{n_{i+1}\}$. The inverse of f_i' gives a bijection $Y_i' \to X_i'$ and we extend this to an injection $Y_{n+1} \to M$; let the image be X_{n+1} and let f_{n+1} be the inverse of the induced bijection $Y_{n+1} \to X_{n+1}$.

(ii) non-zero torsion-free DAGs. By torsion-freeness and divisible-ness these gadgets are easily checked to be canonically **Q**-vector spaces. So now any models are infinite and if $\kappa > \aleph_0$ then a trivial argument involving counting bases shows that the theory is κ-categorical.

(iii) Ambrus claimed that ACF_p and ACF_0 were κ-categorical for any uncountable κ, which is true by a transcendence basis argument. Ambrus tried to give a proof "by construction" but it seemed to have a hole in it when κ was a limit cardinal. However it worked for $\aleph_1$ and this is all that one needs to deduce completeness. The idea was: given two algebraically closed fields of cardinality $\aleph_1$ and the same characteristic, enumerate all the elements, indexing by ω_1, and then construct recursively an isomorphism between ever-larger subfields a la the dense linear order proof. The problem with the proof the way Ambrus presented it is that at limit ordinals one wants to take unions and be sure that one doesn't use up all of one field but not all of the other! This problem doesn't arise in ω_1 though. □

Homework. A total order is *discrete* if every element is either the top or has a unique successor, and if every element is either the bottom or has a unique predecessor. Ambrus said that he thought that the theory of discrete total orders with a min but no max is not κ-categorical for any $\kappa > \aleph_0$. I think that a proof of this might be: take the positive integers and then glue κ copies of **Z** on the end; there's one model. Now glue one more **Z** on the end and there's a second model. I don't think these can be isomorphic: one has a cofinal set of the form $x, s(x), s(s(x)), \ldots$ (with s the successor function) and the other doesn't.

Theorem 13. *(Ax) For any sentence in the language of rings, TFAE:*

(i) $\mathrm{ACF}_0 \models \phi$

(ii) $\mathbf{C} \models \phi$

(iii) $\overline{\mathbf{F}}_p \models \phi$ *for infinitely many primes p*

(iv) $\mathrm{ACF}_p \models \phi$ *for all but finitely many p.*

Proof. (i) implies (ii) trivially. Conversely (ii) implies (i) because ACF_0 is complete—we just proved this.

We'll show (i) implies (iv) implies (iii) implies (i). The first and last arrow require a little work; the middle one is trivial.

(i) implies (iv) uses

Lemma 14. *If T is an $\mathcal{L}$-theory and ϕ is an $\mathcal{L}$-sentence with $T \models \phi$ then there's a finite subset $S \subseteq T$ with $S \models \phi$.*

Proof. This is just the compactness theorem. If the lemma were false then apply the compactness theorem to $T \cup \{\neg\phi\}$ for an immediate contradiction: the negation of the conclusion of the lemma implies that any finite subset of $T \cup \{\neg\phi\}$ has a model. □

Using this lemma, (i) implies (iv) is easy: if $\mathrm{ACF}_0 \models \phi$ then there's a finite subset S of ACF_0 which models ϕ. But ACF_0 is ACF union countably many axioms of the form $1+1+1+\cdots+1 \neq 0$, and S will only mention finitely many of these axioms, so there's N such that no axiom of the form $N = 0$ is true in S, so any algebraically closed field K of characteristic $p > N$ will have $K \models S$ and hence $K \models \phi$. So $\mathrm{ACF}_p \models \phi$ for all $p > N$.

Finally we must show (iii) implies (i). But somehow this is trivial. We do it by contradiction. Say ϕ is a sentence such that (iii) holds for ϕ but (i) is false for ϕ. By completeness of ACF_0 and a "useful remark" above, we must have $\mathrm{ACF}_0 \models \neg\phi$. But then by (i) implies (iv) applied to $\neg\phi$ we deduce that $\mathrm{ACF}_p \models \neg\phi$ for all sufficiently large p, and hence that $\overline{\mathbf{F}}_p \models \neg\phi$ for all sufficiently large p, which contradicts (iii). □

9 Axiomatisation.

A class $\mathcal{C}$ of $\mathcal{L}$-structures is *axiomatisable* if there's a theory T such that $\mathcal{M} \in \mathcal{C}$ iff $\mathcal{M} \models T$. We say T *axiomatises* $\mathcal{C}$.

We say $\mathcal{C}$ is *finitely axiomatisable* if there's a finite T that axiomatises it.

If one wants to prove that a certain axiomatisable theory is not finitely axiomatisable, here's a powerful approach which is a consequence of compactness:

Lemma 15. *If the class of all models of a theory T is finitely axiomatisable then there's a finite subset S of T which axiomatises the class.*

Proof. Take any finite set of sentences that axiomatises the class. Now $\wedge$ them all together and you get one sentence ϕ that axiomatises the class! If $\mathcal{M} \models T$ then $\mathcal{M} \models \phi$, by definition. So $T \models \phi$ by definition. So by compactness there's a finite subset S of T such that $S \models \phi$ (we proved this earlier). The claim is that this S works and the proof is easy: if $\mathcal{M} \models T$ then certainly $\mathcal{M} \models S$. On the other hand if $\mathcal{M} \models S$ then $\mathcal{M} \models \phi$ and hence $\mathcal{M} \models T$! ☐

Theorem 16. *The following classes are axiomatisable but not finitely axiomatisable.*
(i) The class of torsion-free abelian groups.
(ii) The class of fields of characteristic zero.
(iii) The class of algebraically closed fields.

Proof. (i) is axiomatisable: write down the axioms for an abelian group (finitely many) and then just add an axiom $phi_n : nx = 0 \implies x = 0$ for every $n \geq 1$. But if this class were finitely axiomatisable then there would be a finite subset that did the job, and such a finite set would only mention ϕ_n for a finite set of n, and if it didn't mention ϕ_p for p prime then $\mathbf{Z}/p\mathbf{Z}$ kills you.

(ii) Just the same: there are finite fields obeying any finite set of equations of the form $n \neq 0$ for n an integer.

(iii) Just the same. It suffices to show that for any prime p there is a field which is not algebraically closed because it has an extension of degree p, but which has no extension of degree $1 < n < p$. There's an extension of $\mathbf{Z}/\ell\mathbf{Z}$ which works, by Galois theory: its absolute Galois group is $\mathbf{Z}_p \times H$ with H the product of all the other q-adic integers, q not p, so there's a field whose algebraic closure has Galois group $\mathbf{Z}_p$ and this will do. ☐

Remark 17. It is not the case in general that theories that have only infinite models will not be finitely axiomatisable: for example atomless Boolean algebras, or (non-empty) total orders without endpoints only have infinite models.

10 Boolean rings and algebras.

A *Boolean ring* is a commutative ring R with a 1 such that $x^2 = x$ for all x. If I is a maximal ideal of R then R/I is a field with $x^2 = x$ and hence $R/I = \mathbf{Z}/2\mathbf{Z}$. If $x \in R^\times$ then there's a y with $xy = 1$, so $x = x^2y = xy = 1$ and hence $R^\times = \{1\}$. If $0 \neq x \in R$ then $1 - x$ isn't 1 and hence isn't a unit, so there's a maximal ideal containing $1 - x$ and this maximal ideal does of course not contain x. Hence the intersection of all the maximal ideals is zero, and hence R injects into a product of $\mathbf{Z}/2\mathbf{Z}$s. Conversely any subring of a product of $\mathbf{Z}/2\mathbf{Z}$s (for example the subring of the countable product of $\mathbf{Z}/2\mathbf{Z}$s consisting of functions which are constant away from a finite set) is a Boolean ring.

A *Boolean algebra* is a set X equipped with constants 0 and 1, binary operators $\wedge$, $\vee$, and a unitary operator $\neg$, and satisfying a bunch of axioms (each binary operator distributes over the other, for example). It's a rather tedious but completely elementary check to see that given a Boolean algebra one can put a Boolean ring structre on it by defining $x+y = (x \wedge (\neg y)) \vee (y \wedge (\neg x))$, and $xy = xy$, and in fact the categories of Boolean rings and Boolean algebras are isomorphic. A basic example of a Boolean algebra is the set of subsets of a set, with $\wedge$ being intersection, $\vee$ being union and $\neg$ being complement. The fact proved above that every Boolean ring embeds into a product of $\mathbf{Z}/2\mathbf{Z}$s just says that in fact any Boolean algebra is a subalgebra of a power set algebra.

One can define a partial ordering on a Boolean algebra: $x \leq y$ iff $xy = x$. From the point of view of sets this just says $x \subseteq y$. An *atom* in a Boolean algebra is x with $x \neq 0$ but $y \leq x$ implies that either $y = x$ or $y = 0$.

Finite Boolean algebras are all checked to be equal to power set algebras (one can embed it into a power set algebra using maximal ideals and then use basic facts about coprime maximal ideals to construct all the idempotents one needs). In particular all finite Boolean algebras of size greater than 1 contain atoms (but this was clear anyway). More generally, a Boolean algebra B is said to be *finitely-generated* if there's a finite subset X of B such that B is the smallest sub-Boolean-algebra of B containing X, and the theorem is that a finitely-generated Boolean algebra is isomorphic to the power set of a finite set.

A useful fact is: if B is generated by the finite set G then the map from maximal ideals of B to finite subsets of G given by sending I to $\{g \in G : g \in I\}$ is an injection. Indeed, $g \notin I$ iff $1 - g \in I$ because the residue field is $\mathbf{Z}/2\mathbf{Z}$, and hence for I and J maximal ideals with the same intersection with G, the induced maps $B \to \mathbf{Z}/2\mathbf{Z}$ agree on a generating set and hence agree everywhere, so the kernels are the same.

Proposition 18. *If B is a Boolean algebra generated by a set G then every $b \in B$ can be written as*

$$b = c_1 \vee c_2 \vee \ldots \vee c_n$$

for some $n \in \mathbf{Z}_{\geq 0}$, where each $c_j \in B$ is of the form

$$c_j = d_{1j} \wedge d_{2j} \wedge \ldots \wedge d_{ij}$$

where each d_{ij} is either in G or of the form $\neg g$ for some $g \in G$.

Proof. For $H \subseteq G$ write $B(H)$ for the subalgebra of B generated by H. Then $B = \cup_H B(H)$ where H runs through the *finite* subsets of G (because the union is a subalgebra containing G) and hence WLOG G is finite. Hence WLOG B is the power set of a finite set. Any element of B is a finite union of singletons, so it suffices to prove that any singleton is of the form c_j above. Well, we have $B = P(X)$ generated by G, so, for $x \in X$, let $Y \in B$ be the obvious intersection: for each $g \in G$, throw in g if $x \in g$ and $\neg g$ if $x \notin g$. If we've recovered $\{x\}$ then we're done. Clearly $x \in Y$. Now if $z \in Y$ and $z \neq x$ then, because the maximal ideals of B biject with X, the maximal ideals corresponding to x and z do not coincide, so there's some $g \in G$ with either $x \in g$ and $z \notin g$ or $x \notin g$ and $z \in g$ (the subsets of G separate the maximal ideals). But this is impossible by definition. □

Similarly every $b \in B$ is an $\wedge$ of some c'_j, with the c'_j's all the $\vee$s of d'_{ij}s, each of which is either in G or its complement is in G.

The theory of atomless Boolean algebras (which is easily shown to be finitely axiomatisable) is $\aleph_0$-categorical. In fact one can prove this using back-and-forth. The basic thing you need is of course that if you have two atomless Boolean algebras X and Y, and finitely-generated subalgebras A and B with a given isomorphism between them, and you choose $x \in X$ with $x \notin A$, and let A' denote the smallest subalgebra of X containing A and x, then there's a subalgebra B' of Y such that the given isomorphism $A = B$ extends to an isomorphism $A' = B'$. The proof is this. We know that A and A' are finitely-generated Boolean algebras and are hence isomorphic to power sets: say $A = P(U)$ and $A' = P(U')$. The injection $A \to A'$ sends orthogonal idempotents to orthogonal idempotents and hence sends distinct elements of U to disjoint subsets of U'. The fact that 1 goes to 1 means that the subsets of U' cover U' and hence the injection $A \to A'$ is induced by a surjection $U' \to U$. For every $u \in U$ with more than one pre-image the corresponding idempotent in B isn't an atom, so choose something less than it and throw it in to B'; this way one can build a B' which naturally contains a copy of A'; now cut B' down until it's exactly A'.

An interesting remark is that the theory of atomless Boolean algebras is not κ-categorical for any $\kappa > \aleph_0$. But I don't really understand why not.[2] 2

[2]Why not?

11 Quantifier elimination.

If ϕ is a formula with free variables $u_1, u_2, \ldots, u_n$, and if $\mathcal{M}$ is an $\mathcal{L}$-structure then we write $\mathcal{M} \models \phi$ if for all $\vec{a} \in M^n$ we have $\mathcal{M} \models \phi(\vec{a})$. If ϕ is a sentence then this new definition coincides with the old. If ϕ is a formula as above then let its *closure* $\overline{\phi}$ be the sentence $\forall u_1 \forall u_2 \ldots \forall u_n \phi$. Then one checks $\mathcal{M} \models \phi$ iff $\mathcal{M} \models \overline{\phi}$.

Note however that $\overline{\neg\phi} \neq \neg\overline{\phi}$! In other words, for a random formula ϕ we might have $\mathcal{M} \not\models \phi$ and $\mathcal{M} \not\models \neg\phi$: this occurs when ϕ is "sometimes true, sometimes false".

If T is a set of $\mathcal{L}$-formulae and $\mathcal{M}$ is an $\mathcal{L}$-structure then we say $\mathcal{M} \models T$ if $\mathcal{M} \models \phi$ for all $\phi \in T$.

If T is a set of $\mathcal{L}$-formulae and ϕ is an $\mathcal{L}$-formula then we say $T \models \phi$ if for any $\mathcal{M}$ with $\mathcal{M} \models T$ we have $\mathcal{M} \models \phi$.

If T and Γ are sets of formulae then we write $T \models \Gamma$ if $T \models \phi$ for all $\phi \in \Gamma$.

We say that a set of formulae is *satisfiable* if there's an $\mathcal{L}$-structure $\mathcal{M}$ with $\mathcal{M} \models T$.

Let me (KB, not Ambrus) go through these definitions again. If T is a set of formulae, let $\overline{T}$ denote the set $\{\overline{\phi} : \phi \in T\}$, that is, just replace every formula that has some free variables with its closure. Ambrus noted that $\mathcal{M} \models \phi$ iff $\mathcal{M} \models \overline{\phi}$. I think that the same sort of thing will be true for these other notions. If T is a set of $\mathcal{L}$-formulae and $\mathcal{M}$ is an $\mathcal{L}$-structure then $\mathcal{M} \models T$ iff $\mathcal{M} \models \phi$ for all $\phi \in T$ iff $\mathcal{M} \models \overline{\phi}$ for all $\phi \in T$ iff $\mathcal{M} \models \psi$ for all $\psi \in \overline{T}$ iff $\mathcal{M} \models \overline{T}$. Similarly T is a set of formulae and ϕ is a formula then $T \models \phi$ iff ($\mathcal{M} \models T \implies \mathcal{M} \models \phi$) iff ($\mathcal{M} \models \overline{T} \implies \mathcal{M} \models \overline{\phi}$) iff $\overline{T} \models \overline{\phi}$. Next, if T and Γ are sets of formulae then $T \models \Gamma$ iff $T \models \phi$ for all $\phi \in \Gamma$ iff $\overline{T} \models \overline{\phi}$ for all $\phi \in \Gamma$ iff $\overline{T} \models \psi$ for all $\psi \in \overline{\Gamma}$ iff $\overline{T} \models \overline{\Gamma}$. Finally, T is satisfiable iff $\overline{T}$ is satisfiable.

Lemma 19 (deduction lemma)**.** *If ϕ is a sentence (note: not a formula!) and if Γ is a set of formulae such that $\Gamma \cup \{\phi\}$ is not satisfiable then $\Gamma \models \neg\phi$.*

Proof. Trivial proof by contradiction. If the statement $\Gamma \models \neg\phi$ is false, then there's $\mathcal{M}$ such that $\mathcal{M} \models \Gamma$ and $\mathcal{M} \not\models \neg\phi$. But ϕ is a sentence, and hence $\mathcal{M} \models \phi$. So $\mathcal{M} \models \Gamma \cup \{\phi\}$, contradicting our assumption. □

Lemma 20. *If Δ is a theory (i.e., a bunch of sentences in a language $\mathcal{L}$) and if Γ is a set of formulae, and if $\Delta \cup \Gamma$ is not satisfiable then there is $\phi_1, \phi_2, \ldots, \phi_n \in \Delta$ such that $\Gamma \models \neg(\phi_1 \wedge \phi_2 \wedge \ldots \wedge \phi_n)$.*

Proof. Easy. If $\Gamma \cup \Delta$ is not satisfiable then by compactness there's a finite subset $\phi_1, \phi_2, \ldots, \phi_n$ of Δ such that $\Gamma \cup \{\phi_1, \phi_2, \ldots, \phi_n\}$ is not satisfiable. So $\Gamma \cup \{\phi_1 \wedge \phi_2 \wedge \ldots \wedge \phi_n\}$ is not satisfiable. So $\Gamma \models \neg(\phi_1 \wedge \phi_2 \wedge \ldots \wedge \phi_n)$ by the previous lemma. □

If $\mathcal{L}$ is a language and ϕ is a formula with free variables $x_1, x_2, \ldots, x_n$, and if $c_1, c_2, \ldots, c_n$ are constants not appearing in $\mathcal{L}$, and if $\mathcal{L}'$ is the language obtained from $\mathcal{L}$ by adding these n new constants, then let $\phi(\vec{c})$ denote the sentence in $\mathcal{L}'$ obtained by substituting c_i for x_i.

Lemma 21. *If Γ is a set of $\mathcal{L}$-formulae and if ϕ and $\mathcal{L}'$ are as above, and if $\Gamma \models_{\mathcal{L}'} \phi(\vec{c})$ then $\Gamma \models_{\mathcal{L}} \phi$.*

Remark: $\Gamma \models_{\mathcal{L}} \phi$ iff $\Gamma \models_{\mathcal{L}'} \phi$.

Proof. If $\mathcal{M}$ is an $\mathcal{L}$-structure with $\mathcal{M} \models_{\mathcal{L}} \Gamma$ and if $\vec{a} \in M^n$ then we need to show that $\mathcal{M} \models_{\mathcal{L}} \phi(\vec{a})$. But if $\mathcal{M}'$ is the $\mathcal{L}'$-structure obtained from $\mathcal{M}$ by defining $(c_i)^{\mathcal{M}'} = a_i$ for all $i \leq n$ then $\mathcal{M}' \models_{\mathcal{L}'} \Gamma$ and hence by assumption $\mathcal{M}' \models_{\mathcal{L}'} \phi(\vec{c})$, so by an easy formula induction $\mathcal{M}' \models_{\mathcal{L}'} \phi(\vec{a})$, and hence $\mathcal{M}' \models_{\mathcal{L}} \phi(\vec{a})$. □

Here's the final lemma before we state the main theorem of this section. Say that a formula ψ is *open* if it has no quantifiers. The big definition: an $\mathcal{L}$-theory T has *quantifier elimination* if for every formula ϕ there's an open formula ψ with $V(\phi) = V(\psi)$ (i.e. the same free variables) and such that $T \models (\phi \iff \psi)$. Note: even though ϕ and ψ might have $n > 0$ free variables, the statement $T \models (\phi \iff \psi)$ is super-strong: it says that for any $\vec{a} \in M^n$, $T \models (\phi(\vec{a}) \iff \psi(\vec{a}))$ and hence $\phi(\vec{a})$ is true in T if and only if $\psi(\vec{a})$ is.

An example: Cramer's rule is an example of quantifier elimination! Cramer's rule says that a square matrix has an inverse if and only if its determinant is non-zero. The existence of an inverse is n^2 existence statements satisfying n^2 equations; the determinant is just one assertion about something being non-zero.

Stupid remark: if ϕ is a sentence, for example $(\exists x)(x = x)$, then the formula ψ above would have to be an open formula with no free variables and hence with no variables at all. If furthermore the language has no constants then there arguably are no possibilities for ψ at all. To fix this one can demand that $\mathcal{L}$ always contains at least one constant, or one can drop the constraint that $V(\phi) = V(\psi)$. Because we're always demanding that our models are non-empty, we could always demand that our theories have a constant element and this wouldn't change anything we've said.

We have seen that completeness is a very desirable property; you can prove a theorem for all models at once by just checking it for one special model that might have lots of extra nice properties. It turns out that quantifier elimination sometimes implies completeness, so quantifier elimination should also be thought of as a desirable property of a theory.

Lemma 22. *Let T be an $\mathcal{L}$-theory. Assume that for all open formulae ϕ with free variables $V(\phi) = \{u_1, u_2, \ldots, u_n, x\}$, there's an open formula ψ with $V(\psi) = \{u_1, u_2, \ldots, u_n\}$ and such that $T \models ((\exists x \phi) \iff \psi)$. Then T has quantifier elimination.* {22}

I guess that this is somehow unsurprising: to remove all the quantifiers from a general formula you "start in the middle", removing the first $\exists$, and then work out.

Proof. Straightforward formula induction. If ϕ is atomic then it's open so $\psi = \phi$ will do. If $\phi = \neg\theta$ then by induction there's open α such that $V(\theta) = V(\alpha)$ and $T \models \theta \iff \alpha$. Now it's not hard to check that $\psi := \neg\alpha$ is open, $V(\psi) = V(\phi)$, and $T \models \phi \iff \psi$ (but this last point does need checking! The point is that if $T \models \theta \iff \alpha$ then for any model for T we have $\theta(\vec{a})$ true iff $\alpha(\vec{a})$ true, so $\theta(\vec{a})$ false iff $\alpha(\vec{a})$ false.)

Now $\wedge$ and $\vee$ are obvious. What's left is $\exists$ and this is of course where we use the assumption. If ϕ is $\exists x\theta$ then by induction there's open α with $T \models (\theta \iff \alpha)$. By the assumption of the lemma applied to α, there's an open ψ with $T \models ((\exists x\alpha) \iff \psi)$. The claim is that ψ will do, that is, that $T \models (\phi \iff \psi)$. To prove this, take $\mathcal{M}$ with $\mathcal{M} \models T$. Then $\mathcal{M} \models (\theta \iff \alpha)$ and $\mathcal{M} \models ((\exists x\alpha) \iff \psi)$. We claim that $\mathcal{M} \models (\phi \iff \psi)$. To verify this, take an arbitrary $\vec{a} \in M^n$. We need to check that $\mathcal{M} \models (\exists x\theta(\vec{a}, x))$ iff $\mathcal{M} \models \psi(\vec{a})$. Well, $\mathcal{M} \models (\exists x\theta(\vec{a}, x)$ if and only if there's $b \in M$ with $\mathcal{M} \models \theta(\vec{a}, b)$. This is true if and only if there's $b \in M$ with $\mathcal{M} \models \alpha(\vec{a}, b)$. And this is true if and only if $\mathcal{M} \models \exists x\alpha(\vec{a}, x)$. And this is true if and only if $\mathcal{M} \models \psi(\vec{a})$. So we're home. □

So we're working towards the proof of the "big theorem about quantifier elimination". But to even explain the *statement* of the theorem we need the notion of "algebraically prime models".

Let T be an $\mathcal{L}$-theory. Let $T_\forall$ denote the set of all $\mathcal{L}$-sentences $\overline{\phi}$, where ϕ runs through the open $\mathcal{L}$-formulae such that $T \models \phi$. Recall that $T \models \phi$ iff $T \models \overline{\phi}$. Vaguely, for a sentence to be in $T_\forall$, it must be deducible from T, it can only have $\forall$s in, no $\exists$s, and all the $\forall$s must be at the very beginning.

The theory $T_\forall$ is called the set of all *universal consequences* of T. By definition if $\mathcal{M} \models T$ then $\mathcal{M} \models T_\forall$. But the really neat thing about the definition is that the truth of $T_\forall$ is inherited by substructures.

Lemma 23. *If $\mathcal{M} \subseteq \mathcal{N}$ are $\mathcal{L}$-structures and if $\mathcal{N} \models T$ then $\mathcal{M} \models T_\forall$.*

Proof. Choose $\psi \in T_\forall$. Then $\psi = \overline{\phi}$ with ϕ an open formula and $T \models \phi$. So if $\mathcal{N} \models T$ then $\mathcal{N} \models \phi$, and ϕ has no quantifiers, so $\mathcal{M} \models \phi$, so $\mathcal{M} \models \overline{\phi}$. □

So here's examples of this notion. If $\mathcal{L} = \{0, 1, +, -, *\}$ is the language of rings and T is the theory of fields, then I claim that $T_\forall$ is the theory of integral domains. For if R is an integral domain, then it is a subring of its field of fractions K, and $K \models T$ and hence $R \models T_\forall$. So any integral domain models $T_\forall$. Conversely if S is a ring and $S \models T_\forall$ then we need to check that S is an integral domain, so we need to check $0 \neq 1$, that $xy = yx$ and that $xy = 0 \implies x = 0 \vee y = 0$. So it suffices to check that these three statements are in $T_\forall$. Well, $0 \neq 1$ is certainly a consequence of the theory of fields, and it's an open formula, so $0 \neq 1 \in T_\forall$. Similarly $xy = yx$ is true for all fields and is an open formula, so $\forall x \forall y xy = yx$ is in $T_\forall$. Finally $x = 0 \vee y = 0 \vee xy \neq 0$ is true for all fields, so again we're done.

Similarly if $\mathcal{L}$ is the language of rings and T is the theory of algebraically closed fields, then $T_\forall$ is again the theory of integral domains.

Similarly if $\mathcal{L} = \{0, *\}$ is the language of groups (let's not throw in inverses into the language, let's make them part of the theory) and T is the theory of non-zero torsion-free divisible abelian groups, then $T_\forall$ is the theory of commutative monoids with cancellation and no torsion (note that for a monoid to have no torsion is just the infinitely many axioms (one for each $n \geq 1$) that $nx = ny$ implies $x = y$). Note that $\{0\} \models T_\forall$ (because it's a substructure of a non-zero torsion-free divisible abelian group) but it's not non-zero.

If T is a theory in a language $\mathcal{L}$ then we say that T has *algebraically prime models* if for every model $\mathcal{M}$ of $T_\forall$ there's a model $\mathcal{N}$ of T and an $\mathcal{L}$-embedding $\eta : \mathcal{M} \to \mathcal{N}$ such that for every $\mathcal{L}$-embedding $\mathcal{M} \subseteq \mathcal{N}'$ with $\mathcal{N}' \models T$ there's an embedding (not assumed unique) $\mathcal{N} \to \mathcal{N}'$ making the obvious diagram commute. We say that $\mathcal{N}$ is an *algebraically prime extension* of $\mathcal{M}$.

As examples, the theory of fields has algebraically prime extensions (given an integral domain, take its field of fractions) and in this case the morphism $\mathcal{N} \to \mathcal{N}'$ is unique. Similarly the theory of algebraically closed fields has algebraically prime extensions, but here the algebraically prime extension of an integral domain is an algebraic closure of its field of fractions, and in particular the embedding $\mathcal{N} \to \mathcal{N}'$ (with notation as above) will not in general be unique.

As a final example, the theory of non-zero torsion-free divisible abelian groups has algebraically prime models too: if M is a commutative torsion-free monoid with cancellation and $K(M)$ is the abelian group generated by M then $K(M) \otimes \mathbf{Q}$ will be an algebraically prime extension of M, if this is non-zero. If however it is zero then $M = 0$ and in this case $\mathbf{Q}$ is an algebraically prime extension of M and again there is severe non-uniqueness in making the diagram commute.

Another definition: if $\mathcal{M} \subseteq \mathcal{N}$ are $\mathcal{L}$-structures then we say $\mathcal{M}$ is *simply closed in* $\mathcal{N}$ (written $\mathcal{M} \prec_s \mathcal{N}$) if for all open formulae ϕ with free variables $u_1, u_2, \ldots, u_n, x$ and for all $\vec{a} \in M^n$, if $\mathcal{N} \models \exists x \phi(\vec{a}, x)$ then $\mathcal{M} \models \exists x \phi(\vec{a}, x)$.

Note that if $\mathcal{M} \prec \mathcal{N}$ then (applying the definition of $\prec$ to the formula $\exists x \phi(\vec{u}, x)$) we deduce that $\mathcal{M} \prec_s \mathcal{N}$. I think Ambrus said the converse was false though; $\prec_s$ is strictly weaker than $\prec$. He didn't give any examples though, and that's a shame because my favourite example $\mathbf{Z} \not\prec \mathbf{Q}$ also has the property that $\mathbf{Z}$ is not simply closed in $\mathbf{Q}$ either, for the usual reasons: let ϕ be $x + x = 1$ and set $n = 0$. Ambrus said, when I asked him, that he didn't know an example offhand either.

The theorem which we'll prove in this section is about a theory T in a language $\mathcal{L}$ which is *assumed to have a non-empty set of constant symbols*. This is for the silly reason mentioned earlier in the definition of quantifier elimination.

{24}

Theorem 24. *If T is an $\mathcal{L}$-theory ($\mathcal{L}$ assumed to have a constant) such that*
(i) T has algebraically prime models, and
(ii) If $\mathcal{M} \subseteq \mathcal{N}$ and $\mathcal{M} \models T$ and $\mathcal{N} \models T$ then $\mathcal{M} \prec_s \mathcal{N}$.
Then T has quantifier elimination.

The proof is very long and we need sublemmas. However the sublemmas are all proved under the same assumptions as the theorem and use notation which can only be defined using things appearing in the theorem, so it's probably best to formally begin the proof now. Of course the strategy is to use Lemma 22. The mystery (to me, at this point) is how you're going to use the assumptions of the theorem to get rid of an $\exists$. I can't make any more sensible comments because I haven't internalised the proof yet.

Proof. (of Theorem). We use Lemma 22. So let $\phi(\vec{u}, x)$ is an open formula with $n+1$ free variables. Adjoin n new constants $c_1, c_2, \ldots, c_n$ to $\mathcal{L}$, to get a language $\mathcal{L}'$. Here's an elementary observation about $\mathcal{L}$ and $\mathcal{L}'$.

Sublemma 25. *Let U be an $\mathcal{L}$-theory. (a) If U has quantifier elimination as an $\mathcal{L}$-theory then U has quantifier elimination as an $\mathcal{L}'$-theory.*

(b) If U has algebraically prime models as an $\mathcal{L}$-theory then U has algebraically prime models as an $\mathcal{L}'$-theory.

Note: I made this up. I hope it's OK. Perhaps the converse of the lemma is true too.

Proof. (a) Say ϕ' is is an $\mathcal{L}'$-formula. Replace each occurrence of a constant c_i with a variable x_i, to get an $\mathcal{L}$-formula ϕ. By QE for U, we know that there's an open ψ with $V(\phi) = V(\psi)$ and $U \models_{\mathcal{L}} \phi \iff \psi$. Let ψ' be $\psi(\vec{c})$. Then $V(\phi') = V(\psi')$ and ψ' is open. The claim is that $U \models_{\mathcal{L}'} \phi' \iff \psi'$. To check this, let $\mathcal{M}$ be an $\mathcal{L}'$-structure with $\mathcal{M} \models_{\mathcal{L}'} U$. Then $\mathcal{M} \models_{\mathcal{L}} U$ and hence $\mathcal{M} \models_{\mathcal{L}} \phi \iff \psi$. Hence $\mathcal{M} \models_{\mathcal{L}} \phi(\vec{c}) \iff \psi(\vec{c})$. Hence $\mathcal{M} \models_{\mathcal{L}'} \phi' \iff \psi'$, which is what we wanted.

(b) Let U' denote U regarded as an $\mathcal{L}'$-theory. If ϕ is an open $\mathcal{L}$-formula with $U \models_{\mathcal{L}} \phi$ then ϕ can also be regarded as an open $\mathcal{L}'$-formula, and I claim that $U' \models_{\mathcal{L}'} \phi$, and this is because any $\mathcal{L}'$-structure which is a model for U' is naturally a $\mathcal{L}$-structure and is naturally a model for U, so it models ϕ in $\mathcal{L}$, so it models ϕ in $\mathcal{L}'$. Hence $U_\forall \subseteq U'_\forall$.

Now say $\mathcal{M}'$ is a model for the $\mathcal{L}'$-theory $U'_\forall$. Let $\mathcal{M}$ denote the obvious $\mathcal{L}$-structure associated to $\mathcal{M}'$. Then clearly $\mathcal{M} \models_{\mathcal{L}} U_\forall$. Because U has algebraically prime models, there's an $\mathcal{L}$-model $\mathcal{N}$ of U that contains $\mathcal{M}$ and is "universal" in the weak sense of the definition of algebraically prime. Now N contains M which is an $\mathcal{L}'$-structure, so $\mathcal{N}$ can naturally be made into an $\mathcal{L}'$-structure $\mathcal{N}'$ (use the same constants as in M). My claim is that this model is an algebraically prime extension of $\mathcal{M}'$. For if $\mathcal{N}''$ is any extension of $\mathcal{M}'$ which is an $\mathcal{L}'$-structure and a model for U', it's a model for U, so there's a map $\mathcal{N}' \to \mathcal{N}''$ of $\mathcal{L}$-structures making the obvious diagram commute. I claim that this is a map of $\mathcal{L}'$-structures, and this is clear because the constants are in $\mathcal{M}'$. ☐

Let Δ be the set of all variable-free formulas in $\mathcal{L}'$: that is, no free variables and no bound variables either! Note that if $\alpha \in \Delta$ then α is closed (that is $\alpha = \overline{\alpha}$) and open (no quantifiers).

Sublemma 26. *If $\mathcal{L}'$ and ϕ are as above, and if T satisfies the assumptions of the theorem, and if $\mathcal{M}$ and $\mathcal{N}$ are two $\mathcal{L}'$-structures with $\mathcal{M} \models T$ and $\mathcal{N} \models T$, and if furthermore*

(i) $\mathcal{M} \models \exists x \phi(\vec{c})$, and

(ii) For all $\psi \in \Delta$, if $\mathcal{M} \models \psi$ then $\mathcal{N} \models \psi$.

Then $\mathcal{N} \models \exists x \phi(\vec{c})$.

Note that if T satisfies the assumptions of the theorem and if the theorem is correct, then T will have quantifier elimination as an $\mathcal{L}$-theory and hence as an $\mathcal{L}'$-theory. Hence there's an open $\mathcal{L}'$-sentence ψ with $T \models_{\mathcal{L}'} (\exists x \phi(\vec{c})) \iff \psi$. Now ψ has no bound variables and no free variables, so it has no variables at all, so it's in Δ, and $\mathcal{M} \models_{\mathcal{L}'} \psi$, so by the assumptions of the lemma we have $\mathcal{N} \models \psi$, and hence $\mathcal{N} \models \exists x \phi(\vec{c})$. In particular the theorem implies the sublemma.

Proof. We are going to use the trick of writing down some kind of structure-theoretic analogue of "the prime subfield of a field" (using the existence of algebraically prime models). First we do the natural thing "without axioms", which is purely formal. Indeed, one checks easily that for a given $\mathcal{L}'$-structure, the intersection of all the $\mathcal{L}'$-substructures is a $\mathcal{L}'$-structure, the smallest $\mathcal{L}'$-substructure of the original structure. Let $\mathcal{M}'$ and $\mathcal{N}'$ denote the smallest sub-$\mathcal{L}'$-structures of $\mathcal{M}$ and $\mathcal{N}$ (of course, if $\mathcal{M} \models T$ there's no reason to suspect that $\mathcal{M}' \models T$). Now assumption (ii) of the lemma is enough to deduce that $\mathcal{M}'$ and $\mathcal{N}'$ are isomorphic as $\mathcal{L}'$-structures. This is because M' can be thought of as the set of "interpreted terms in M" in the obvious sense: an element of M' is of the form $t^{\mathcal{M}}$ with t an $\mathcal{L}'$-term with no variables. So the obvious way to attempt to define an isomorphism $\mathcal{M}' \to \mathcal{N}'$ is to send $t^{\mathcal{M}}$ to $t^{\mathcal{N}}$. One needs to check this is a well-defined bijection and an isomorphism of $\mathcal{L}'$-structures. But this is easy. For example to check i is well-defined, note that $t_1^{\mathcal{M}} = t_2^{\mathcal{M}}$ implies $\mathcal{M} \models (t_1 = t_2)$ and $t_1 = t_2$ is in Δ, so (ii) implies $\mathcal{N} \models (t_1 = t_2)$, so $t_1^{\mathcal{N}} = t_2^{\mathcal{N}}$. The fact that the map is an injection is proved similarly, as is the fact that it's an embedding of structures and that it's a surjection.

Now by assumption T has algebraically prime models (as a $\mathcal{L}$-theory and hence as a $\mathcal{L}'$-theory). Moreover, $\mathcal{M} \models T$ and hence $\mathcal{M}'$ models $T_\forall$. So let $\mathcal{M}''$ be an algebraically prime extension of $\mathcal{M}'$. Because $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{M} \models T$, we deduce, by definition of algebraically prime extension, that we may assume $\mathcal{M}'' \subseteq \mathcal{M}$. Also by definition, the isomorphism $\mathcal{M}' = \mathcal{N}'$ means that there's an $\mathcal{L}'$-embedding $\mathcal{M}'' \to \mathcal{N}$. Define $\mathcal{N}''$ to be the image of $\mathcal{M}''$ under this embedding. Now $\mathcal{N}'' \models T$ because $\mathcal{N}'' \cong \mathcal{M}''$ and $\mathcal{M}'' \models T$ by definition of an algebraically prime extension. Now we can prove the lemma: By assumption $\mathcal{M} \models \exists x\phi(\vec{c})$. Now one of our assumptions about T is that sub-T-models of T-models are simply closed (as $\mathcal{L}$-structures). Hence $\mathcal{M}'' \prec_s \mathcal{M}$ (as $\mathcal{L}$-structures). So by definition of simply closed, $\mathcal{M}'' \models_{\mathcal{L}} \exists x\phi(\vec{c})$ (because ϕ is open by assumption). Hence $\mathcal{N}'' \models \exists x\phi(\vec{c})$ (because $\mathcal{M}''$ and $\mathcal{N}''$ are isomorphic). Hence $\mathcal{N} \models \exists x\phi(\vec{c})$ (use the x in N''!). □

Recall that we seek an open $\psi(\vec{u})$ such that $T \models (\exists x\phi) \iff \psi$. Let's introduce even more notation: recall $\mathcal{L}'$ was this new language with n extra constants and Δ was the variable-free formulas in $\mathcal{L}'$. Let Δ' denote $\{\delta \in \Delta | T \cup \{\exists x\phi(\vec{c}, x)\} \models \delta\}$. In words, Δ' is the variable-free formulas we can deduce from T and the assumption $\exists x\phi(\vec{c})$. It's visibly true that $T \cup \{\exists x\phi(\vec{c}, x)\} \models \Delta'$. Note also that $\exists x\phi(\vec{c}, x)$ isn't in Δ' because it isn't in Δ. But the funny thing is

Sublemma 27. $T \cup \Delta' \models \exists x\phi(\vec{c}, x)$.

In words, we're saying that we can deduce the existence statement from T and a bunch of statements which have no variables. The proof basically deduces it from the previous sublemma.

Proof. If $T \cup \Delta'$ has no models then the lemma is vacuously true. So let's assume it has models. Say $\mathcal{M} \models T \cup \Delta'$. Let Δ'' denote $\{\delta \in \Delta | \mathcal{M} \models \delta\}$. The claim is that $T \cup \{\exists x\phi(\vec{c})\} \cup \Delta''$ is also satisfiable. In fact this is clear by compactness: if it weren't true then there would be $\delta_1, \delta_2, \ldots, \delta_m \in \Delta''$ with $T \cup \{\exists x\phi(\vec{c})\} \models \neg(\delta_1 \wedge \ldots \wedge \delta_m)$. But then if $\delta := \neg(\delta_1 \wedge \delta_2 \wedge \ldots \wedge \delta_m)$ then $T \cup \{\exists x\phi(\vec{c})\} \models \delta$ and $\delta \in \Delta$. So, by definition, $\delta \in \Delta'$. Now $\mathcal{M} \models \Delta'$ and hence $\mathcal{M} \models \delta$. But $\mathcal{M} \models \delta_i$ for all i too, and this is a contradiction. So the claim above is true and we can take an $\mathcal{L}'$-model $\mathcal{N}$ for $T \cup \{\exists x\phi(\vec{c})\} \cup \Delta''$. Now for any $\delta \in \Delta$, either $\delta \in \Delta''$ or $\neg\delta \in \Delta''$ by definition of Δ''. If $\mathcal{N} \models \delta$ then, because $\mathcal{N} \models \Delta''$, we must have $\delta \in \Delta''$. And hence $\mathcal{M} \models \delta$. We've just shows that $\mathcal{N} \models \delta$ implies $\mathcal{M} \models \delta$. So by the previous lemma, applied with $\mathcal{M}$ and $\mathcal{N}$ the other way around, we may conclude that $\mathcal{M} \models \exists x\phi(\vec{c})$. □

The nifty thing about this lemma is that it seems to say that if T is true, then $\exists x\phi(\vec{c}, x)$ is equivalent to a possibly infinite bunch of variable-free formulas. We can now finally prove the theorem; it's going to follow from compactness. We have $\phi(\vec{u}, x)$ open and need to construct $\psi(\vec{u})$ open with $T \models (\exists x\phi) \iff \psi$. By the previous sublemma and compactness we may find $\delta_1, \delta_2, \ldots, \delta_m$ with $T \cup \{\delta_1, \delta_2, \ldots, \delta_m\} \models \exists x\phi(\vec{c})$. Now set $\delta = \delta_1 \wedge \delta_2 \wedge \ldots \wedge \delta_m$. Then $T \cup \{\delta\} \models \exists x\phi(\vec{c})$. Now all the δ_i are in Δ', and hence (by definition of Δ') $T \cup \{\exists x\phi(\vec{c})\} \models \delta_i$ for all i and hence $T \cup \{\exists x\phi(\vec{c})\} \models \delta$. We deduce that $T \models_{\mathcal{L}'} (\exists x\phi(\vec{c}) \iff \delta)$. Recall that δ is a variable-free formula in $\mathcal{L}'$ and in particular has no quantifiers. So if δ' is the $\mathcal{L}$-formula obtained by replacing all the c_i with variables u_i then δ' is an open $\mathcal{L}$-formula. Of course, to prove the theorem, it suffices to show that $T \models_{\mathcal{L}} \exists x\phi \iff \delta'$. But this is clear: given an $\mathcal{L}$-model $\mathcal{M}$ for T and $\vec{a} \in M^n$ we make $\mathcal{M}$ an $\mathcal{L}'$-structure in the obvious way and now everything is clear. □

12 Logical equivalence.

In this rather short and easy section there's a super result which enables us to "simplify" an open formula, by finding one in a particularly nice form that's "equivalent" to it.

Let $\mathcal{L}$ be a language and let $\mathcal{M}$ be an $\mathcal{L}$-structure. An *admissible equivalence relation* on M is an equivalence relation that respects the structure: that is, if f is a function of n variables and $\vec{a}, \vec{b} \in M^n$ and $a_i \equiv b_i$ for all i, then $f(\vec{a}) \equiv f(\vec{b})$, and similarly for relations ($R(\vec{a})$ iff $R(\vec{b})$). This is enough to ensure that $M/\equiv$ (the equivalence classes) has a unique natural $\mathcal{L}$-structure, called the *quotient* $\mathcal{L}$-structure. Natural examples: quotient groups, quotient rings and so on. But here comes a really fun example!

Let $\mathcal{L}$ be a language, and say ϕ is a formula. We say that ϕ is a *logical truth*, and write $\models \phi$, if $\mathcal{M} \models \phi$ for every $\mathcal{L}$-structure $\mathcal{M}$. If ϕ and ψ are formulae we say that they are *logically equivalent* if $\models (\phi \iff \psi)$. For example $x = x$ is a logical truth.

Let $\mathcal{W}^0$ denote the set of open formulas for $\mathcal{L}$. Then $\mathcal{W}^0$ is a structure for the language of Boolean algebras (but not a Boolean algebra!). The language of Boolean algebras is two constants 0 and 1, and three operators $\wedge$, $\vee$ and $\neg$, so we let $\mathcal{W}^0$ be a structure by letting 0 be $\neg(x = x)$ and and letting 1 be $(x = x)$. Note that this is not a Boolean algebra: for example $\neg\neg\phi$ is two more characters longer than ϕ! But of course this is easily fixed:

Proposition 28. *Logical equivalence is an admissible equivalence relation on $\mathcal{W}^0$ in the language of Boolean algebras, and the quotient structure is a boolean algebra.*

Proof. Obvious, basically. □

Motivated by the "structure" of Boolean algebras, we make the following definitions. We say that the *conjunction* of a bunch of formulas $\phi_1, \phi_2, \ldots, \phi_n$ is $\phi_1 \wedge \phi_2 \wedge \ldots \phi_n$. We say that the *disjunction* of them is $\phi_1 \vee \phi_2 \vee \ldots \vee \phi_n$. Finally, we say that an open formula ϕ is in *conjunctive normal form* if it's the conjunction of a bunch of open formulae each of which is the disjunction of a bunch of formulae each of which is either an atomic formula or its negation. Similarly an open formula is in *disjunctive normal form* if it's the disjunction of a bunch of conjunctions of a bunch of atomic formulae or their negations.

Theorem 29. *Every open formula ϕ is logically equivalent to an open formula ψ with $V(\psi) \subseteq V(\phi)$ and such that ψ is in conjunctive normal form.*

Proof. If ϕ is open then take the atomic formulas that appear in ϕ and consider the subalgebra of the Boolean algebra of open formulae modulo logical equivalence generated by these atomic formulas. Now use standard results about Boolean algebras. □

13 More on $\mathcal{M} \prec_s \mathcal{N}$, and a practical test for quantifier elimination.

I want to put together some results we've already proved in order to give a practical test to show that a theory has QE! We already have Theorem 24. What are the hurdles in applying it? The two hurdles are obviously checking the two assumptions of the theorem. In some cases we can get algebraically prime models rather easily (for example I already explained how to get them for *ACF* and *DAG*), so it's the "simply closed" thing we need to deal with. Here's a useful criterion for checking simply-closedness, which uses the apparently innocuous previous section crucially.

Proposition 30. *If $\mathcal{M} \subseteq \mathcal{N}$ are $\mathcal{L}$-structures, and if for every $\phi \in \mathcal{W}^0$ which is the conjunction of things which are either atomic formulas or their negations, and for every $x \in V(\phi) = \{x, u_1, u_2, \ldots, u_n\}$, and for every $a \in M^n$, we have $\mathcal{M} \models \exists x \phi(x, a)$ if $\mathcal{N} \models \exists x \phi(x, a)$. Then $\mathcal{M} \prec_s \mathcal{N}$.* {30}

Proof. If $\psi \in \mathcal{W}^0$ is an arbitrary open formula with $n+1$ free variables including x, and if $a \in M^n$, then (by definition) we have to check that $\mathcal{N} \models \exists x \psi(x, a)$ implies $\mathcal{M} \models \exists x \psi(x, a)$. To verify this, choose ϕ in disjunctive normal form which is logically equivalent to ψ. Write $\phi = \phi_1 \vee \phi_2 \vee \ldots \vee \phi_m$ with the ϕ_i all conjunctions. Say $\mathcal{N} \models \exists x \psi(\vec{a})$. We know $\models \phi \iff \psi$. It's not hard to deduce that $\models \exists x \phi \iff \exists x \psi$. So $\mathcal{N} \models \exists x \phi(\vec{a})$. So there's an i such that $\mathcal{N} \models \exists x \phi_i(\vec{a})$. So by assumption $\mathcal{M} \models \exists x \phi_i(\vec{a})$. So $\mathcal{M} \models \exists x \phi(\vec{a})$. So $\mathcal{M} \models \exists x \psi(\vec{a})$ which is what we wanted. □

We now put Proposition 30 and Theorem 24 together, to get a criterion for quantifier elimination which is frequently applicable in practice.

Corollary 31. *If $\mathcal{L}$ is a language with a constant, and if* {31}

(i) T has algebraically prime models, and

(ii) If $\mathcal{M} \subseteq \mathcal{N}$ are $\mathcal{L}$-structures and both models for T, and if for every open formula ϕ which is the conjunction of things which are either atomic or the negation of atomic, with $V(\phi) = \{x, u_1, u_2, \ldots, u_n\}$, and for every $a \in M^n$, we have $\mathcal{M} \models \exists x \phi(x, a)$ if $\mathcal{N} \models \exists x \phi(x, a)$,

then T has QE.

This is just Proposition 30 and Theorem 24.

As I say, this result is actually "usable" in practice. Sometimes (i) is hard to check but sometimes it comes for free. Our strategy in practice for checking that a given theory satisfies (ii) is very "concrete": for the relatively simple theories we'll see in the next section we will in some sense be able to completely explicitly write down the sets $\{x \in N : \phi(x, a)\}$ for $a \in M^n$ as above, and observe that if such sets are non-empty then their intersections with M will also be non-empty. We'll see concrete examples of this in the next section. Note however that it's crucial that we're only using one quantifier here! If $k \subseteq K$ are two algebraically closed fields then the subsets of K defined by polynomial equations with variables in k are easily checked to be non-empty in K iff they have k-points. But the corresponding statement for k^n and K^n is the Nullstellensatz! The fact that we only have one variable in the corollary is very useful!

14 Examples of complete theories and theories with QE.

I was bigging up quantifier elimination in the last section; I was saying that it "almost implied completeness", and completeness is a Good Thing. But in some sense, what I've just said (QE being close to completeness) is nonsense! Given a theory T there's something called a *conservative extension* of T, which is, vaguely speaking, defined by taking all the "functions which are implied by the axioms" and actually turning them into functions in a new language and then constructing a new theory T' in this beefed-up language such that models of T and T' basically coincide (we'll see an explicit example of this in a second, involving discrete total orders). It turns out that T' will then have quantifier elimination! So in fact, in some sense, quantifier elimination is very weak because any theory at all is in some sense equivalent to one where QE holds. Moreover, I think Ambrus also said that if T' is complete then T is complete.

But now let me try and explain why one might hope to use the notion of QE to prove completeness sometimes. To show a theory is complete, we need to check that every sentence is either true in every model or false in every model. But if a theory has QE then a sentence will be equivalent to a formula with no variables at all. Now if one is lucky one can in some sense "write down" all such formulae, and observe that each one is either visibly true in every model or visibly false in every model; this happens for example in cases when you can isolate some kind of minimal substructure of any model, and that the minimal substructures of all models are isomorphic; this happens for example in ACF_0 and ACF_p. We'd already proved that these theories were complete using Vaught's test (the theories are $\aleph_1$-categorical), but right now we'll see an example of a theory where Vaught's test provably doesn't apply, but which we can prove is complete via a QE argument.

Let $\mathcal{L}$ be the language $\{<\}$ of orderings and let T be the theory of total orderings which have a bottom, no top, and are discrete (that is, every element has a unique successor, and every element that isn't the bottom has a unique predecessor). Now Vaught's test won't tell us anything about completeness of T, because T is not κ-categorical for any infinite κ (consider the positive integers, plus κ copies of $\mathbf{Z}$, and that's one model, and if you bung an extra $\mathbf{Z}$ on top then that's another model of the same cardinality but it's not isomorphic to the first model, as one can see by considering the quotient orders defined by the equivalence relation $x\tilde{y}$ iff $x = y + n$ for some
$n \in \mathbf{Z}$ [this can easily be made to make sense: the ordering is discrete].[3] However, 3

Theorem 32. *T is complete.*

---

[3]Does this sound right?

We won't start the proof yet. We need to introduce this trick (basically, conservative extensions) first. We are going to use the theory of quantifier elimination to prove the theorem. Unfortunately, T doesn't have quantifier elimination! The standard example $\mathbf{Z}_{\geq 1} \subseteq \mathbf{Z}_{\geq 0}$ shows that substructures might not be simply closed, and this implies (as we'll see later) that T can't have quantifier elimination (we'll see in a few pages that QE implies model completeness, and T isn't model complete).

We fix this by basically finding a theory "equivalent" to T which does have QE. Let $\mathcal{L}'$ denote the language obtained from $\mathcal{L}$ by adding a constant 0 (bottom) and two unary functions S and P (successor and predecessor). Now let T' denote the $\mathcal{L}'$-theory obtained by taking T and then adding in the three axioms "0 is the bottom", "S is the successor function" and "P is the predecessor function, except $P(0) = 0$". Now any $\mathcal{L}$-structure which is a model for T naturally becomes an $\mathcal{L}'$-structure which is a model for T', and conversely any $\mathcal{L}'$-structure which is a model for T' is naturally an $\mathcal{L}$-structure which is a model for T; the two theories are completely equivalent, in some sense. But given a model, the sub-$\mathcal{L}$-structures don't coincide with the sub-$\mathcal{L}'$-structures: indeed $\mathbf{Z}_{\geq 1}$ isn't a sub-$\mathcal{L}'$-structure of $\mathbf{Z}_{\geq 0}$ because it doesn't contain the "bottom" constant. Now completeness of T' will imply completeness of T, because the two theories are sufficiently "equivalent" to make this assertion trivial. And the trick is

Proposition 33. *T' has quantifier elimination.*

Proof. We use Corollary 31. Usually the hard thing to verify is the existence of algebraically prime models. But in this case this is easy! Indeed, the cute observation is that $T'_\forall \models T'$! Why is this? Well, let's look at what T' is. It is firstly the theory of total orders, so $x < y \wedge y < z \implies x < z$, and $x < y \vee y < x \vee x = y$ and furthermore "only one of these is true", that is, $\neg(x < y \wedge y < x)$ and $\neg(x < y \wedge x = y)$. Now each of these are axioms for total orders, and they're open formulae, so their closures are in $T'_\forall$. Next are the axioms of T which said that there was a bottom and no top and a successor and a predecessor-apart-from-when-you're-at-the-bottom. Now these look rather troublesome on the face of it. As an example, the "there is a bottom" axiom looks like $\exists x \forall y(x = y \vee x < y)$ and the "there is a successor" axiom looks like $\forall x \exists y((x < y) \wedge (\forall z(z < x) \vee (z = x) \vee (z = y) \vee y < z))$. It doesn't look like these axioms are in $T'_\forall$. But let's ignore them for the minute and do the last axioms. The last axioms say that 0 is the bottom, S is the successor function, and P is the predecessor. So the first is $\forall x(0 = x) \vee (0 < x)$, which is going to be in $T_\forall$, the second is that $\forall x \forall z x < S(x) \wedge ((z < x) \vee (z = x) \vee z = S(x) \vee S(x) < z)$, which is in $T'_\forall$, and the last is $\forall x \forall z(x = 0 \wedge P(x) = 0) \vee ((P(x) < x) \wedge (z < P(x) \vee z = P(x) \vee z = x \vee x < z))$. So this is in $T'_\forall$ too. And the messy axioms are logical consequences of these (because they assert the existence of things and P and S and 0 give these things). So any model of $T'_\forall$ satisfies all the axioms of T' (because we just checked each one) and hence $T'_\forall \models T'$. As a consequence we see that T' has algebraically prime models for free!

So to apply Corollary 31 all we have to do is to check the second assumption. We basically do it by "explicitly" writing down what $\{x \in N : \phi(x, a)\}$ can look like, for $a \in M^n$. So say $\mathcal{M} \subseteq \mathcal{N}$ are both models for T', and $\phi = \phi_1 \wedge \phi_2 \wedge \ldots \wedge \phi_m$ is an open formula as in that corollary (so the ϕ_i are either atomic or the negation of an atomic), and with $V(\phi) = \{u_1, u_2, \ldots, u_n, x\}$. We have our model $\mathcal{N}$, and it will contain the substructure $N_0 := \{0, S(0), S(S(0)), \ldots\}$. For $c \in N$ define the "closed half-lines" $I_+(c) := \{b \in N : b \geq c\}$ and $I_-(c) := \{b \in N : b \leq c\}$. Say $\vec{a} \in M^n$.

Sublemma 34. *Say $\vec{a} \in M^n$. If $J := \{b \in N : \mathcal{N} \models \phi(\vec{a}, b)\}$ then there are finitely many $c_i \in M$ such that $J = \bigcap_i I_\pm(c_i)$ (the signs chosen appropriately), up to a finite error (that is, the two sets might not be exactly equal but there's only a finite number of elements of N in one but not in the other), and this finite error is completely contained within N_0.*

Proof. WLOG ϕ is atomic (by our strong assumptions on ϕ earlier). If ϕ is of the form $t_1 = t_2$ for two terms then we can replace ϕ by the form $\neg(t_1 < t_2) \wedge \neg(t_2 < t_1)$ so we can even assume that ϕ is of the form $t_1 < t_2$ for t_1 and t_2 terms. Now $\vec{a} \in M^n$ and we need to consider $J = \{b \in N : t_1(\vec{a}, b) < t_2(\vec{a}, b)\}$, with t_1 and t_2 built only from S and P and 0. But S and P are unary, and 0 is nullary, so in fact t_1 and t_2 can only be of the form $S(P(S(S(\ldots(S(a))\ldots))))$ or

$P(P(S(\ldots(P(b))))\ldots)$, or $S(P(S(S(P(\ldots,P(P(S(0)))))))\ldots)$ and so on. Now anything of the form $S(P(S(\ldots P(a)))\ldots)$ is still in M, so we may as well just replace it with some element of M. Similarly anything of the form $S(P(S(S(P(\ldots(P(0))))))\ldots)$ is in M too. Finally anything of the form $P(P(S(\ldots(P(b))))\ldots)$, well, you have to be a bit careful here, because for example $S(S((P(P(P((b)))))))$ is usually $b-1$, but if $b \leq 2$ then it's just 2. This is where the noise comes in! One checks that the general term of this form is going to be of the form $\max\{b+n,m\}$ with $n \in \mathbf{Z}$ and $m \in \mathbf{Z}_{\geq 0}$. So we have to check four cases:

(i) $a_1 < a_2$: either empty or the whole thing.

(ii) $a_1 < \max\{b+n,m\}$. For $b+n > m$, which happens for b sufficiently large, this is just $b+n > a_1$ so it's $b > \max\{a_1 - n, 0\}$ and the right hand side is in M,

(iii) $a > \max\{b+n,m\}$. Again for b sufficiently large this is just $b < a-n$

(iv) $\max\{b+n_1,m_1\} < \max\{b+n_2,m_2\}$ and up to finite noise this is either everything or nothing. □

We now finish the proof that T' has quantifier elimination. Recall that we're using Corollary 31. Let ϕ be as in the statement of Corollary 31(ii) and say $\mathcal{N} \models \exists x \phi(\vec{a})$. Then there exists $b \in N$ with $\mathcal{N} \models \phi(\vec{a}, b)$. We want to find $b' \in M$ such that $\phi(\vec{a}, b')$. We've seen that up to a finite amount of noise the set $X := \{b \in N : \phi(a,b)\}$ is an interval with endpoints in M. If X contains elements of N_0 then we're home, because $N_0 \subseteq M$. If it doesn't but it's non-empty then by our structure theorem it must have a bottom element and this bottom element will be in M, so again we're home. We deduce that $\mathcal{M} \models \exists x \phi(\vec{a})$. Hence by Corollary 31 T' has QE. □

Note that ACF is not complete: the statement $1+1=0$ is true in some algebraically closed fields but not in others. On the other hand ACF has quantifier elimination—we'll see this later on. So it's not true that QE implies completeness in general. Indeed, as I mentioned, any theory has what is known as a conservative extension, which satisfies QE.

In fact, in the example above, the (conservative) extension T' of T has QE and we'll now use this to show that T' is complete.

Lemma 35. *T' is complete.*

Proof. Given a sentence in $\mathcal{L}'$ we need to check that it is either true in every model, or false in every model. But by QE, the sentence will be equivalent (under T') to an open formula with no variables and no quantifiers. Hence it suffices to show that every atomic formula without variables is either true in every model or false in every model. But an atomic formula without variables is either of the form $t_1 = t_2$ or $t_1 < t_2$ for t_1 and t_2 terms, and all terms will evaluate to one of $0, S(0), S(S(0)), \ldots$. Hence $\mathcal{M} \models \phi$ iff $\mathbf{Z}_{\geq 0} \models \phi$, where $\mathbf{Z}_{\geq 0}$ is the minimal substructure of $\mathcal{M}$! □

Ambrus says that essentially every reasonable countable complete theory has a *unique* minimal substructure.

Finally we can prove that T is complete (recall this is the first theorem of this section: the thing we've been aiming for).

Proof. Any $\mathcal{L}$-sentence is an $\mathcal{L}'$-sentence, so has a universal truth value for any model of T'. But any model of T is a model of T'![4] □ 4

We now turn to non-zero torsion-free divisible abelian groups (the theory we're calling DAG). We've already shown that the theory is κ-categorical for any uncountable κ. Hence the theory is complete. But in fact we'll also show

Theorem 36. *DAG has QE.*

[4] check this with Ambrus.

Proof. We use Corollary 31 again. An easy algebraic argument (c.f. $ACF_\forall$ being the theory of integral domains) shows that if we let the language of groups be $\mathcal{L} := \{0, +\}$, that is, we don't throw in inverses, then the theory $DAG_\forall$ is just the theory of torsion-free monoids with cancellation. Hence ACF has algebraically prime models. So it suffices to check condtion (ii) of Corollary 31. So let ϕ be a conjunction of things which are either atomic or the negation of atomics, if $V(\phi) = \{x, u_1, \ldots, u_n\}$. if $a \in M^n$ and if $X = \{b \in N : \mathcal{N} \models \phi(a, b)\}$, then we need to check that $X \neq \emptyset$ implies $X \cap M \neq \emptyset$. We do this, as in the case of orderings, by explicitly identifying what the possibilities are for X. Well, $\phi = \phi_1 \wedge \phi_2 \wedge \ldots \wedge \phi_m$, and each $\phi_i(x, a)$ is an evaluation of an atomic formula in the language of groups, so it can only be of the form $nx = c$ for some $c \in M$ and $n \in \mathbf{Z}$ or $nx \neq c$ for some $c \in M$ and $n \in \mathbf{Z}$. Any equation of the form $0.x \neq 0$ or $0.x = c$ for some $c \neq 0$ will imply that X is empty, and any equation of the form $0.x = 0$ or $0.x \neq c$ for some non-zero c will not change X at all, so we may assume that either X is everything (in which case we're done), or empty (in which case we're done), or is defined by finitely many formulae of the form $n.x = c$ and $n.x \neq c$ with $n \neq 0$. Now the equation $nx = c$ has a unique solution when $n \neq 0$, and hence if one of these equations is mentioned then either X is empty or has one element, of the form c/n, which will be in M because $c \in M$. Finally if all the equations are of the form $nx \neq c$ with $n \neq 0$ then X will be the complement of a finite set in N, and hence will intersect M non-trivially because M is a model of DAG and is hence infinite. So we're done! □

There was a typo in Ambrus' notes on proof because Ambrus was preparing the lecture in a CASLAT talk.

Corollary 37. *DAG is complete!*

Note that we already saw a proof of this, because DAG is $\aleph_1$-categorical. But this is a totally different proof now.

Proof. We need to check that every sentence in the language of groups is either true in every non-zero torsion-free divisible abelian group, or false in all of them. By QE for DAG, DAG shows that a given sentence is equivalent to a formula with no quantifiers and no free variables, and hence with no variables at all. Now what could such a formula look like? Well, the only variable-free terms in the language of groups are 0 and $0 + 0$ and $(0 + 0) + 0$ and $0 + (0 + 0)$ and so on, and in the theory of groups all of these evaluate to zero. So the only atomic formulae all look like things like $(0 + 0) + 0 = 0 + (0 + (0 + 0))$, and all of these will be true, and so every formula is either going to be equivalent in the theory DAG to either the statement $0 = 0$ or the statement $0 \neq 0$, and the first is always true and the second is always false! □

Here's a key concept that we have already implicitly used. Let $\mathcal{M}$ be an $\mathcal{L}$-structure. We say that a set $X \subseteq M^n$ is *definable* if there's a formula ϕ with $V(\phi) = \{u_1, u_2, \ldots, u_n, v_1, v_2, \ldots, v_m\}$ and $b \in M^m$ such that $X = \{a \in M^n : \mathcal{M} \models \phi(a, b)\}$. In this case we say that $\phi(u, b)$ *defines* x.

Noddy examples: any finite subset of M^n is definable, because if the finite set is $\{b_1, b_2, \ldots, b_r\}$ then ϕ can be $u = b_1 \vee u = b_2 \vee \ldots \vee u = b_r$, where $u = b_1$ is shorthand for $(u_1 = b_{11} \wedge u_2 = b_{12} \wedge \ldots \wedge u_n = b_{1n})$. Similarly any cofinite subset is definable, because you can just stick a $\neg$ in front of the previous example. In fact, while we're at it, we may as well observe the obvious

Lemma 38. *The definable subsets of M^n form a Boolean subalgebra of the power set of M^n.*

Proof. The whole of M^n is defined by the formula $u_1 = u_1 \wedge u_2 = u_2 \wedge \ldots \wedge u_n = u_n$. If X is defined by ϕ and Y by ψ then $X \cap Y$ is defined by $\phi \wedge \psi$, and so on, and it's trivial. □

Caveat 1: if you know the definable subsets of M you might not know much about the definable subsets of M^2. For example, I suspect that in the language of rings, the definable subsets of $\mathbf{C}$, the complexes, will be just the subsets which are either finite or cofinite. But the definable subsets of $\mathbf{C}^2$ will contain loads of complicated affine curves, like $y^7 = x^9 + 1$ and so on.

Caveat 2: in DAG, what are the definable subsets of $\mathbf{Q}^2$? This is an ill-defined question! If $M = \mathbf{Q}$ then the line $\{(x,x) : x \in \mathbf{Q}\}$ is going to be a definable subset of $M^2 = \mathbf{Q}^2$. But if $M = \mathbf{Q}^2$ then we'll see below that any definable subset of $M = \mathbf{Q}^2$ is either finite or cofinite.

Fun examples: in the language of rings, the set $\mathbf{R}_{\geq 0}$ is a definable subset of $\mathbf{R}$, because it's defined by the formula $\phi(x) = (\exists y : y.y = x)$. This is easy. But can one define $\mathbf{Q}_{\geq 0}$ in $\mathbf{Q}$? Yes, but one needs a theorem! $\phi(x) = (\exists a \exists b \exists c \exists d x = a.a + b.b + c.c + d.d)$ will work, but only because every positive integer is the sum of four squares. Finally, is $\mathbf{Z}$ a definable subset of $\mathbf{Q}$ in the language of rings? The answer is yes but this is a theorem of Julia Robinson! So in fact the definable subsets of $\mathbf{Q}$ as a ring are rather subtle. But in a second we'll see that the definable subsets of $\mathbf{Q}$ as a DAG are all either finite or cofinite.

Note that the concept of definability can go much further, but we won't take it much further. For example a function is definable if its graph is definable, and one can check that if f and X are definable then so is $f(X)$ and $f^{-1}(X)$, and so on. If I knew enough algebraic geometry to know what a constructible set was then I would now be able to prove a theorem of Chevalley, I think Ambrus said.

We've seen that any finite or cofinite subset of M is definable. We'll now show that in DAG these are all the definable sets. In fact we'll do better: we'll even classify the definable subsets of M^n. If M is a DAG then say that a *hyperplane* in M^n is a set of the form $\{a \in M^n : \sum_i r_i a_i = b\}$ for some fixed $r \in \mathbf{Z}^n$ and $b \in M$. For example, the hyperplanes in M are M ($r = b = 0$), the empty set $r = 0$ and $b \neq 0$) and the 1-element sets ($r = 1$ and b arbitrary), and the hyperplanes in $\mathbf{Q}^2$ contain things like $x + 2y = 7$.

Proposition 39. *The definable subsets of M^n are the Boolean algebra generated by the hyperplanes.*

Note the following trivial consequence:

Corollary 40. *The definable subsets of M are precisely the sets which are finite or cofinite.*

Proof. (of proposition). Say X is defined by $\phi(u,b)$. Now DAG has QE, so there's ψ an open formula with the same free variables as ϕ with $DAG \models \phi \iff \psi$. In particular X is defined by $\psi(u,b)$ too. Now an open formula is logically equivalent to one in conjunctive normal form, so X is in the Boolean algebra generated by the subsets of M^n defined by atomic formulae. But DAG is a theory in the language of groups, so the only atomic formulae are of the form $t_1 = t_2$ with the t_i terms, and the only terms look like $b_1 + u_3 + b_2 = u_7 + u_7 + b_3$ and so on, which in the theory of DAG will be equivalent to terms defining hyperplanes. Conversely every hyperplane is clearly definable so we're done. $\square$

So $\mathbf{Q}$ as a group is easy but $\mathbf{Q}$ as a ring is very very complicated. In fact one corollary of Gödel's theorem is that there is no Turing machine which can tell you whether an arbitrary subset of $\mathbf{Q}$ (as a ring) is definable or not.

We'll now prove that ACF has QE (and note that ACF isn't complete, because the statement $1 + 1 = 0$ is sometimes true and sometimes not), and our proof will as ever be an application of Corollary 31.

Theorem 41. *ACF has QE.*

Proof. Note first that $\mathcal{M} \models ACF_\forall$ iff $\mathcal{M}$ is an integral domain, so ACF has algebraically prime models (take the algebraic closure of the field of fractions). So as ever we just have to check condition (ii) of Corollary 31. Say we have two algbraically closed fields $M \subseteq N$, and an open formula ϕ which is the conjunction of things which are either atomic or the negation of atomic, and $V(\phi) = \{x, u_1, u_2, \ldots, u_n\}$, and $a \in M^n$. We need to check that if X is $\{x \in N : \phi(x,a)\}$ then X is non-empty iff $X \cap M$ is non-empty. So now we need to think a little about what the atomic formulas look like in the language of rings, and we see that (after substituting the elements a_i of M for the variables u_i) they are (equivalent to formulas) of the form $P(x) = Q(x)$ with P and Q in $M[x]$. So we are reduced to checking the following. Say P_i and Q_i are in $M[x]$ for $1 \leq i \leq r + s$, and say X is $\{x \in N : P_i(x) = Q_i(x), 1 \le i \le r\} \cap \{x \in N : P_i(x) \ne Q_i(x), r+1 \le i \le r+s\}$. Then X is non-empty iff $X \cap M$ is non-empty. Well, the equalities are either always true (and can hence be dropped), always false (in which case we're done), or true for a non-zero finite number of elements of N, all of which are in M, in which case we're done. So we may as well assume that there are no equalities at all. And the inequalities are also either always true (so drop them), always false (so we're done), or true away from a finite set, so again we're done because M is infinite. □

Corollary 42. *ACF_0 and, for p a prime, ACF_p, are complete.*

Note that we proved this already: the theories are $\aleph_1$-categorical so are complete by Vaught's test.

Proof. If $\mathcal{M} \models ACF_p$ then the smallest substructure of $\mathcal{M}$ (in the language of rings) is $\mathbf{Z}/p\mathbf{Z}$. Similarly if $\mathcal{M} \models ACF_0$ then the smallest substructure is $\mathbf{Z}$. Now let v be either 0 or p. Then $ACF_v \supseteq ACF$ and hence ACF_v has QE (this follows trivially from the definition of QE; given a formula ϕ just write down an open ψ with $ACF \models (\phi \iff \psi)$; then ACF_v models it too). So we just have to check that every sentence with no variables at all has an "absolute truth value" independent of every model. It suffices to check for atomic formulae, so it suffices to check that terms are either always equal or always unequal, and the interpretation of these terms will always be in $\mathbf{Z}$ or $\mathbf{Z}/p\mathbf{Z}$ so at the end of the day the result is clear. □

15 Model completeness.

An $\mathcal{L}$-theory T is *model complete* if for every $\mathcal{M} \subseteq \mathcal{N}$ (models of T) we have $\mathcal{M} \prec \mathcal{N}$. Non-examples: the theory of rings ($\mathbf{Z} \subseteq \mathbf{Q}$) (because $existsx : x+x = 1$ causes trouble) and the theory of orders ($\mathbf{Z}_{\ge 1} \subseteq \mathbf{Z}_{\ge 0}$) (because $\exists x : x < 1$ causes trouble; note that we're allowed to refer to 1 even though it's not a constant). On the positive side we have

Proposition 43. *If T has QE then T is model complete.*

Proof. Let ϕ be a formula with n free variables $x_1, x_2, \ldots, x_n$ and say $a \in M^n$. We need to check that $\mathcal{N} \models \phi(a)$ iff $\mathcal{M} \models \phi(a)$. Well, by QE we can find an open $\psi(x_1, \ldots, x_n)$ with $T \models \phi \iff \psi$. In particular $T \models \phi(a) \iff \psi(a)$. So in fact it's now easy isn't it. First, $\mathcal{N} \models \phi(a)$ iff $\mathcal{N} \models \psi(a)$. But the interpretation (i.e. the truth value) of $\psi(a)$ in $\mathcal{M}$ and $\mathcal{N}$ coincides, because ψ is open, and the interpretation of terms coincides and so on and so on ("the truth of open formulas is preserved by extensions and substructures"). So $\mathcal{N} \models \psi(a)$ iff $\mathcal{M} \models \psi(a)$. And because $\mathcal{M}$ is a model of T this is true iff $\mathcal{M} \models \phi(a)$. And that's what we wanted. □

Hence ACF_0 and ACF_p and DAG are model complete. In particular this means that in the language of groups $\mathbf{Q} \subseteq \mathbf{Q}^2$ is an elementary embedding (whereas $\mathbf{Z} \subseteq \mathbf{Z}^2$ isn't, because one can write down a sentence saying "there exists t such that every x is either twice something, or twice something plus t").

Corollary 44. *The theory of discrete linear orders with a bottom but no top does not have quantifier elimination!*

Proof. $\mathbf{Z}_{\ge 1} \subseteq \mathbf{Z}_{\ge 0}$ is not an elementary substructure, becasue if $\phi(y)$ is $\exists x : x < y$ then $\phi(1)$ is true in $\mathbf{Z}_{\ge 0}$ but and makes sense in $\mathbf{Z}_{\ge 1}$ but is not true in $\mathbf{Z}_{\ge 1}$. □

Here's an awesome application! The Nullstellensatz!

Theorem 45. *Let K be an algebraically closed field, and let I be a proper ideal of $K[x_1, x_2, \ldots, x_n]$. Then there's a point a in K^n such that $f(a) = 0$ for all $f \in I$.*

Proof. WLOG I is maximal (Zorn). Let L be the residue field and let $\overline{L}$ be an algebraic closure of L. Choose a finite set of generators $I = (f_1, \ldots, f_r)$ for I. Then $\overline{L} \models \exists \vec{x}(f_1(\vec{x}) = 0 \wedge \ldots \wedge f_r(\vec{x}) = 0)$. But by model completeness we know $K \prec \overline{L}$. So K models the above sentence too. And that's it! ☐

I *think* that Ambrus said that one of Hilberts problems (the one about whether every $f \in \mathbf{R}[X_1, X_2, \ldots, X_n]$ which maps $\mathbf{R}^n$ to the non-negative reals is automatically a sum of finitely many squares of polynomials) can be answered affirmatively this way using model theory.[5] 5

16 Types.

A *type* is just a set Σ of formulas ϕ with $V(\phi) = \{x_1, x_2, \ldots, x_n\}$ (that is, each ϕ has exactly the same free variables). Let $\mathcal{M}$ be an $\mathcal{L}$-structure. We say that $a \in M^n$ *realises* Σ if $\mathcal{M} \models \sigma(a)$ for all $\sigma \in \Sigma$. We also say that $\mathcal{M}$ realises Σ. We say that $\mathcal{M}$ *omits* Σ if it doesn't realise Σ.

We say that Σ as above is a *complete type* if (a) it can be realised in some $\mathcal{L}$-structure, and (b) it's maximal with respect to sets with this property. I guess maximality will be equivalent to the condition that for every formula ϕ with $V(\phi) = \{x_1, x_2, \ldots, x_n\}$, either $\phi \in \Sigma$ or $\neg\phi \in \Sigma$. So it seems to me that an equivalent definition would be that a type Σ is a complete type if there's an $\mathcal{L}$-structure $\mathcal{M}$ and $a \in M^n$ such that Σ is the set of ϕ with $V(\phi) = \{x_1, x_2, \ldots, x_n\}$ with $\mathcal{M} \models \phi(a)$. We say that Σ formed in this way is the *type of* a. We say that a type is an n-type if all of the formulae in it have n free variables.

Example: in the language $\{0, +\}$ of groups, the statements $\neg(x = 0)$, $\neg(x+x = 0)$, $\neg(x+x+x = 0)$ and so on, just say x is not torsion. Any element realising this type will be a non-torsion element.

Example: if we beef up the language of rings by adding in constants for every rational, so $\mathcal{L} = \{+, -, *\} \cup \mathbf{Q}$ then Ambrus says that the statements $\{f(x) \neq 0\}$ as f runs through all the non-constant elements of $\mathbf{Q}[x]$ gives us an (incomplete) type and all transcendentals in $\mathbf{C}$ satisfy it. Because there are automorphisms of $\mathbf{C}$ that send an arbitrary transcendental to another one, all the types associated to transcendentals are the same. In particular there are only countably many types realised by $\mathbf{C}$.

Example: in the language of ordered rings $\{+, -, *, 0, 1, <\}$ the reals with its usual structure have the following property: if $x < y$ then there's a rational r with $x < r < y$ and if $r = a/b$ and $b > 0$ then $bx < a$ and $a < by$, so the types of x and y are distinct. In particular there are uncountably many types realised by $\mathbf{R}$.

Proposition 46. *Let T be a theory, and let Σ be a type in variables $x_1, x_2, \ldots, x_n$ Then TFAE:*
(i) T has a model which realises Σ
(ii) Every finite subset of Σ is realised in some model of T
(iii) if X is the set of all sentences of the form $\exists x_1 \exists x_2 \ldots \exists x_n(\phi_1 \wedge \phi_2 \wedge \ldots \wedge \phi_m)$ for $m \in \mathbf{Z}_{\geq 1}$ and the ϕ_i running through all the m-element subsets of Σ, then $T \cup X$ is satisfiable.

This is easy. (i) iff (ii) is, I think, just compactness applied to some beefed-up language. I think this is the hardest part of the proposition. (i) implies (iii) is trivial, because if $a \in M$ realises Σ then a is a witness to the truth of all the elements of X. And (iii) also trivially implies (ii) because the model for $T \cup X$ will realise any finite subset of Σ.

In short, the compactness theorem is a machine which enables us to make models which realise types. But making models which omit types is harder!

Definition: Let Σ be a type in $(x_1, x_2, \ldots, x_n)$. A theory T *locally omits* Σ if there is no formula ϕ with $V(\phi) \subseteq \{x_1, x_2, \ldots, x_n\}$ such that (i) $T \cup \{\exists x \phi\}$ (here x is a vector) is satisfiable, and (ii) $T \models \phi \implies \psi$ for every $\psi \in \Sigma$.

Theorem 47 (Omitting types theorem)**.** *Let T be a satisfiable theory in a countable language, and let Σ be a type in $(x_1, x_2, \ldots, x_n)$. If T locally omits Σ then T has a model which omits Σ.*

[5]check!

Before we start the proof, here's a (counter-)example to show that the assumption of a countable language is necessary. If $\mathcal{L}$ is the language $\{c_i : i \in \omega\} \cup \{d_j : j \in \omega_1\}$ (constants) then $\mathcal{L}$ is uncountable, and if T is the $\mathcal{L}$-theory which says that the interpretations of d_j are pairwise distinct, that is, the sentences $\neg(d_i = d_j)$ for all $i < j \in \omega_1$, then $\mathcal{M} \models T$ implies that $|M| \geq \aleph_1$. So the 1-type Σ whose elements are just $\{x \neq c_i : i \in \omega\}$ is never omitted by models of T: any model will have an element which realises Σ. But T locally omits this type! Here's why. Firstly, observe that the standard argument (Corollary 31) shows that T has QE: indeed $T_\forall \models T$ so T has algebraically prime models, the terms of $\mathcal{L}$ are just constants and variables, the atomic formulas are all of the form $x = x$ or $x = y$ or $x = c$ or $c = c'$ for c and c' constants, so if $\mathcal{M} \subseteq \mathcal{N}$ and X is of the form $\{x \in N : \phi(x, a)\}$ for ϕ as in Corollary 31 then X is either empty, an element of M, a constant, or the whole of N minus a finite set, and in every case it's non-empty iff it has non-empty intersection with M. So T has QE. Now let's show that T locally omits Σ, by contradiction. Say it didn't. Then choose $\phi(x)$ with $V(\phi) \subseteq \{x\}$ with $T \cup \{\exists x\phi(x)\}$ satisfiable. Next choose an open ψ with $V(\phi) = V(\psi)$ and $T \models \phi \iff \psi$. Then any model of $T \cup \{\exists x\phi(x)\}$ will also model $\exists x\psi(x)$, so $T \cup \{\exists x\psi\}$ is also satisfiable. Moreover, if $T \models \phi \implies \sigma$ for every $\sigma \in \Sigma$ then $T \models \psi \implies \sigma$ too. But Σ mentions infinitely many constants (all the c_i) and the claim is that ψ as above cannot model the statement $x \neq c_i$ under T if ψ doesn't mention c_i, because an open formula is logically equivalent to one in conjunctive normal form so is in the Boolean algebra generated by the constants and if it only mentions finitely many constants then it's in the Boolean algebra generated by these constants!

Before we embark on the proof of the theorem, let's make some remarks about it. If $\mathcal{L}$ is a countable language and T is an $\mathcal{L}$-theory with a model $\mathcal{N}$ that omits Σ, then it has a countable model $\mathcal{M}$ which omits Σ. For either N is countable (or finite), in which case we're done, or N is uncountable, in which case we can choose a countable elementary substructure $\mathcal{M}$ of $\mathcal{N}$ by the going down theorem and the claim is that this still omits Σ. Indeed, if $\mathcal{M}$ realises Σ then choose $a \in M$ realising Σ; then $\mathcal{M} \models \phi(a)$ for all $\phi \in \Sigma$ and hence (by definition of elementary embedding) $\mathcal{N} \models \phi(a)$ too.

The final remark before we prove the theorem is that if T is a *complete* and satisfiable theory, and if Σ is a type, and if T has a model which omits Σ, then T locally omits Σ. Indeed, if T does not locally omit Σ then there's some formula ϕ with $T \cup \{\exists x\phi\}$ satisfiable and $T \models \phi \implies \sigma$ for all $\sigma \in \Sigma$. Now $T \cup \{\exists x\phi\}$ satisfiable means there's some model of T where $\exists x\phi$ is true (where here x is a vector), and completeness of T means that $T \models \exists x\phi(x)$. So if $T \models \phi \implies \sigma$ then for any model of T we just choose some a with $\phi(a)$ true and then $\sigma(a)$ will also be true, so all models of T realise Σ.

OK so finally let's prove the theorem. We will be lazy and just do it for 1-types.

Proof. Let $\mathcal{L}'$ denote $\mathcal{L}$ plus new constants $\{c_i : i \in \omega\}$. Then $\mathcal{L}'$ is still countable. Let's assume that our stock of variables is countably infinite, but furthermore that there are countably infinitely many variables in our stock that aren't mentioned at all in T. Then certainly the set of all sentences of $\mathcal{L}'$ that don't mention any of the bound variables in T is also countable! So list them as $\{\phi_i : i \in \omega\}$. We have our satisfiable theory T, which we can view as a satisfiable $\mathcal{L}'$-theory, and we're going to recursively build an increasing sequence of satisfiable $\mathcal{L}'$-theories $T \subseteq T_1 \subseteq T_2 \subseteq \ldots$, with each T_i formed by taking T_{i-1} and adding one new sentence. Let's say T_{i-1} is T plus $i-1$ new sentences. Let θ be the conjunction of these new sentences. Now θ might of course mention some of the c_j. Let $\psi(x, y)$ denote the formula which you get by changing c_i (same i as in T_i) to x and changing all the other c_j to new variables y_j. Recall we're assuming that Σ is a 1-type. Let $\Sigma(c_i)$ denote the set of sentences $\{\sigma(c_i) : \sigma \in \Sigma\}$. By assumption, T locally omits Σ. The claim is that $T \cup \{\theta\} \not\models \Sigma(c_i)$, and we check this by contradiction: by assumption $T \cup \{\theta\}$ is satisfiable, so $T \cup \{\exists x\exists y\psi(x, y)\}$ (with x a variable and y a vector) is satisfiable. Moreover, we're assuming for a contradiction that $T \cup \{\theta\} \models \sigma(c_i)$ for all $\sigma \in \Sigma$, and hence $T \models \theta \implies \sigma(c_i)$ for all $\sigma \in \Sigma$. It follows that $T \models_{\mathcal{L}} \exists y\psi \implies \sigma$, because for any $\mathcal{L}$-model $\mathcal{M}$ of T we want to check that for all $a \in M$ we have $\exists y\psi(a, y)$ implies $\sigma(a)$, and we do this by noting that whenever $\exists y\psi(a, y)$ is true we can make $\mathcal{M}$ an $\mathcal{L}'$-structure in the obvious way (let c_i be a and let all the other c_j be the y's) and then noting that θ has become true in $\mathcal{M}$, and hence $\mathcal{M} \models_{\mathcal{L}'} \sigma(c_i)$, so

$\mathcal{M} \models_{\mathcal{L}} \sigma(a)$. But this is a contradiction! Because $\exists y \psi(x, y)$ shows that T doesn't locally omit Σ.

So take a model of $T \cup \{\theta\}$. We've just seen that there must be some σ, say $\sigma = \sigma_i \in \Sigma$, such that $\sigma_i(c_i)$ is false in the model. Hence $T \cup \{\theta, \neg\sigma_i(c_i)\}$ is $\mathcal{L}'$-satisfiable. We are finally ready to define θ_i! There are three cases. First, we take a model of $T \cup \{\theta, \neg\sigma_i(c_i)\}$. Now we ask whether ϕ_i is true or not in the model. If ϕ_i is false then we set $\theta_i = \neg\sigma_i(c_i) \wedge \neg\phi_i$. If ϕ_i is true then there are two cases. If ϕ_i is *not* precisely of the form $\exists t \xi_i(t)$ then we let θ_i be $\neg\sigma_i(c_i) \wedge \phi_i$. Note that in both of these cases the resulting theory $T_i = T_{i-1} \cup \{\theta_i\}$ is trivially satisfiable. The 3rd case is the most interesting. Say ϕ_i is true and furthermore that ϕ_i happens to be of the form $\exists t \xi_i(t)$, with $V(\xi_i) = \{t\}$, then choose $p \in \mathbf{Z}_{\geq 0}$ the smallest index such that c_p is not mentioned in any of θ, $\sigma(c_i)$ or ϕ_i, and let θ_i be the sentence $\neg\sigma_i(c_i) \wedge \phi_i \wedge \xi_i(c_p)$. Again it's clear that the resulting set T_i will be satisfiable. So by compactness $T_\omega := \bigcup_i T_i$ will be satisfiable. Moreover, for every $\mathcal{L}'$-sentence ρ, either $T_\omega \models \rho$ or $T_\omega \models \neg\rho$, because ρ is logically equivalent (after possible renaming of bound variables) to one of the ϕ_i. In particular T_ω is complete. Let $\mathcal{M}$ be a model.

The final claim is that the subset of M consisting of the interpretations of the c_i is actually the underlying set of an elementary substructure of $\mathcal{M}$. Well, first let's check that it's a substructure (that is, closed under constants and functions). For the functions we note that if f is a function then we want to check that $f(c_{i_1}, c_{i_2}, \ldots, c_{i_r})$ interprets to one of our named constants. But $\mathcal{M} \models \exists t(t = f(c_{i_1}, c_{i_2}, \ldots))$ and this latter statement is one of our ϕ_i, so there will be some p such that c_p interprets to $f(c_{i_1}, \ldots)$ because we will exactly be in that funny 3rd case. Similarly for constants: $\exists t(t = c)$, for c one of the original constants in $\mathcal{L}$, will be one of the ϕ_i.

So the set of interpretations of constants is a substructure $\mathcal{M}_0$ of $\mathcal{M}$. To check it's an elementary substructure we use Tarski-Vaught. So assume $\phi(u, x)$ is a formula with $n+1$ free variables (vector u and variable x). Say $a \in M_0^n$ and say that there's $c \in M$ with $\mathcal{M} \models \phi(a, c)$. We need to find $b \in M_0$ with $\mathcal{M} \models \phi(a, b)$. Because $a \in M_0^n$ it's just a bunch of constants! Let ϕ' be ϕ with a replaced by the relevant constants. Then $\mathcal{M} \models \exists x \phi'(x)$. But T_ω is complete, so $T_\omega \models \exists x \phi'(x)$. Moreover that latter sentence appeared in our list! So again there's some c_p such that $\phi'(c_p)$ is true in $\mathcal{M}$. And that is what we needed. □

Finally, here's an extension of the previous theorem. If $\mathcal{L}$ is a countable language and T is a satisfiable theory which omits countably many types $\Sigma_1, \Sigma_2, \ldots$. Then T has a countable model which omits all of the Σ_i. The proof is "the same" but it's messier.

16.1 Appendix: the point of types?

In a verbal conversation with Ambrus he pointed out the type $\{x > 1, x > 1+1, x > 1+1+1, \ldots\}$ in the language of totally ordered groups, say. This type is certainly realised, and moreover any finite subset of it is realised in $\mathbf{Z}$, but it itself isn't realised in $\mathbf{Z}$, but it will be realised in an ultraproduct of $\mathbf{Z}$'s. Conversely, if Σ is a type and it's realisable in an elementary extension of a model $\mathcal{M}$, then any finite subset of Σ will be realised in a model of the full theory of $\mathcal{M}$, and the full theory of $\mathcal{M}$ is complete! So the finite subset of Σ can be re-interpreted as a sentence $(\exists x \sigma_1 \wedge \sigma_2 \wedge \ldots)$ which is true in an elementary extension of $\mathcal{M}$ and is hence true in $\mathcal{M}$ itself. On the other hand Σ might not be realised in $\mathcal{M}$. This is a bit vague, but I'm not really on top of types at the minute.

17 Saturated models.

Say $\mathcal{M}$ is an $\mathcal{L}$-structure and $A \subseteq M$. Let $\mathcal{L}_A$ denote the language $\mathcal{L}$ plus a constant c_a for each $a \in A$. Then $\mathcal{M}$ is naturally an $\mathcal{L}_A$-structure. Let $S_A^n(\mathcal{M})$ denote the set of n-types for the language $\mathcal{L}_A$ which happen to be realised by a model of the full theory of $\mathcal{M}$ in $\mathcal{L}_A$. I'll remark that if Σ is such an n-type, then any finite subset of Σ will be, I think, realisable in $\mathcal{M}$ itself, because the full theory of $\mathcal{M}$ is a complete theory, so any finite subset can be recast as a sentence $\exists x(\sigma_1(x) \wedge \sigma_2(x) \wedge \ldots)$ which, if it's true in an extension of $\mathcal{M}$ must be a consequence of the full theory of $\mathcal{M}$, as this theory is complete! Hence the sentence is true in $\mathcal{M}$ and hence any finite subset of Σ is actually realised by $\mathcal{M}$. But a complete type might of course not be realised by $\mathcal{M}$ itself.

Let $S_A(\mathcal{M}) = \cup_{n\geq 1} S^n_A(\mathcal{M})$. Let κ be a cardinal. Say $\mathcal{M}$ is *κ-saturated* if for all $A \subseteq M$ with $|A| < \kappa$, the $\mathcal{L}_A$-structure $\mathcal{M}$ itself realises every type in $S_A(\mathcal{M})$. Say $\mathcal{M}$ is *saturated* if it's $|M|$-saturated.

I think we're going to be trying to construct saturated models.

Lemma 48. *(i) $\prec$ is a partial order on models (so $\mathcal{M} \prec \mathcal{M}$ and $\mathcal{A} \prec \mathcal{B} \prec \mathcal{C}$ implies $\mathcal{A} \prec \mathcal{C}$).*
(ii) If $\mathcal{A} \prec \mathcal{C}$ and $\mathcal{B} \prec \mathcal{C}$ and $A \subseteq B$ then $cA \prec \mathcal{B}$.

Proof. Trivial unravelling of definitions. □

Now let α be an ordinal and let $\{\mathcal{M}_\beta : \beta < \alpha\}$ be a collection of $\mathcal{L}$-structures with $\beta < \gamma \implies M_\beta \subseteq M_\gamma$. Let $\mathcal{M}$ denote the union of the $\mathcal{M}_\beta$; this is also naturally an $\mathcal{L}$-structure.

We say that the $\mathcal{M}_\beta$ form an *elementary chain* if $\mathcal{M}_\beta \prec \mathcal{M}_\gamma$ for all $\beta < \gamma < \alpha$.

Lemma 49. *If $\{\mathcal{M}_\beta : \beta < \alpha\}$ is an elementary chain with union $\mathcal{M}$ then $\mathcal{M}_\beta \prec \mathcal{M}$ for all β.*

The lemma follows immediately from

Lemma 50. *For every $\mathcal{L}$-formula ϕ, the following is true: for every $\beta < \alpha$ and for every $a \in M^n_\beta$ (where n is the number of free variables of ϕ), we have $\mathcal{M}_\beta \models \phi(a) \iff \mathcal{M} \models \phi(a)$.*

Note that the point is that we'll use formula induction to prove the result for all β at once; this is crucial in the argument.

Proof. An easy formula induction. Trivial for atomic formulae, and $\neg$ and $\wedge, \vee$. The only time when we have to use the assumption of the lemma is when ϕ is of the form $\exists x\psi$, and we know $\mathcal{M} \models \phi(a)$ and we want to deduce $\mathcal{M}_\beta \models \phi(a)$ (the other way of this is also easy). Here's how this unique non-formal part goes. We know $\mathcal{M} \models \phi(a)$, so there's $b \in M$ with $\mathcal{M} \models \psi(a, b)$. Now choose $\gamma \geq \beta$ with $b \in M_\gamma$. We know $\mathcal{M} \models \psi(a, b)$ so by our inductive hypothesis applied to ψ, we deduce $\mathcal{M}_\gamma \models \psi(a, b)$. So certainly $\mathcal{M}_\gamma \models \exists x\psi(a, x)$, which is $\phi(a)$. But $\mathcal{M}_\beta \prec \mathcal{M}_\gamma$ and hence $\mathcal{M}_\gamma \models \phi(a)$, which is what we wanted. □

We now need, for some reason, a long technical lemma.

Lemma 51. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure. Let κ be a cardinal with $\aleph_0 \leq |M| \leq 2^\kappa$ and $|\mathcal{L}| \leq \kappa$. Then $\mathcal{M}$ has an elementary extension $\mathcal{N}$ with $|N| = 2^\kappa$ and with the following amazing property: for any $A \subseteq M$ with $|A| \leq \kappa$, $\mathcal{N}$ (interpreted as an $\mathcal{L}_A$-structure in the obvious way) realises every type in $S_A(\mathcal{M})$!*

Recall $S_A(\mathcal{M})$ is the types realised by some model of the full theory of $\mathcal{M}$ in $\mathcal{L}_A$. So I guess that in some sense this isn't so amazing: you can envisage $\mathcal{N}$ as being some huge ultraproduct. This wouldn't be an elementary extension, but you can probably fix this by some trick of enlarging the language.

In fact before we start let's do some counting. A type in $S_A(\mathcal{M})$ is a bunch of formulae in $\mathcal{L}_A$, and $|A| \leq \kappa$ and $|\mathcal{L}| \leq \kappa$, and the assumptions imply κ is infinite, so $|\mathcal{L}_A| \leq \kappa$, so the number of formulae in $\mathcal{L}_A$ is at most κ, so the size of $S_A(\mathcal{M})$ is at most 2^κ. Furthermore the number of subsets A of M of size at most κ is at most $\kappa.|M|^\kappa \leq \kappa.(2^\kappa)^\kappa \leq 2^\kappa$, so in fact the total number of types mentioned in the lemma is at most 2^κ. So in fact this lemma isn't surprising at all!

Proof. The trick, as ever, is to adjoin a new variable for every Σ in every $S_A(\mathcal{M})$. So let the language $\mathcal{L}^*$ denote $\mathcal{L}_M$ ($\mathcal{L}$ and a constant for every $m \in M$) plus a constant vector d_Σ of length n for every $A \subseteq M$ of size at most κ and every $n \geq 1$ and every $\Sigma \in S^n_A(\mathcal{M})$ (where by a constant vector of length n I just mean n constants). The point is that d_Σ will be a witness to the realisation of Σ.

Now consider the full theory of $\mathcal{M}$ as an $\mathcal{L}_M$-structure (that is, the elementary diagram of $\mathcal{M}$). Consider this theory as an $\mathcal{L}^*$-theory. Now let T be the following $\mathcal{L}^*$-theory: firstly, take the full theory of $\mathcal{M}$ as an $\mathcal{L}_M$-structure. And now also add the sentences $\sigma(d_\Sigma)$ for Σ running through all of $S_A(\mathcal{M})$ (and A running through all subsets of M of size at most κ), and σ running through Σ.

The claim is that T is satisfiable. If we can prove this claim then we're done, because if $\mathcal{N}$ is a model of T then WLOG it has cardinality 2^κ (by going down and then going up), and $\mathcal{M} \subseteq \mathcal{N}$ (because of the constants) and in fact $\mathcal{M} \prec \mathcal{N}$ (because T contains the elementary diagram of M), and for any Σ the interpretation of the constant vector d_Σ realises Σ.

So it suffices to prove that T is satisfiable, and by compactness it suffices to prove that all finite subsets of T are satisfiable. So say $S \subseteq T$ is a finite subset. Let's denote all the elements of S not in the elementary diagram as $\phi_1, \phi_2, \ldots, \phi_r$, with $\phi_i = \sigma_i(d_{\Sigma_i})$, where of course some of the Σ_i might coincide for different i. But what's for sure is that there is a *finite* subset $A \subseteq M$ such that all the Σ_i are in $S_A(\mathcal{M})$. Moreover, all the elements of S that *are* in the elemtary diagram of $\mathcal{M}$ can, after enlarging A if necessary, be assumed to be sentences in $\mathcal{L}_A$. So now in fact let's throw in all of the full theory of $\mathcal{M}$ as an $\mathcal{L}_A$-structure into S and prove that the resulting thing is satisfiable. And now here's the trick: we assumed that each Σ_i was realised by some model of the theory of $\mathcal{M}$, so each statement of the form $\exists x \sigma_i(x)$ is realised by some model of the full theory of $\mathcal{M}$, and because the full theory is complete it must be a consequence of the full theory of $\mathcal{M}$, so it's realised by $\mathcal{M}$ itself. If we fix Σ and let I be the finite set of i such that $\Sigma_i = \Sigma$ then $\exists x(\sigma_{i_1}(x) \wedge \sigma_{i_2}(x) \wedge \ldots)$ is also true in $\mathcal{M}$, and we deduce that $\mathcal{M}$ is a model for S! So all finite subsets of T have a model, so T has a model, and we're done. □

So now we use this lemma to prove

Proposition 52. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure and let κ be a cardinal such that $\aleph_0 \leq |\mathcal{M}| \leq 2^\kappa$ and $|\mathcal{L}| \leq \kappa$. Then $\mathcal{M}$ has an elementary extension $\mathcal{N}$ with $|\mathcal{N}| = 2^\kappa$ and such that $\mathcal{N}$ is κ^+-saturated.*

Note that this is stronger than the previous lemma because we have to allow $A \subseteq N$ rather than $A \subseteq M$.

Proof. We'll build an elementary chain $\{\mathcal{M}_\alpha : \alpha < \kappa^+\}$ with $\mathcal{M}_0 = \mathcal{M}$, $|\mathcal{M}_\alpha| = 2^\kappa$ if $\alpha > 0$ and if $A \subseteq M_\alpha$ with $|A| \leq \kappa$ then $\mathcal{M}_{\alpha+1}$ realises all types in $S_A(\mathcal{M}_\alpha)$. If we can do this then we're done: let $\mathcal{N}$ be the union of the $\mathcal{M}_\alpha$; then $|\mathcal{N}| \leq \kappa^+.2^\kappa = 2^\kappa$, and $\mathcal{M} = \mathcal{M}_0 \prec \mathcal{N}$, and if $A \subseteq N$ with $|A| \leq \kappa$ there's some $\alpha < \kappa^+$ with $A \subseteq M_\alpha$ (because the cofinality of κ^+ is κ^+), and any $\Sigma \in S_A(\mathcal{M}_\alpha)$ is realised in $\mathcal{M}_{\alpha+1}$ and hence in $\mathcal{N}$.

We construct the elementary chain by ordinal induction, surprise surprise. At limit ordinals just take the union (the only thing we have to check is that the cardinality doesn't blow up, which is easy to check). At successor ordinals we just use the previous lemma, which is exactly what we need to make the argument work. So we're done! □

Corollary 53. *Assume GCH. If $\kappa \geq \aleph_0$ is a cardinal with $|\mathcal{L}| \leq \kappa$ then any $\mathcal{L}$-theory T which has infinite models has saturated models of cardinality κ^+.*

Proof. Take $\mathcal{M} \models T$ with $|\mathcal{M}| = \kappa$ (by going up). By the previous proposition there's an elementary extension $\mathcal{N}$ of $\mathcal{M}$ with $|N| = 2^\kappa = \kappa^+$, and $\mathcal{N}$ being κ^+-saturated. Hence $\mathcal{N}$ is saturated. Moreover $\mathcal{M} \prec \mathcal{N}$ and hence $\mathcal{M} \equiv \mathcal{N}$, so $\mathcal{N} \models T$. □

Proposition 54. *If $\mathcal{M}$ and $\mathcal{N}$ are two saturated $\mathcal{L}$-structures with $\mathcal{M} \equiv \mathcal{N}$ and $|M| = |N| \geq \aleph_0$, then $\mathcal{M} \cong \mathcal{N}$.*

That's a pretty neat application of the notion of saturated models! The proof is going to be some crazy back-and-forth argument. Just before we start, let's show

Lemma 55. *If $|\mathcal{M}| \geq \aleph_0$ and $\mathcal{M}$ is κ-saturated, then $|\mathcal{M}| \geq \kappa$.*

Proof. By contradiction. Say $|\mathcal{M}| < \kappa$. Consider the 1-type $\{\neg(c = c_a) : a \in M\}$. Now M is infinite, so all finite subsets of Σ are realised by a model of $\mathrm{Th}_{\mathcal{L}_M}(\mathcal{M})$ (namely $\mathcal{M}$), so Σ is realised by some model of $\mathrm{Th}_{\mathcal{L}_M}(\mathcal{M})$, so $\Sigma \in S_M(\mathcal{M})$. By definition of saturated, $\mathcal{M}$ itself realises Σ—but this is clearly false. □

We'll now start the proof of this saturated-models-are-isomorphic proposition.

Proof. Let $\mathcal{M}$ and $\mathcal{N}$ be as in the proposition. Say $A \subseteq M$ and let's say we're given an injective map $f : A \to N$. Then $\mathcal{N}$ naturally becomes a $\mathcal{L}_A$-structure. Let $\mathcal{M}_A$ and $\mathcal{N}_{f(A)}$ denote $\mathcal{M}$ and $\mathcal{N}$ as $\mathcal{L}_A$-structures. Now count $M = \{m_\alpha : \alpha \in \kappa\}$ and $N = \{n_\alpha : \alpha \in \kappa\}$. We'll now use the back and forth idea.

Grr. I never typed this up; the rest of this proof (a standard back and forth) is in 2 sides of photocopied notes in my hard-copy notes on this course (which also contains the example sheets). □

18 Appendix: real fields.

A field is *formally real* if -1 isn't a sum of squares. The theorem is that F is formally real iff it's orderable. Better: if F is formally real and $x \in F$ such that $-x$ isn't a sum of squares, there's an ordering on F such that $x > 0$.

A field F is *real closed* if it's formally real and if the only finite formally real extension of F is F itself. Note that there will surely be massively infinite formally real extensions of F. I think that Ambrus said that a slick definition of real closed is simply that a field is real closed iff it's a structure in the language of fields that models the full theory of the reals!

The theorem is that F is real closed iff (a) every polynomial of odd degree has a root, and (b) for any non-zero $f \in F$, either $+f$ or $-f$ has a square root.

Examples: the reals. The hyperreals (a countable ultraproduct of copies of the reals)—this isn't the reals because the element $(1, 2, 3, 4, 5, \ldots)$ is bigger than every integer. On the other hand if F is formally real and if you adjoin a square root of -1 then you get something algebraically closed! So the hyperreals embed into the complexes, giving an element of $\mathrm{Aut}(\mathbf{C})$ of order 2 which isn't conjugate to complex conjugation within $\mathrm{Aut}(\mathbf{C})$.

A much easier construction: start with $\mathbf{R}$, adjoin a transcendental x and decree that x is bigger than every real. Now take the real closure!