INTRODUCTION TO MODEL THEORY

MOSHE KAMENSKY

1. Intro and review

1.1. **Boolean algebras.** Let $\mathcal{B}$ be a boolean algebra. We let $\mathbb{S}(\mathcal{B})$ be the set of boolean algebra maps $\mathcal{B} \to \mathbf{2}$, where $\mathbf{2} = \{0, 1\}$ is the unique 2-element boolean algebra. Any element $b \in \mathcal{B}$ determines a map $\underline{b} : \mathbb{S}(\mathcal{B}) \to \mathbf{2}$, via $\underline{b}(x) = x(b)$. Note that the set $\mathrm{Hom}(X, \mathbf{2})$, for any set X, is a boolean algebra with pointwise operations. Viewing the elements of this algebra as characteristic functions, this is identified with the power set of X, with the boolean algebra structure given as usual.

Exercise 1.1.1. *The map* $b \mapsto \underline{b}$ *is an injective homomorphism of boolean algebras* $r : \mathcal{B} \to \mathrm{Hom}(\mathbb{S}(\mathcal{B}), \mathbf{2})$

Hence, any boolean algebra can be viewed as a subalgebra of a power set. To describe the image, we use a topological language.

Definition 1.1.2. The *Stone space* of a boolean algebra $\mathcal{B}$ is the set $\mathbb{S}(\mathcal{B})$, with the weak topology making all functions $\underline{b}$ for $b \in \mathcal{B}$ continuous (where $\mathbf{2}$ is discrete). **Stone space**

By definition, each function $\underline{b}$ is continuous, so the image of r lies in the subalgebra $C(\mathbb{S}(\mathcal{B}), \mathbf{2})$ consisting of continuous functions.

Theorem 1.1.3 (Stone representation theorem)**.** *The space* $\mathbb{S}(\mathcal{B})$ *is Hausdorff and compact. The map* $r : \mathcal{B} \to C(\mathbb{S}(\mathcal{B}), \mathbf{2})$ *is an isomorphism from* $\mathcal{B}$ *to the algebra of continuous functions.*

Proof. Compactness is an exercise (this is the compactness theorem for propositional logic). Assume $f : \mathbb{S}(\mathcal{B}) \to \mathbf{2}$ is continuous, and let $\mathcal{B}_f = \{b \in \mathcal{B} | \underline{b} \geqslant f\}$. By compactness, this filter is generated by finitely many elements, hence by one element. □

If $t : \mathcal{B}_1 \to \mathcal{B}_2$ is a homomorphism of boolean algebras, composition determines a continuous map $t^* : \mathbb{S}(\mathcal{B}_2) \to \mathbb{S}(\mathcal{B}_1)$. Conversely, if $u : \mathbb{S}(\mathcal{B}_2) \to \mathbb{S}(\mathcal{B}_1)$ is continuous, composition with u determines a boolean algebra homomorphism from $\mathcal{B}_1 = C(\mathbb{S}(\mathcal{B}_1), \mathbf{2})$ to $\mathcal{B}_2 = C(\mathbb{S}(\mathcal{B}_2), \mathbf{2})$, and one easily checks that these two processes are inverse. In other words, $\mathcal{B} \mapsto \mathbb{S}(\mathcal{B})$ is a fully faithful (contravariant) functor. In particular, if $\mathbb{S}(\mathcal{B}_1)$ is homeomorphic to $\mathbb{S}(\mathcal{B}_2)$ then $\mathcal{B}_1$ is isomorphic to $\mathcal{B}_2$. In checking this condition it is useful to recall the following

Exercise 1.1.4. *If* $t : X \to Y$ *is a continuous bijection, where* X *is compact and* Y *is Hausdorff, then* t *is a homeomorphism*

We thus have the following corollary, which is the main goal of this discussion.

{cor:stonebij}

Corollary 1.1.5. *If* $t : \mathcal{B}_1 \to \mathcal{B}_2$ *is a boolean algebra homomorphism such that* $t^* : \mathbb{S}(\mathcal{B}_2) \to \mathbb{S}(\mathcal{B}_1)$ *is a bijection, then* t *is an isomorphism.*

totally disconnected

To complete the picture, it might be interesting to find the image of the Stone space functor. A space is *totally disconnected* if it has basis of clopen (closed and open) sets. By definition, $\mathbb{S}(\mathcal{B})$ is such.

Exercise 1.1.6. *Let* X *be a totally disconnected space, and let* $\mathcal{B} = C(X, \mathbf{2})$ *be the algebra of continuous functions. There is a natural continuous map* $X \to \mathbb{S}(\mathcal{B})$*, which is injective if* X *is Hausdorff, and is a homeomorphism if* X *is (in addition) compact. Thus, the category of boolean algebras is anti-equivalent (via the Stone space functor) to that of totally disconnected Hausdorff compact spaces.*

End of lecture 1, Mar 14

2. Quantifier elimination and Model companions

We would like to analyze the definable sets in particular theories of interest, for example theories of fields. What makes the collection of definable sets complicated is the occurrence of quantifiers: for example, any quantifier-free definable subset of any field is finite or co-finite, but the set of squares in $\mathbb{Q}$ is already fairly more complicated. On the other hand, the set of squares in $\mathbb{C}$ is also very simple: it is just $\mathbb{C}$. In fact, we will see below that every definable subset of $\mathbb{C}^n$ definable by a boolean combination of polynomial equations. This property is named quantifier elimination:

quantifier elimination

Definition 2.0.1. A theory $\mathcal{T}$ admits *quantifier elimination* (QE) if for any formula $\phi(\bar{x})$ there is a quantifier-free formula $\psi(\bar{x})$ equivalent to it, i.e., such that $\mathcal{T} \models \forall \bar{x}(\phi \iff \psi)$

As a first observation, we note:

{ex:onequant}

Exercise 2.0.2. $\mathcal{T}$ *admits QE iff every formula of the form* $\exists x \phi(x, \bar{y})$*, where* ϕ *is a conjunction of basic formulas, is equivalent a quantifier free one.*

{exa:pureinf}

Example 2.0.3. Let $\mathcal{T}$ be the theory in the signature with one sort (and equality) saying that the sort is infinite. We show that it admit QE. Using the last remark, we need to show that a formula of the form

$$\exists x(x = y_1 \wedge \cdots \wedge x = y_k \wedge x \neq z_1 \wedge \cdots \wedge x \neq z_m) \tag{2.1}$$

If equivalent to a quantifier free one. There are two cases: if $k > 0$, the formula is equivalent to $y_1 = \cdots = y_k \wedge y_1 \neq z_1, \ldots, y_1 \neq z_m$. Otherwise, the formula asserts that there is an element distinct from the given m elements, and this is true since the models are infinite. □

The last example, despite its simplicity, was a bit messy, so we look for better tools for proving QE. We fix a theory $\mathcal{T}$. For any set V, let $\mathcal{F}(V)$ be the set of formulas with free variables in V. Logical equivalence (with respect to $\mathcal{T}$) forms an equivalence relation on $\mathcal{F}(V)$, and we denote by $\mathcal{B}(V)$ the quotient set. We will often blur the difference between $\mathcal{F}(V)$ and $\mathcal{B}(V)$, but $\mathcal{B}(V)$ has the advantage that (under the obvious operations) it is a boolean algebra.

Definition 2.0.4. An element of the Stone space of $\mathcal{B}(V)$ is called a (complete) *type* (in the variables V). The Stone space itself is called the *type space*, denoted $\mathbb{S}(V)$.

type
type space

Again, we identify a type with the set of formulas representing its elements. Unpacking the definitions, a type is a maximal consistent set of formulas (with respect to $\mathcal{T}$). For example, for $V = \emptyset$, this is the space of complete theories extending $\mathcal{T}$. In general, a type can be viewed as describing an "ideal" element. Indeed:

Exercise 2.0.5. *If $\mathbb{M}$ is a model of $\mathcal{T}$ and $A \subseteq M$, the set $\{\phi(\overline{a}) | \overline{a} \in \phi^{\mathbb{M}}\}$ is a complete type in the variables A (here, A serves in the double duty of a set of variables and the assignment in $\mathbb{M}$ to them given by the inclusion). This type is denoted* $\text{tp}_{\mathbb{M}}(A)$.

Conversely, for any type $p \in \mathbb{S}(V)$ there is a model $\mathbb{M}$ of $\mathcal{T}$ containing V, such that $p = \text{tp}_{\mathbb{M}}(V)$.

The set $\mathcal{F}(V)$ contains some interesting subsets: the subset $\mathcal{F}^{qf}(V)$ of quantifier free formulas, the subset $\mathcal{F}^{\exists}(V)$ of formulas of the form $\exists x(\phi(x, \overline{y}))$ with ϕ quantifier free, etc. Each generates a sub-algebra $\mathcal{B}^{qf}(V)$, $\mathcal{B}^{\exists}(V)$ of $\mathcal{B}(V)$, and thus determines type spaces $\mathbb{S}^{qf}(V)$ and $\mathbb{S}^{\exists}(V)$. An element of $\mathbb{S}^{qf}$ is called a *quantifier free type*. So far we assumed nothing about the set V, but all our statements will immediately reduce to the case when V is finite, which we will assume from now on. Combining these remarks with the previous results we get:

quantifier free type

Proposition 2.0.6. *The following are equivalent:*

1. $\mathcal{T}$ *admits QE*
2. *Every quantifier-free type extends uniquely to a complete type*
3. *Every quantifier-free type extends to a unique element in* $\mathbb{S}^{\exists}$

Proof. If every quantifier free type extends to a unique element of $\mathbb{S}^{\exists}$, the induced map on type spaces (coming from the inclusion $\mathcal{B}^{qf} \subseteq \mathcal{B}^{\exists}$) is a bijection. By Cor 1.1.5, it follows that the inclusion $\mathcal{B}^{qf} \subseteq \mathcal{B}$ is an isomorphism. This implies QE by 2.0.2. The other directions are obvious. □

Example 2.0.7. Consider $\mathcal{T}$ from Example 2.0.3. A quantifier-free type essential consists of formulas $y_k \neq y_j$ for some variables y_i. An existential formula $\exists x(\psi_1(x, \overline{y}) \wedge \cdots \wedge \psi_m(x, \overline{y}))$ (with ψ_j basic) can be added to this type precisely when at most one ψ_j is a equation. Hence the extension is unique. □

The last example shows that we gained a slight improvement. A more substantial improvement is achieved through parameters.

2.1. **Adding constants.** If A is a structure for the signature Σ, a structure over A is a structure for Σ along with a homomorphism from A. If $\mathcal{T}$ is a theory in Σ, we denote by Σ_A the signature with additional constant symbols for the elements of A, and by $\mathcal{T}_A$ the extension of $\mathcal{T}$ to the signature Σ_A by all quantifier free sentences in Σ_A that hold in A (i.e., by the "multiplication table" of A). Any model of $\mathcal{T}_A$ can be viewed as a model $\mathbb{M}$ of $\mathcal{T}$, and determines in addition a homomorphism $A \to \mathbb{M}$, and this process is an equivalence between models of $\mathcal{T}_A$ and models of $\mathcal{T}$ over A.

type over A A (complete) *type over* A for a theory $\mathcal{T}$ is, by definition, a type for $\mathcal{T}_A$, and likewise for quantifier free types, formulas, etc. We have

Proposition 2.1.1. *The following are equivalent*

(1) $\mathcal{T}$ *has QE*
(2) *For every structure* A*,* $\mathcal{T}_A$ *is complete (if consistent)*
(3) *For every structure* A *and every quantifier free* $\phi(x, \bar{a})$ *in* Σ_A*, either* $\mathcal{T}_A \models \exists\phi$ *or* $\mathcal{T}_A \models \neg\exists x\phi$ *(i.e.,* $\mathcal{T}_A$ *is complete for existential sentences)*
(4) *For any models* $\mathbb{M}_1$ *and* $\mathbb{M}_2$ *with a common substructure* A*, and any quantifier free* $\phi(x, \bar{a})$*, if* ϕ *is satisfied in* $\mathbb{M}_1$*, then it is satisfied in* $\mathbb{M}_2$*.*

Proof. If $\mathcal{T}$ has QE, every sentence over A is equivalent to a quantifier-free one, hence it or its negation is included in $\mathcal{T}_A$. The equivalence of the last two conditions is just a translation using the description of models of $\mathcal{T}_A$.

Assume the last condition, and let p be a quantifier-free type for $\mathcal{T}$. If p_1 and p_2 are two elements of $\mathbb{S}^{\exists}$ extending it, let a_1 and a_2 be realisations in some models $\mathbb{M}_1$ and $\mathbb{M}_2$. Since a_1 and a_2 realise the same quantifier-free type, the structures generated by them are isomorphic. By the last condition, $p_1 = p_2$. Hence each quantifier-free type has a unique extension in $\mathbb{S}^{\exists}$, so $\mathcal{T}$ admits QE. □

Note that since every formula uses a finite number of elements of A, it suffices to check the condition for finite sets A. Going back to the example, we have a quantifier-free formula $\phi(x, \bar{a})$ realised in an infinite set containing $\bar{a}$. We need to show that it can be realised in any other infinite set containing $\bar{a}$. This is clear. We now consider some more complicated examples.

End of lecture 2, Mar 21

2.2. **Dense linear orders.**

{prp:dlo}

Proposition 2.2.1. *The theory* $\mathcal{DLO}$ *of dense linear orders with no minimum or maximum eliminates quantifiers, and is therefore complete.*

Proof. Let L_1 and L_2 be two such orders, and let A be a common (ordered) subset. Let $\phi(x, \bar{a})$ be a quantifier free formula satisfied by $b \in L_1$. Let A_0 be the set of parameters in $\bar{a}$ which are smaller than b, and A_1 the ones bigger (if $b \in A$ we are done). So $a_0 < a_1$ for all $a_0 \in A_0$ and $a_1 \in A_1$. By density and no endpoints, we may find $b' \in L_2$ with $a_0 < b'$ and $b' < a_1$ for all $a_i \in A_i$. So b' must satisfy $\phi(x, \overline{a})$. □

Exercise 2.2.2. *Show that the theory of dense linear orders with (distinct) minimum and maximum has QE in the language with constants for the minimum and maximum, but not without (and similarly for just one endpoint).*

2.3. **Equivalence relations.** Let $\mathcal{T}$ be the 1-sorted theory with one binary relation E, stating that E is an equivalence relation. $\mathcal{T}$ is not complete: for example, it is consistent with the statement that there are 3 classes of size 7 and with its negation.

Exercise 2.3.1. *With $\mathcal{T}$ as above*

(1) *For every $n \in \mathbb{N}$, the set of elements whose equivalence class has size n is definable.*
(2) *Let $f : \mathbb{N} \to \mathbb{N} \cup \{\infty\}$ be any function. Show that there is a theory $\mathcal{T}_f \supseteq \mathcal{T}$ whose models are the equivalence relations with $f(n)$ classes of size n, for each n.*
(3) *Assume that f is eventually 0 (i.e., $f(n) = 0$ for all but finitely many n). Show that for every $m \in \mathbb{N} \cup \{\infty\}$ there is a theory $\mathcal{T}_{f,m}$ extending $\mathcal{T}_f$, whose models are models of $\mathcal{T}_f$ with m infinite classes. Show that this theory is complete.*
(4) *If f is not eventually 0, can you prove that $\mathcal{T}_f$ is complete?*

The theories above do not eliminate quantifiers: The set union of all classes of size 7 is definable but not (in general) without quantifiers. To fix this, we add to the signature unary predicates P_k ($k \geqslant 1$), and to the theory for eack k the sentence that says that $P_k(x)$ holds if and only if the class of x has size k.

Proposition 2.3.2. *Let $\mathcal{T}$ be any of the theories $\mathcal{T}_{f,m}$ or $\mathcal{T}_f$ (when f is not eventually 0), as above, extended to the signature with predicates P_k. Then $\mathcal{T}$ admits quantifier elimination.*

Proof. Let $\mathbb{M}_1$ and $\mathbb{M}_2$ be two models, with a common subset A (note that any subset is a substructure here), and let $\phi(x, \overline{a})$ be a quantifier-free formula satisfied by an element $m_1 \in \mathbb{M}_1$. The truth value of $\phi(x, \overline{a})$ is determined by the truth of finitely many basic formulas Ψ of the form $x = a$, xEa or $P_k(x)$. If $m_1 = a$ for some $a \in A$, there is nothing to prove. We consider the two other cases:

(1) For some $a \in A$, $m_1 E a$ in $\mathbb{M}_1$. It follows that for all k, $P_k(m_1)$ holds if and only if $P_k(a)$ does. So we are looking for $m_2 \in \mathbb{M}_2$ satisfying $m_2 E a$ there, and also different from a finite number k of elements of A. Since $m_1 E a$ and is distinct from all these elements, we have $\neg P_l(a)$ for all $l \leqslant k$. Hence such an m_2 can be found (note that all formulas of the form xEa' for other $a' \in A$ are determined by the relationship between a and a').

(2) $\neg(m_1 E a)$ for all $a \in A$. We either have $P_k(x) \in \Psi$ for some k, or $\neg P_k(x) \in \Psi$ for finitely many k. Hence, we need to find an element of $\mathbb{M}_2$ whose class is of size k or different from some $k_1, \dots, k_m$, and which is in-equivalent to finitely many $a_1, \dots, a_n$. The fact that we managed to do this in $\mathbb{M}_1$ means that $\mathcal{T}$ allows this.

□

Corollary 2.3.3. *For* f *not eventually* 0*, the theory* $\mathcal{T}_f$ *is complete. Such theories, along with the theories* $\mathcal{T}_{f,m}$ *(as in the exercise) are all the completions of the theory of an equivalence relation.*

Proof. Considering the expansion by the P_i, QE implies that every sentence is equivalent to a quantifier-free one, and these are only the trivial ones. Since the additional predicates are defined also in the original signature, the original theory is also complete. □

Corollary 2.3.4. *For each* $m \in \mathbb{N} \cup \{\infty\}$*, the theory* $\mathcal{T}_{0,m}$ *that says that there are* m *infinite classes and no finite classes is complete, and eliminates quantifiers in the original language.*

Proof. In this case the predicates P_k are empty, and so are definable without quantifiers □

End of lecture 3, Mar 28

2.4. **Modules.** Let R be a ring (associative with unit). We consider the theory $\mathcal{T}_R$ of R-modules, in a signature with unary function $\underline{r}$ for each $r \in R$. In general, this theory does not eliminate quantifiers:

Example 2.4.1. Let $R = \mathbb{Z}$, so that $\mathcal{T}_R$ is the theory of abelian groups. Viewing $\mathbb{Z}$ as such, the set of even elements is definable, but not without quantifiers □

An equation is a formula of the form $T\bar{x} = 0$, where T is a matrix over R. Clearly, every quantifier-free formula is a boolean combination of equations. Further, a (finite) conjunction of equations is an equation. A *positive-primitive (pp)* formula has the form $\exists \bar{y}(T\overline{xy} = 0$. We note that if $\phi(\bar{x}, \bar{y})$ is a pp-formula and M is an R-module, then $\phi(M, 0)$ is a subgroup, and for any $m \in M$, $\phi(M, m)$ is a coset of it, or is empty. If ϕ and ψ are two pp-formulas and n is a natural number, the statement "the index of ψ in ϕ is n" is first-order (the index of a subgroup B of a group A in another subgroup C is the index of $B \cap C$ in C; further, if the index is not finite, we will say it is ∞, i.e., won't distinguish different infinite cardinalities). We denote a sentence stating this by $\theta(\psi, \phi, n)$.

positive-primitive (pp)

{prp:modqe}

Proposition 2.4.2. *Let* R *be a ring, and let* $\mathcal{T}$ *be a theory extending* $\mathcal{T}_R$ *that decides all sentences* $\theta(\psi, \phi, n)$ *(as above) for pp-formulas* ψ, ϕ *in one variable. Then* $\mathcal{T}$ *eliminates quantifiers to the level of pp-formulas.*

Proof. Let $\mathbb{M}_1$ and $\mathbb{M}_2$ be two models, with a common sub-structure A (so A is a submodule). Given pp-formulas $\phi_0(x, \bar{y}), \dots, \phi_k(x, \bar{y})$, a tuple $\bar{a} \in A$

and an element $b \in \mathbb{M}_1$ satisfying $\phi_0(b, \overline{a})$ and $\neg\phi_i(b, \overline{a})$ for $i > 0$, our task is to find a similar element in $\mathbb{M}_2$.

As noted above, each $\phi_i(x, \overline{a})$ defines a coset $b_i + B_i$ of a subgroup B_i of $\mathbb{M}_1$, or the empty set. Which of the two cases it is is determined by the pp-formula $\exists x \phi_i$, so is determined in the substructure A, is thus the same in $\mathbb{M}_2$. Hence, we may assume it is non-empty, and similarly defines a coset $c_i + C_i$ in $\mathbb{M}_2$. Note that B_i is defined by a formula without parameters (which also defines C_i), so the index of B_i in B_0 is determined by $\mathcal{T}$, and thus is equal to the index of C_i in C_0.

So we need to show that $c_0 + C_0$ is not covered by the finitely cosets $c_i + C_i$, for $i > 0$. By Neumann's Lemma, below, we may assume that the index of each C_i in C_0 is finite. By the inclusion-exclusion principle, this depends only on the index of $\cap_{i \in I} C_i$ in C_0, where I runs over all subsets of $\{1, \dots, k\}$. Again, by assumption these numbers are determined. □

To complete the proof, we need to show:

Lemma 2.4.3 (Neumann's Lemma)**.** *No group is a finite union of cosets of infinite index subgroups*

Proof. By induction on the number of subgroups, the base is by definition. Assume that $G = h_1 H_1 \cup \cdots \cup h_n H_n$, where each H_i is of infinite index. If some H_i has finite index in H_n, then we can get rid of H_n and we are done by induction. Otherwise, H_n is covered by finitely many cosets of the infinite index subgroups $H_i \cap H_n$, and again we are done by induction. □

Corollary 2.4.4. *Each theory as in Prop. 2.4.2 is complete. These are precisely the completions of* $\mathcal{T}_R$ *(not necessarily consistent)*

We note that any pp-formula may be written as $\exists \overline{y}(A\overline{x} = B\overline{y})$ for matrices A and B over R. For the purpose of quantifier elimination, we may assume that A is the identity (perhaps in more variables, since we may substitute). In general, not much more simplification may be achieved. However, in special cases the situation is simpler:

Corollary 2.4.5. *Assume that* R *is a (commutative) PID. Then the theory of an* R*-module is determined by the sizes of* M_r/sM *for* $r, s \in R$*, where* M_r *is the* r*-torsion* $M_r = \{m \in M | rm = 0\}$ *and* sM *is the image of* s *on* M*. It eliminates quantifiers in the language enriched by predicates* D_r *for the image of multiplication by* r *(i.e., divisibility by* r*). Furthermore, it suffices to take* r *primary (power of a prime).*

Proof. For any matrix B over a PID there are invertible matrices L and R such that LBR is a diagonal matrix D (use elementary row and column operations; if B is not square, the extra rows or columns are 0). Hence the pp-formula $\exists \overline{y}(\overline{x} = B\overline{y})$, is equivalent to $\exists \overline{y}(L\overline{x} = D\overline{y})$, so it suffices include predicates for the image of multiplication by $d \in R$. Since any such d is a product of primary elements, the primary elements suffice. □

Corollary 2.4.6. (1) *The complete theory of $\mathbb{Z}$ as an abelian group is the extension stating that $\mathbb{Z}$ is torsion-free, and $\mathbb{Z}/n\mathbb{Z}$ has size n.*
(2) *The theory of non-zero vector spaces over an infinite field is complete and eliminates quantifiers*
(3) *More generally, the theory of torsion-free divisible non-zero modules over an infinite PID is complete and eliminates quantifiers.*

The last part extends to more general rings. Considering again a system of equations $B\bar{y} = \bar{a}$, for the equation to have a solution it is clearly necessary that if $\bar{c}B = 0$ for some row vector $\bar{c}$, then $\overline{ca} = 0$. When R is Noetherian (any ideal is finitely generated), a module is called *injective* if these necessary conditions are also sufficient (for general rings, a similar condition with possibly infinite number of variables is required). For a given matrix B, the set of $\bar{c}$ for which $\bar{c}B = 0$ is a submodule of R^k, so is finitely generated. Since it suffices to state the condition for the generators, injectivity is a first order property.

Corollary 2.4.7. *If R is a Noetherian ring, the theory of any injective module eliminates quantifiers (in the module language)*

In a PID, a module is injective if and only if it is divisible, so this is a generalisation of the previous corollary.

Proof. The pp-formula $\exists\bar{y}(\bar{x} = B\bar{y})$ is equivalent to the formula $A\bar{x} = 0$, where A is a maximal-rank matrix satisfying $AB = 0$. □

2.5. **Algebraically closed fields.** Let $\mathcal{F}$ be the theory of fields in its natural signature. The quantifier-free one-variable sets are all finite or co-finite. In contrast, the set of squares $\exists y(x = y^2)$, for example, depends very much on the field in question, and need not be either (for example, in $\mathbb{R}$). So $\mathcal{F}$ does not eliminate quantifiers. It is definitely not complete.

Let $\mathcal{ACF}$ be the theory of algebraically closed fields. We know by categoricity that $\mathcal{ACF}$ becomes complete once we specify the characteristic. We will now see that even without specifying it, the theory eliminates quantifiers. Before doing this, we formulate another slight variant of the QE criterion:

Lemma 2.5.1. *A theory $\mathcal{T}$ eliminates quantifiers if and only if for any two models $\mathbb{M}_1$ and $\mathbb{M}_2$ of $\mathcal{T}$ admitting a common substructure A, and any $b \in \mathbb{M}_1$, the quantifier-free type of b over A is realised in some elementary extension of $\mathbb{M}_2$.*

Exercise 2.5.2. *Prove it*

Proposition 2.5.3. *The theory $\mathcal{ACF}$ of algebraically closed fields eliminates quantifiers*

Proof. Let K_1 and K_2 be two algebraically closed fields, with a common substructure A. Being a substructure means, in this case, being a subring, hence a subdomain. The fraction field of A must then also be included in both K_1 and K_2, so we may assume that A is a field.

Likewise, the algebraic closure of A in both K_1 and K_2 is an (absolute) algebraic closure of A (since the K_i are algebraically closed), and any two such algebraic closures are isomorphic over A. Picking such an isomorphism, we may assume that A is algebraically closed.

Let $b \in K_1$ be any element. If $b \in A$ we are done, so we may assume otherwise, i.e., that b is transcendental over A. This fact determines the quantifier-free type of b over A, so we need to find a transcendental element in an elementary extension of K_2. Since we are allowed to pass to an elementary extension, we may assume $K_2 \neq A$, but any element not in A is transcendental. □

As a corollary, we obtain another proof of the completeness of $\mathcal{ACF}_p$, using the following observation:

Exercise 2.5.4. *Let $\mathcal{T}$ be a theory with quantifier elimination such that all models of $\mathcal{T}$ have a common substructure. Then $\mathcal{T}$ is complete*

Corollary 2.5.5. *For p prime or 0, the theory $\mathcal{ACF}_p$ of algebraically closed fields of characteristic p is complete.*

End of lecture 4, Apr 11

2.5.6. *Digression: affine varieties.* Let X be a set, and let k be a field. The set k^X of k-valued functions on X form a commutative k-algebra, under pointwise operations. Let A be a sub-algebra. Then each $x \in X$ determines a k-algebra map $\underline{x} : A \to k$, given by evaluation: $\underline{x}(a) = a(x)$. The pair (X, A) is called an *affine algebraic variety* if this map is a bijection, and A is finitely generated as an algebra over k. If (Y, B) is another affine variety, a map function $f : X \to Y$ determines, by composition, a k-algebra map $f^* : k^Y \to k^X$. We call f a map of affine varieties if f^* restricts to a map from B to A (the restriction is also denoted f^*). Algebraic geometry can be viewed as the study of this category.

affine algebraic variety

Example 2.5.7. A k-algebra map from the polynomial algebra $k[x_1, \dots, x_n]$ to an algebra A is determined by the images of the generators x_i, which can take any value in A. Hence, the set of such maps can be canonically identified with A^n. Taking $A = k$, we see that the pair $(k^n, k[x_1, \dots, x_n])$ is an affine variety, called the *affine n-space*, denoted $\mathbb{A}^n$. □

affine n-space
$\mathbb{A}^n$

If (X, A) and (Y, B) are affine varieties, any homomorphism $t : B \to A$ is of the form f^*, where $f(x)(b) = t(b)(x)$ (we identify x with $\underline{x}$). Hence both the affine varieties and the maps between them are determined by the algebra alone, so a natural question is: which algebras correspond to affine varieties?

By definition, if A is the algebra of some affine variety, the set of points is given by $X_A = \mathrm{Hom}_{k-Alg}(A, k)$. The evaluation map, given by $a(x) = x(a)$ for $x \in X_A$ and $a \in A$ determines a k-algebra map $A \to k^{X_A}$, and so the question is, for which algebras A is this map injective? In other words, given a non-zero $a \in A$, is there a k-algebra map $x : A \to k$ such that $x(a) \neq 0$? Clearly, for this to hold a must not be nilpotent ($a^n \neq 0$ for all n). Given that, the localisation $A_a = A[t]/(ta - 1)$ is non-zero, and any map from it maps a to an invertible element, so our question becomes

Question 2.5.8. Suppose that A is a finitely generated non-zero k-algebra. Is there a k-algebra map $A \to k$? □

The assumption that A is finitely generated means that there is a surjective k-algebra map from some polynomial algebra to A, and the surjectivity means that the induced map from X_A to $\mathbb{A}^n$ is injective. If $I \subseteq k[x_1, \dots, x_n]$ is the kernel of the map, the image of X_A is precisely the set $Z(I)$ of common zeroes of all elements in I. Such subsets are called *Zariski closed*; they form a topology on $\mathbb{A}^n$ called the *Zariski topology*. Thus, every algebraic variety can be embedded as a Zariski closed subset of some affine space (but not uniquely). The assumption that A is non-zero means that I is a proper ideal, and Hilbert's *Basis Theorem* asserts that any finitely generated algebra is Noetherian, i.e., any ideal is finitely generated. Thus, the question above can be reformulated as:

Zariski closed

Zariski topology

Question 2.5.9. Suppose that $p_1(\bar{x}), \dots, p_m(\bar{x})$ are polynomials having no common zero in k. Does it follow that they generate the unit ideal? □

We now note that the answer to the question is, in general, "no": If k is not algebraically closed, it has some proper finite field extension L, and there are can be no maps from L to k. In terms of the second formulation, there is an irreducible polynomial $p(x)$ (one variable) that has no zero in k, but which is non-constant, and therefore does not generate the unit ideal. Hence, to have a reasonable theory, we need k to be algebraically closed. Hilbert's Nullstellensatz assert that this is sufficient:

Theorem 2.5.10 (Hilbert's Nullstellensatz)**.** *The following are equivalent for a field* k*:*

1. k *is algebraically closed*
2. *Any finitely generated* k*-algebra admits a map to* k
3. *Any polynomials* $p_1, \dots, p_m$ *over* k *that do not have a common zero in* k *generate the unit ideal.*

We had already explained all direction, except the difficult one, that k algebraically closed implies the other conditions. This will follow easily from quantifier elimination:

Proof. If A is a non-zero algebra, it has map to some field L extending k (by taking a maximal ideal). Hence, polynomials that do not generate the unit ideal have a common zero in some extension L, which we may further extend and assume to be algebraically closed. Viewing the polynomial equations as formulas in $\mathcal{ACF}$, we obtain by quantifier elimination that they also have a zero in k. □

2.6. **Real closed field.** The field $\mathbb{R}$ of real numbers does not eliminate quantifiers: the set of non-negative elements can be defined as the set of all squares, but this is neither finite nor co-finite. It is natural to view $\mathbb{R}$ as an ordered field. What is the rest of the theory?

Definition 2.6.1. An ordered field K is *real closed* if any positive $x \in K$ has a square root, and every polynomial of odd degree over K has a root in K. **real closed**

The field of reals is real closed by the intermediate value theorem. Thus, to show that this is the complete theory, we need to prove

Theorem 2.6.2. *The theory $\mathcal{RCF}$ of real closed fields eliminates quantifiers in the language of ordered fields. It is therefore complete.*

The proof is similar to the $\mathcal{ACF}$ case, but using less familiar algebraic facts, that we (partially) prove below.

Proof. Assume that K_1 and K_2 are real closed, with a common sub-domain A. The order on A extends uniquely to the fraction field, so we may assume that A is a field. By 2.6.3, the algebraic closures of A in K_1 and K_2 are canonically isomorphic. Therefore we may assume that A is real-closed.

Let $b \in K_1 \setminus A$. Then b is transcendental. Choose $b' \in K_2$ satisfying the same cut over A as b, like in the proof of Prop. 2.2.1. To show that b' and b satisfy the same quantifier-free type, it suffices to show that for any polynomial $p(x)$ over A we have $p(b) > 0$ if and only if $p(b') > 0$. If $p(x) = (x - a)q(x)$, the signs of $b - a$ and $b' - a$ are the same, since they satisfy the same cut, so we reduce to the case where p has no roots. Then the result follows from Prop 2.6.7. □

To complete the proof, we need the following two results:

{thm:rcl}

Theorem 2.6.3. *Any ordered field K has an algebraic real closed extension K^r. If L is any real closed field extending K, there is a unique embedding of K^r in L over K.*

It follows that the real closed field above is unique up to a unique isomorphism, and is called the *real closure* of K. **real closure**

To prove the theorem, we need to ask the question: which fields can be ordered at all? We phrase the question in terms of the set of positive elements in a potential partial order:

Definition 2.6.4. Let K be a field. A *semi-positive cone* in K is a subset P containing all sums of squares of elements in K, not containing -1, and closed under multiplication and addition. **semi-positive cone**

Lemma 2.6.5. *A field can be ordered if and only if it contains a semi-positive cone*

Proof. If K can be ordered, the set $P = \{x \in K | x \geqslant 0\}$ is a semi-positive cone. Conversely, assume P is a semi-positive cone. Since positive semi-cones are closed under unions of chains, we may assume that P is maximal. If $-b \notin P$, one easily checks that $P + bP$ is a semi-positive cone, hence $b \in P$. Thus P is the set of non-negative elements of an ordering. □

formally real

Since the set of sums of squares is closed under the field operations, the condition that a cone exists amounts to saying that -1 is not a sum of squares. A field with this property is called *formally real.*

Proof of 2.6.3. Let K be an ordered field. By the lemma, a field extension L can be ordered precisely if -1 cannot be written as $\sum a_i b_i^2$, where $b_i \in L$ and $a_i \in K$ positive. This condition is preserved under unions of chains, so we may take a maximal such L. We claim that L is real closed.

Assume that $0 < a \in L$ has no square root. Then $L(\sqrt{a})$ is a proper extension, and therefore cannot be ordered. Hence $-1 = \sum a_i(b_i + c_i\sqrt{a})^2$ for some $a_i, b_i, c_i \in L$ with $a_i > 0$. It follows that $-1 = \sum a_i(b_i^2 + ac_i^2)$ in L, but this is impossible, since a_i, a and the squares are positive.

The proof that odd-degree polynomials have a square root is similar. To prove the second part, one needs to show that if a polynomial $p(x)$ over K has a root in K^r then it has one in L, i.e., that the existence of a root can be inferred from the coefficients. This is a classical result, and is omitted. The uniqueness of the map follows from the ordering. □

The statement we need will follow directly from the "fundamental theorem of algebra".

{thm:fta}

Theorem 2.6.6. *A field K is real closed if and only if $C = K(\sqrt{-1})$ is algebraically closed*

Proof. Assume C is algebraically closed, let $p(x)$ be a polynomial over K. The set of roots is closed under complex conjugation, so if p is of odd degree, there is at least one in K, and if p has the form $x^2 - a$ then either the roots or their product by $\sqrt{-1}$ are fixed by conjugation, so are in K.

For the converse, let $G = \mathrm{Aut}(E/K)$, where E is a finite extension of C, and let H be a 2-sylow subgroup of G. The fixed field of H is an odd-degree extension of K, so must be trivial, since there are no irreducible polynomials of odd degree over K. It follows that $G = H$ is a 2-group, hence solvable, so E can be obtained from C by extracting square roots. However, it is easy to see that C is closed under all square roots. □

{cor:intermediate}

Corollary 2.6.7. *The intermediate value theorem holds for polynomials over real closed fields K: If $p(x)$ is a polynomial and $a, b \in K$ are such that $p(a) < 0 < p(b)$, then p has a root in K.*

Proof. It suffices to prove this for irreducible p, but by the last theorem, such a polynomial must be quadratic, of the form $p(x) = (x - a)^2 + b$, with $b > 0$. Such polynomials are always positive, by definition of real closed fields. □

End of lecture 5, Apr 16

Exercise 2.6.8. *Show that $\mathcal{RCF}$ can be defined as an ordered field satisfying the intermediate value theorem.*

As a corollary of QE we get:

Corollary 2.6.9. *Every definable subset of each model of $\mathcal{RCF}$ is a finite union of intervals*

A theory (expanding the theory of linear orders) that satisfies the above property is called an o-*minimal theory* ($\mathcal{DLO}$ is another example). **o-minimal theory**

Corollary 2.6.10. *The full theory of the reals (as an ordered field) is $\mathcal{RCF}$*

Similar to Hilbert's Nullstellensatz, a solution to Hilbert's seventeenth problem follows from QE for the reals. A real rational function $f(\bar{x})$ is called *non-negative* if its image consists of non-negative numbers.

Corollary 2.6.11 (Hilbert's problem 17)**.** *Any non-negative real rational function is a sum of squares of rational functions*

Proof. Such a function f can be viewed as an element of the field $K = \mathbb{R}(\bar{x})$ of rational functions. This field is formally real: otherwise, -1 is a sum of squares of rational functions, and evaluating at any point where they are defined we obtain a contradiction. We saw above that if f is not a sum of squares (in K), then K be ordered so that f is negative, compatibly with the order on $\mathbb{R}$, and embedded with this order into a real-closed field L. Hence, the formula $f(\bar{x}) < 0$ is satisfiable in a real-closed extension of $\mathbb{R}$. By QE, the extension $\mathbb{R} \subseteq L$ is elementary, so it is also satisfiable in $\mathbb{R}$. ☐

2.7. **Model completeness and model companions.** Reviewing the proofs of Hilbert's Nullstellensatz and of Hilbert 17th problem reveals a similar pattern: A field extension is embedded in a model of the theory in question, and then QE is used to show that if an extension contains a point of a quantifier-free formula over the original field, then so does the original field. We formalise the last condition in the following definition:

Definition 2.7.1. A structure $\mathbb{M}$ is *existentially closed (ec)* in a class $\mathcal{C}$ of structures if whenever $\phi(\bar{x}, \bar{y})$ is a quantifier-free formula, and $\bar{a} \in \mathbb{M}$ is such that $\phi(\bar{x}, \bar{a})^{\mathbb{M}} = \emptyset$, then $\phi(\bar{x}, \bar{a})^{\mathbb{N}} = \emptyset$ for all $\mathbb{N} \in \mathcal{C}$. If the class $\mathcal{C}$ is not mentioned, we take it to be (the class of models of) $\mathcal{T}_\forall$, the universal part of the theory $\mathcal{T}$ of $\mathbb{M}$. **existentially closed (ec)**

A theory $\mathcal{T}$ is *model complete* if all its models are existentially closed. **model complete**

{prp:modcomp}

Proposition 2.7.2. *The following are equivalent for a theory $\mathcal{T}$*

(1) *$\mathcal{T}$ is model complete*
(2) *$\mathcal{T}$ is the theory of all existentially closed models of $\mathcal{T}_\forall$*
(3) *Every extension of models of $\mathcal{T}$ is elementary*
(4) *Every universal formula is equivalent (with respect to $\mathcal{T}$) to an existential one*
(5) *Every formula is equivalent to an existential formula*
(6) *For every model $\mathbb{M}$, the theory $\mathcal{T}_{\mathbb{M}}$ is complete*

To prove that every definable set is existential, we would like to proceed as with quantifier elimination, and consider the space of *existential* types $\mathbb{S}^e(V)$, defined as $\{p^e | p \in \mathbb{S}(V)\}$, where p^e is the subset of p consisting of existential formulas. We topologise it by taking the existential formulas to be a basis of closed sets. Note that in general, the elements of $\mathbb{S}^e(V)$ need not be maximal among consistent sets of existential formulas, and the topology need not be Hausdorff. In fact we have:

Lemma 2.7.3. *A theory $\mathcal{T}$ is model complete if and only if every element of $\mathbb{S}^e(V)$ is maximal (equivalently, if it is T_1).*

Proof. Assume that $\mathcal{T}$ is model complete, let p be an existential type, and let $\phi(\bar{x})$ be an existential formula consistent with p. If $\phi \notin p$, there are models $\mathbb{M}$ and $\mathbb{N}$ of $\mathcal{T}$ and tuples $\overline{m}$ and $\bar{n}$, respectively, realising p, so that ϕ is satisfied by $\bar{n}$ but not by $\overline{m}$. We claim that $\mathbb{N}$ may be chosen to include $\mathbb{M}$, so that $\overline{m} = \bar{n}$. This will complete the proof, since $\mathbb{M}$ is assumed to be existentially closed.

To show the claim, let $\mathbb{N}$ be arbitrary as above, and consider the theory $\mathcal{T}$ expanded by the types of $\mathbb{M}$ and $\mathbb{N}$ (in a language with appropriate constant symbols), and with the additional axiom $\overline{m} = \bar{n}$. It suffices to show that this theory is consistent, since then we may take any model of it. The fact that it is consistent follows easily from compactness and the fact that both $\overline{m}$ and $\bar{n}$ satisfy p.

In the other direction, if $\mathbb{M}$ is a model of $\mathcal{T}$ which is not existentially closed, let $\phi(\bar{x}, \overline{m})$ be a formula satisfied in an extension but not in $\mathbb{M}$. Then the existential type of $\overline{m}$ in $\mathbb{M}$ is not maximal. □

We now assume that any extension of models of $\mathcal{T}$ is elementary, aiming to prove that every existentially closed model of $\mathcal{T}_\forall$ is a model of $\mathcal{T}$. The main tool will be the following observation, which is of independent interest.

{prp:elemchain}

Proposition 2.7.4. *If $\mathbb{M}_i$ is a filtering system of elementary extensions, then $\lim_i \mathbb{M}_i$ is an elementary extension of each*

The terminology of filtering systems and their limits is explained elsewhere. A central example is: $\mathbb{M}_i$ are indexed by an ordinal α, and for $i \leqslant j \in \alpha$, $\mathbb{M}_i \preceq \mathbb{M}_j$. The limit is the union. The proof of this statement is an exercise.

This observation allows us to find a semantic condition for ∀∃-axiomatizability.

{prp:inductive}

Proposition 2.7.5. *The following are equivalent for a theory $\mathcal{T}$*

(1) $\mathcal{T} \equiv \mathcal{T}_{\forall\exists}$
(2) *The of models of $\mathcal{T}$ is closed under unions of chains*

A theory satisfying the equivalent conditions above is called an *inductive theory.*

inductive theory

Example 2.7.6. Let $\mathcal{T}$ be the theory of $\mathbb{Z}$ as an ordered set. This theory states, in particular, that every element has an immediate successor, which is not $\forall\exists$ in the natural axiomatisation, but perhaps it can somehow be expressed as an $\forall\exists$ axiom?

Consider the chain of models $\frac{1}{2^i}\mathbb{Z}$, with inclusions. The union is dense, hence is not a model of $\mathcal{T}$. By the proposition, it has no $\forall\exists$ axiomatisation ☐

To prove the proposition, we use the following lemma.

Lemma 2.7.7. *Let $\mathbb{M}$ be a model of the $\forall\exists$ part of a theory $\mathcal{T}$. Then there are $\mathbb{M} \subseteq \mathbb{N} \subseteq \mathbb{M}_1$, with $\mathbb{N}$ a model of $\mathcal{T}$ and the inclusion $\mathbb{M} \subseteq \mathbb{M}_1$ elementary*

Proof. Let $\mathcal{T}_0$ be the universal theory of $\mathbb{M}$, in a language with constants for elements of $\mathbb{M}$. We claim that $\mathcal{T}_1 = \mathcal{T} \cup \mathcal{T}_0$ is consistent: otherwise, $\mathcal{T} \models \neg\phi$ where $\phi = \forall\bar{x}\psi(\bar{x}, \bar{m}) \in \mathcal{T}_0$, with $\bar{m}$ constants not in $\mathcal{T}$. Hence, $\mathcal{T} \models \forall\bar{y}\exists\bar{x}\neg\psi(\bar{x}, \bar{y})$. Since $\mathbb{M}$ is a model of $\mathcal{T}_{\forall\exists}$, this sentence holds in $\mathbb{M}$, a contradiction.

Let $\mathbb{N}$ be a model of $\mathcal{T}_1$. Then $\mathbb{M}$ is (isomorphic to) a substructure of $\mathbb{N}$. On the other hand, $(\mathcal{T}_1)_\forall \supseteq \mathcal{T}_0$, so $\mathbb{N}$ can be embedded in a model $\mathbb{M}'$ of the full theory of $\mathbb{M}$ (with constants for $\mathbb{M}$). Such a model is an elementary extension of $\mathbb{M}$. ☐

Proof of 2.7.5. Assume the class of models is closed under unions of systems, and let $\mathbb{M}$ be a model of $\mathcal{T}_{\forall\exists}$, we will prove it must be a model of $\mathcal{T}$. Let $\mathbb{M}_0 = \mathbb{M}$. Applying the lemma repeatedly, we find a chain of models $\mathbb{M}_i \subseteq \mathbb{N}_i \subseteq \mathbb{M}_{i+1}$, with $\mathbb{M}_i \preceq \mathbb{M}_{i+1}$ for all $i \geqslant 0$. The union of this chain is an elementary extension of $\mathbb{M} = \mathbb{M}_0$ by Prop 2.7.4, and a model of $\mathcal{T}$ by assumption. Hence $\mathbb{M}$ is a model of $\mathcal{T}$.

The other direction is left as an exercise ☐

3. Elimination of imaginaries

From now on, we assume quantifier elimination, unless mentioned otherwise. Thus, all embeddings are elementary.

3.1. Saturated models.

Definition 3.1.1. Let $A \subseteq \mathbb{M}$. An element $b \in \mathbb{M}$ is *definable* over A if $\{b\}$ is an A-definable subset. We let $(A) = \{b \in \mathbb{M} | b \text{ is definable over } A\}$ and say that A is *definably closed* if $A = (A)$. **definable** **definably closed**

The notion of a definably closed set is the natural notion of a "structure" in the context where our basic data is the collection of definable sets (rather than a particular signature generating them). The next exercise shows that the definition of a structure depends only on the theory, and not on the enclosing model.

Exercise 3.1.2. *Assume that $A \subseteq \mathbb{M}, \mathbb{N}$ as a structure, where $\mathbb{M} \equiv \mathbb{N}$. Show that ${}_{\mathbb{M}}(A)$ is uniquely isomorphic to ${}_{\mathbb{N}}(A)$*

Exercise 3.1.3. *Show that* $(A) = \{f(\bar{a}) | \bar{a} \in A, f$ *a definable function*$\}$ *(recall that a* definable function *is a function* $f : X \to Y$ *between definable sets, whose graph is a definable subset of* $X \times Y$*)*

definable function

Example 3.1.4. Assume that $\mathbb{M}$ is an algebraically closed field of characteristic p, and let $A \subseteq \mathbb{M}$ be a subset. Since the field operations are definable, we get from the last exercise that (A) contains the subfield $K(A)$ generated by A. Is there anything else? If $p > 0$, the Frobenius $x \mapsto x^p$ is bijective, hence its inverse is a definable function. Thus, (A) contains even the *perfect closure* of $K(A)$ (the smallest perfect subfield containing A).

We claim that this is in fact the whole definable closure. To show that, we recall that if K is perfect, and bK is algebraic, there is an automorphism of K^a that fixes K and moves b. Further, if $b \in \mathbb{M}$ is transcendental, we may extend it to a transcendence basis of $\mathbb{M}$ over K^a, and thus move it to any other transcendental element by an automorphism fixing K^a. In any case, if bK, we found an automorphism moving it while fixing K. The following exercise shows that no such element is contained in the definable closure. □

Exercise 3.1.5. *Show that if* $\sigma : \mathbb{M} \to \mathbb{M}$ *is an automorphism fixing* $A \subseteq \mathbb{M}$ *pointwise, then it also fixes each element of* (A)

In the example, we used a converse to the exercise to compute the definable closure. Can we always do this?

Example 3.1.6. Let $\mathbb{M} = \mathbb{R}$. What is $P = (0)$ in this case? As before, P is a subfield of $\mathbb{R}$. Following the previous example, we may ask if it is equal to the set of fixed points of the automorphism group of $\mathbb{R}$. However, this automorphism group is trivial (exercise!), so that would mean that $P = \mathbb{R}$. This is impossible, for example because P must be countable (in general, the size of the language).

Is it possible that $P = \mathbb{Q}$? The two roots of 2 can be distinguished in $\mathbb{R}$, since exactly one of them is positive. Thus, both belong to (0). Similarly, each algebraic real is definable. Thus P contains all real algebraic numbers. But this is already a model, and (A) must be included in each submodel that contains it, so P is the set of algebraic reals (more generally, (0) is a model in any o-minimal structure expanding a group). □

Exercise 3.1.7. *Show that if* $\mathbb{M}$ *is an algebraically closed field and* $A \subseteq \mathbb{M}$*, then* $b \in (A)$ *if and only if the type* $\mathrm{tp}(b/A)$ *of* b *over* A *is realised only by* b*. Show that the same if false for* $\mathbb{M} = \mathbb{R}$*.*

Definition 3.1.8. Let κ be a cardinal, and let $\mathbb{M}$ be a structure. A subset $A \subseteq \mathbb{M}$ is called κ-small if its cardinality is less than κ. $\mathbb{M}$ is called

- **κ-saturated** κ-*saturated* if any type over a κ-small subset is realised in $\mathbb{M}$
- **κ-homogeneous** κ-*homogeneous* if any partial elementary map $\tau : A \to \mathbb{M}$ on a κ-small subset A can be extended to any element of $\mathbb{M}$
- **κ-strongly homogeneous** κ-*strongly homogeneous* if any partial elementary map on a κ-small subset can be extended to an automorphism of $\mathbb{M}$.

If κ is omitted, we take $\kappa = \mathbb{M}$.

Clearly, $\kappa = \mathbb{M}$ is the largest cardinal for which these definitions make sense. We note

Exercise 3.1.9. *Let $\mathbb{M}$ be a structure. Show:*

(1) *If $\mathbb{M}$ is κ-saturated, then it is κ-homogeneous*
(2) *$\mathbb{M}$ is homogeneous if and only if it is strongly homogeneous*

In particular, a saturated model is strongly homogeneous

Proposition 3.1.10. *Let κ be a cardinal. Any model $\mathbb{M}$ can be embedded in a κ^+-saturated model of size at most $\mathbb{M}^\kappa$*

Proof. Enumerate everything and build elementary chains □

Review of set theory: A subset of $A \subseteq \kappa$ of a cardinal is bounded if it is contained in a strictly smaller ordinal. The cofinality κ is the smallest cardinality of an unbounded subset of κ. κ is called regular if $\kappa = \kappa$. It is called inaccessible if it is regular and is limit (not of the form λ^+).

Call a cardinal κ adequate (non-standard) if $\alpha < \kappa$ implies $2^\alpha \leqslant \kappa$.

Claim 3.1.11. *Assume that $\mathbb{M}$ is κ-strongly homogeneous. Then for any κ-small set A and any α-type over p over A, for $\alpha < \kappa$, the group $\mathrm{Aut}(\mathbb{M}/A)$ acts transitively on $p(\mathbb{M})$. If $\mathbb{M}$ is, in addition, κ-saturated, then the set of α-types over A is the set of orbits of $(\mathbb{M}/A)$ on $\mathbb{M}^\alpha$.*

Proposition 3.1.12. *Assume that $\mathbb{M}$ is $\mathcal{T}$-strongly homogeneous, and let $\mathbb{M}_0$ be a reduct to $\mathcal{T}_0$. If X is a definable set such that $X(\mathbb{M})$ is preserved (setwise) by $\mathrm{Aut}(\mathbb{M}_0)$, then X is definable in $\mathcal{T}_0$.*

Proof. By the claim, $X(\mathbb{M})$ is a union of types in $\mathcal{T}_0$. Apply compactness in $\mathcal{T}$ □

Corollary 3.1.13 (Beth definability). *Let $\mathcal{T}_0$ be a reduct of $\mathcal{T}$ with the same sorts, and assume that every model of $\mathcal{T}_0$ has at most one expansion to $\mathcal{T}$. Then they have the same definable sets*

Corollary 3.1.14. *Let X be an $\mathbb{M}$-definable set, where $\mathbb{M}$ is κ-strongly homogeneous, and let $A \subseteq \mathbb{M}$ be κ-small. Then X is A-definable if and only if it is preserved by $H = \mathrm{Aut}(\mathbb{M}/A)$. In particular, $(A) = \mathbb{M}^H$ (fixed points)*

Corollary 3.1.15. *If $\mathbb{M}$ is κ-saturated and κ-strongly homogeneous, and $A \subseteq \mathbb{M}$ is κ-small, then $b \in \mathbb{M}$ is definable over A if and only if it is the unique $\mathbb{M}$-point of $\mathrm{tp}(b/A)$*

Definition 3.1.16. For $A \subseteq \mathbb{M}$, an $b \in \mathbb{M}$ is an *algebraic element* over A if b belongs to a finite A-definable set. We let (A) be the set of all elements algebraic over A, and say that A is *algebraically closed* if $A = (A)$. **algebraic element** (A) **algebraically closed**

Corollary 3.1.17. *b algebraic over A iff it has a finite orbit under $(\mathbb{M}/A)$. If $\mathbb{M}$ is κ-saturated, iff $\mathrm{tp}(b/A)$ has finitely many solutions.*

3.2. **Canonical parameters.** We work in a sufficiently saturated and sufficiently strongly homogeneous model $\mathbb{M}$. If X is an $\mathbb{M}$-definable set, let $G_X = \{\sigma \in Aut(\mathbb{M}) | \sigma(X(\mathbb{M})) = X(\mathbb{M})\}$, and let $(X) = \mathbb{M}^{G_X}$ and $(X) = \{m \in \mathbb{M} | G_X m < \omega\}$. We saw above that X is A-definable if and only if it is fixed by $(\mathbb{M}/A)$, so $(X) \subseteq A$. If X is definable over (X), then we call (X) the *canonical base* of X. It is, in this case, the smallest subset of $\mathbb{M}$ over which X is defined.

canonical base

4. Definable groups and fields

Department of Math, Ben-Gurion University, Be'er-Sheva, Israel
Email address: mailto:kamenskm@math.bgu.ac.il
URL: http://mkamensky.notlong.com