

Model Theory for Algebra and Algebraic Geometry

David Marker

Spring 2010–Orsay

1 Language, Structures and Theories

In mathematical logic, we use first-order languages to describe mathematical structures. Intuitively, a structure is a set that we wish to study equipped with a collection of distinguished functions, relations, and elements. We then choose a language where we can talk about the distinguished functions, relations, and elements and nothing more. For example, when we study the ordered field of real numbers with the exponential function, we study the structure $(\mathbb{R}, +, \cdot, \exp, <, 0, 1)$, where the underlying set is the set of real numbers, and we distinguish the binary functions addition and multiplication, the unary function $x \mapsto e^x$, the binary order relation, and the real numbers 0 and 1. To describe this structure, we would use a language where we have symbols for $+$, \cdot , \exp , $<$, 0 , 1 and can write statements such as $\forall x \forall y \exp(x) \cdot \exp(y) = \exp(x + y)$ and $\forall x (x > 0 \rightarrow \exists y \exp(y) = x)$. We interpret these statements as the assertions “ $e^x e^y = e^{x+y}$ for all x and y ” and “for all positive x , there is a y such that $e^y = x$.”

For another example, we might consider the structure $(\mathbb{N}, +, 0, 1)$ of the natural numbers with addition and distinguished elements 0 and 1. The natural language for studying this structure is the language where we have a binary function symbol for addition and constant symbols for 0 and 1. We would write sentences such as $\forall x \exists y (x = y + y \vee x = y + y + 1)$, which we interpret as the assertion that “every number is either even or 1 plus an even number.”

Definition 1.1 A *language* \mathcal{L} is given by specifying the following data:

- i) a set of function symbols \mathcal{F} and positive integers n_f for each $f \in \mathcal{F}$;
- ii) a set of relation symbols \mathcal{R} and positive integers n_R for each $R \in \mathcal{R}$;
- iii) a set of constant symbols \mathcal{C} .

The numbers n_f and n_R tell us that f is a function of n_f variables and R is an n_R -ary relation.

Any or all of the sets \mathcal{F} , \mathcal{R} , and \mathcal{C} may be empty. Examples of languages include:

- i) the language of rings $\mathcal{L}_r = \{+, -, \cdot, 0, 1\}$, where $+$, $-$ and \cdot are binary function symbols and 0 and 1 are constants;

- ii) the language of ordered rings $\mathcal{L}_{\text{or}} = \mathcal{L}_r \cup \{<\}$, where $<$ is a binary relation symbol;
- iii) the language of pure sets $\mathcal{L} = \emptyset$;
- iv) the language of graphs is $\mathcal{L} = \{R\}$ where R is a binary relation symbol.

Next, we describe the structures where \mathcal{L} is the appropriate language.

Definition 1.2 An \mathcal{L} -structure \mathcal{M} is given by the following data:

- i) a nonempty set M called the *universe, domain, or underlying set* of \mathcal{M} ;
- ii) a function $f^{\mathcal{M}} : M^{n_f} \rightarrow M$ for each $f \in \mathcal{F}$;
- iii) a set $R^{\mathcal{M}} \subseteq M^{n_R}$ for each $R \in \mathcal{R}$;
- iv) an element $c^{\mathcal{M}} \in M$ for each $c \in \mathcal{C}$.

We refer to $f^{\mathcal{M}}$, $R^{\mathcal{M}}$, and $c^{\mathcal{M}}$ as the *interpretations* of the symbols f , R , and c . We often write the structure as $\mathcal{M} = (M, f^{\mathcal{M}}, R^{\mathcal{M}}, c^{\mathcal{M}} : f \in \mathcal{F}, R \in \mathcal{R}, \text{ and } c \in \mathcal{C})$. We will use the notation A, B, M, N, \dots to refer to the underlying sets of the structures $\mathcal{A}, \mathcal{B}, \mathcal{M}, \mathcal{N}, \dots$

For example, suppose that we are studying groups. We might use the language $\mathcal{L}_g = \{\cdot, e\}$, where \cdot is a binary function symbol and e is a constant symbol. An \mathcal{L}_g -structure $\mathcal{G} = (G, \cdot^{\mathcal{G}}, e^{\mathcal{G}})$ will be a set G equipped with a binary relation $\cdot^{\mathcal{G}}$ and a distinguished element $e^{\mathcal{G}}$. For example, $\mathcal{G} = (\mathbb{R}, \cdot, 1)$ is an \mathcal{L}_g -structure where we interpret \cdot as multiplication and e as 1; that is, $\cdot^{\mathcal{G}} = \cdot$ and $e^{\mathcal{G}} = 1$. Also, $\mathcal{N} = (\mathbb{N}, +, 0)$ is an \mathcal{L}_g -structure where $\cdot^{\mathcal{N}} = +$ and $e^{\mathcal{N}} = 0$. Of course, \mathcal{N} is not a group, but it is an \mathcal{L}_g -structure.

Usually, we will choose languages that closely correspond to the structure that we wish to study. For example, if we want to study the real numbers as an ordered field, we would use the language of ordered rings \mathcal{L}_{or} and give each symbol its natural interpretation.

We will study maps that preserve the interpretation of \mathcal{L} .

Definition 1.3 Suppose that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures with universes M and N , respectively. An \mathcal{L} -embedding $\eta : \mathcal{M} \rightarrow \mathcal{N}$ is a one-to-one map $\eta : M \rightarrow N$ that preserves the interpretation of all of the symbols of \mathcal{L} . More precisely:

- i) $\eta(f^{\mathcal{M}}(a_1, \dots, a_{n_f})) = f^{\mathcal{N}}(\eta(a_1), \dots, \eta(a_{n_f}))$ for all $f \in \mathcal{F}$ and $a_1, \dots, a_{n_f} \in M$;
- ii) $(a_1, \dots, a_{m_R}) \in R^{\mathcal{M}}$ if and only if $(\eta(a_1), \dots, \eta(a_{m_R})) \in R^{\mathcal{N}}$ for all $R \in \mathcal{R}$ and $a_1, \dots, a_{m_R} \in M$;
- iii) $\eta(c^{\mathcal{M}}) = c^{\mathcal{N}}$ for $c \in \mathcal{C}$.

A bijective \mathcal{L} -embedding is called an \mathcal{L} -isomorphism. If $M \subseteq N$ and the inclusion map is an \mathcal{L} -embedding, we say either that \mathcal{M} is a *substructure* of \mathcal{N} or that \mathcal{N} is an *extension* of \mathcal{M} .

For example:

- i) $(\mathbb{Z}, +, 0)$ is a substructure of $(\mathbb{R}, +, 0)$.
- ii) If $\eta : \mathbb{Z} \rightarrow \mathbb{R}$ is the function $\eta(x) = e^x$, then η is an \mathcal{L}_g -embedding of $(\mathbb{Z}, +, 0)$ into $(\mathbb{R}, \cdot, 1)$.

The *cardinality* of \mathcal{M} is $|M|$, the cardinality of the universe of \mathcal{M} . If $\eta : \mathcal{M} \rightarrow \mathcal{N}$ is an embedding then the cardinality of \mathcal{N} is at least the cardinality of \mathcal{M} .

We use the language \mathcal{L} to create formulas describing properties of \mathcal{L} -structures. Formulas will be strings of symbols built using the symbols of \mathcal{L} , variable symbols v_1, v_2, \dots , the equality symbol $=$, the Boolean connectives \wedge, \vee , and \neg , which we read as “and,” “or,” and “not”, the quantifiers \exists and \forall , which we read as “there exists” and “for all”, and parentheses $(,)$.

Definition 1.4 The set of \mathcal{L} -terms is the smallest set \mathcal{T} such that

- i) $c \in \mathcal{T}$ for each constant symbol $c \in \mathcal{C}$,
- ii) each variable symbol $v_i \in \mathcal{T}$ for $i = 1, 2, \dots$, and
- iii) if $t_1, \dots, t_{n_f} \in \mathcal{T}$ and $f \in \mathcal{F}$, then $f(t_1, \dots, t_{n_f}) \in \mathcal{T}$.

For example, $\cdot(v_1, -(v_3, 1))$, $\cdot(+v_1, v_2), +(v_3, 1)$ and $+(1, +(1, +(1, 1)))$ are \mathcal{L}_r -terms. For simplicity, we will usually write these terms in the more standard notation $v_1(v_3 - 1)$, $(v_1 + v_2)(v_3 + 1)$, and $1 + (1 + (1 + 1))$ when no confusion arises. In the \mathcal{L}_r -structure $(\mathbb{Z}, +, \cdot, 0, 1)$, we think of the term $1 + (1 + (1 + 1))$ as a name for the element 4, while $(v_1 + v_2)(v_3 + 1)$ is a name for the function $(x, y, z) \mapsto (x + y)(z + 1)$. This can be done in any \mathcal{L} -structure.

Suppose that \mathcal{M} is an \mathcal{L} -structure and that t is a term built using variables from $\bar{v} = (v_{i_1}, \dots, v_{i_m})$. We want to interpret t as a function $t^{\mathcal{M}} : M^m \rightarrow M$. For s a subterm of t and $\bar{a} = (a_{i_1}, \dots, a_{i_m}) \in M$, we inductively define $s^{\mathcal{M}}(\bar{a})$ as follows.

- i) If s is a constant symbol c , then $s^{\mathcal{M}}(\bar{a}) = c^{\mathcal{M}}$.
- ii) If s is the variable v_{i_j} , then $s^{\mathcal{M}}(\bar{a}) = a_{i_j}$.
- iii) If s is the term $f(t_1, \dots, t_{n_f})$, where f is a function symbol of \mathcal{L} and t_1, \dots, t_{n_f} are terms, then $s^{\mathcal{M}}(\bar{a}) = f^{\mathcal{M}}(t_1^{\mathcal{M}}(\bar{a}), \dots, t_{n_f}^{\mathcal{M}}(\bar{a}))$.

The function $t^{\mathcal{M}}$ is defined by $\bar{a} \mapsto t^{\mathcal{M}}(\bar{a})$.

For example, let $\mathcal{L} = \{f, g, c\}$, where f is a unary function symbol, g is a binary function symbol, and c is a constant symbol. We will consider the \mathcal{L} -terms $t_1 = g(v_1, c)$, $t_2 = f(g(c, f(v_1)))$, and $t_3 = g(f(g(v_1, v_2)), g(v_1, f(v_2)))$. Let \mathcal{M} be the \mathcal{L} -structure $(\mathbb{R}, \exp, +, 1)$; that is, $f^{\mathcal{M}} = \exp$, $g^{\mathcal{M}} = +$, and $c^{\mathcal{M}} = 1$.

Then

$$t_1^{\mathcal{M}}(a_1) = a_1 + 1,$$

$$t_2^{\mathcal{M}}(a_1) = e^{1+e^{a_1}}, \text{ and}$$

$$t_3^{\mathcal{M}}(a_1, a_2) = e^{a_1+a_2} + (a_1 + e^{a_2}).$$

We are now ready to define \mathcal{L} -formulas.

Definition 1.5 We say that ϕ is an *atomic \mathcal{L} -formula* if ϕ is either

- i) $t_1 = t_2$, where t_1 and t_2 are terms, or
- ii) $R(t_1, \dots, t_{n_R})$, where $R \in \mathcal{R}$ and t_1, \dots, t_{n_R} are terms.

The set of \mathcal{L} -formulas is the smallest set \mathcal{W} containing the atomic formulas such that

- i) if ϕ is in \mathcal{W} , then $\neg\phi$ is in \mathcal{W} ,
- ii) if ϕ and ψ are in \mathcal{W} , then $(\phi \wedge \psi)$ and $(\phi \vee \psi)$ are in \mathcal{W} , and

iii) if ϕ is in \mathcal{W} , then $\exists v_i \phi$ and $\forall v_i \phi$ are in \mathcal{W} .

Here are three examples of \mathcal{L}_{or} -formulas.

- $v_1 = 0 \vee v_1 > 0$.
- $\exists v_2 v_2 \cdot v_2 = v_1$.
- $\forall v_1 (v_1 = 0 \vee \exists v_2 v_2 \cdot v_1 = 1)$.

Intuitively, the first formula asserts that $v_1 \geq 0$, the second asserts that v_1 is a square, and the third asserts that every nonzero element has a multiplicative inverse. We would like to define what it means for a formula to be true in a structure, but these examples already show one difficulty. While in any \mathcal{L}_{or} -structure the third formula will either be true or false, the first two formulas express a property that may or may not be true of particular elements of the structure. In the \mathcal{L}_{or} -structure $(\mathbb{Z}, +, -, \cdot, <, 0, 1)$, the second formula would be true of 9 but false of 8.

We say that a variable v *occurs freely* in a formula ϕ if it is not inside a $\exists v$ or $\forall v$ quantifier; otherwise, we say that it is *bound*.¹ For example v_1 is free in the first two formulas and bound in the third, whereas v_2 is bound in both formulas. We call a formula a *sentence* if it has no free variables.

Let \mathcal{M} be an \mathcal{L} -structure. We will see that each \mathcal{L} -sentence is either true or false in \mathcal{M} . On the other hand, if ϕ is a formula with free variables v_1, \dots, v_n , we will think of ϕ as expressing a property of elements of M^n . We often write $\phi(v_1, \dots, v_n)$ to make explicit the free variables in ϕ . We must define what it means for $\phi(v_1, \dots, v_n)$ to hold of $(a_1, \dots, a_n) \in M^n$.

Definition 1.6 Let ϕ be a formula with free variables from $\bar{v} = (v_{i_1}, \dots, v_{i_m})$, and let $\bar{a} = (a_{i_1}, \dots, a_{i_m}) \in M^m$. We inductively define $\mathcal{M} \models \phi(\bar{a})$ as follows.

- i) If ϕ is $t_1 = t_2$, then $\mathcal{M} \models \phi(\bar{a})$ if $t_1^{\mathcal{M}}(\bar{a}) = t_2^{\mathcal{M}}(\bar{a})$.
 - ii) If ϕ is $R(t_1, \dots, t_{n_R})$, then $\mathcal{M} \models \phi(\bar{a})$ if $(t_1^{\mathcal{M}}(\bar{a}), \dots, t_{n_R}^{\mathcal{M}}(\bar{a})) \in R^{\mathcal{M}}$.
 - iii) If ϕ is $\neg\psi$, then $\mathcal{M} \models \phi(\bar{a})$ if $\mathcal{M} \not\models \psi(\bar{a})$.
 - iv) If ϕ is $(\psi \wedge \theta)$, then $\mathcal{M} \models \phi(\bar{a})$ if $\mathcal{M} \models \psi(\bar{a})$ and $\mathcal{M} \models \theta(\bar{a})$.
 - v) If ϕ is $(\psi \vee \theta)$, then $\mathcal{M} \models \phi(\bar{a})$ if $\mathcal{M} \models \psi(\bar{a})$ or $\mathcal{M} \models \theta(\bar{a})$.
 - vi) If ϕ is $\exists v_j \psi(\bar{v}, v_j)$, then $\mathcal{M} \models \phi(\bar{a})$ if there is $b \in M$ such that $\mathcal{M} \models \psi(\bar{a}, b)$.
 - vii) If ϕ is $\forall v_j \psi(\bar{v}, v_j)$, then $\mathcal{M} \models \phi(\bar{a})$ if $\mathcal{M} \models \psi(\bar{a}, b)$ for all $b \in M$.
- If $\mathcal{M} \models \phi(\bar{a})$ we say that \mathcal{M} *satisfies* $\phi(\bar{a})$ or $\phi(\bar{a})$ is *true* in \mathcal{M} .

Remarks 1.7 • There are a number of useful abbreviations that we will use: $\phi \rightarrow \psi$ is an abbreviation for $\neg\phi \vee \psi$, and $\phi \leftrightarrow \psi$ is an abbreviation for $(\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$. In fact, we did not really need to include the symbols \vee and \forall . We could have considered $\phi \vee \psi$ as an abbreviation for $\neg(\neg\phi \wedge \neg\psi)$ and $\forall v\phi$ as an abbreviation for $\neg(\exists v\neg\phi)$. Viewing these as abbreviations will be an advantage

¹To simplify some bookkeeping we will tacitly restrict our attention to formulas where in each subformula no variable v_i has both free and bound occurrences. For example we will not consider formulas such as $(v_1 > 0 \vee \exists v_1 v_1 \cdot v_1 = v_2)$, because this formula could be replaced by the clearer formula $v_1 > 0 \vee \exists v_3 v_3 \cdot v_3 = v_2$ with the same meaning. There are some areas of mathematical logic where one wants to be frugal with variables, but we will not consider such issues here. See [?] for a definition of satisfaction for arbitrary formulas.

when we are proving theorems by induction on formulas because it eliminates the \vee and \forall cases.

We also will use the abbreviations $\bigwedge_{i=1}^n \psi_i$ and $\bigvee_{i=1}^n \psi_i$ for $\psi_1 \wedge \dots \wedge \psi_n$ and $\psi_1 \vee \dots \vee \psi_n$, respectively.

- In addition to v_1, v_2, \dots , we will use w, x, y, z, \dots as variable symbols.
- It is important to note that the quantifiers \exists and \forall range only over elements of the model. For example the statement that an ordering is complete (i.e., every bounded subset has a least upper bound) cannot be expressed as a formula because we cannot quantify over subsets. The fact that we are limited to quantification over elements of the structure is what makes it “first-order” logic.

When proving results about satisfaction in models, we often must do an induction on the construction of formulas. The next proposition asserts that if a formula without quantifiers is true in some structure, then it is true in every extension. It is proved by induction on quantifier-free formulas.

Proposition 1.8 *Suppose that \mathcal{M} is a substructure of \mathcal{N} , $\bar{a} \in M$, and $\phi(\bar{v})$ is a quantifier-free formula. Then, $\mathcal{M} \models \phi(\bar{a})$ if and only if $\mathcal{N} \models \phi(\bar{a})$.*

Proof

Claim If $t(\bar{v})$ is a term and $\bar{b} \in M$, then $t^{\mathcal{M}}(\bar{b}) = t^{\mathcal{N}}(\bar{b})$. This is proved by induction on terms.

If t is the constant symbol c , then $c^{\mathcal{M}} = c^{\mathcal{N}}$.

If t is the variable v_i , then $t^{\mathcal{M}}(\bar{b}) = b_i = t^{\mathcal{N}}(\bar{b})$.

Suppose that $t = f(t_1, \dots, t_n)$, where f is an n -ary function symbol, t_1, \dots, t_n are terms, and $t_i^{\mathcal{M}}(\bar{b}) = t_i^{\mathcal{N}}(\bar{b})$ for $i = 1, \dots, n$. Because $\mathcal{M} \subseteq \mathcal{N}$, $f^{\mathcal{M}} = f^{\mathcal{N}}|M^n$. Thus,

$$\begin{aligned} t^{\mathcal{M}}(\bar{b}) &= f^{\mathcal{M}}(t_1^{\mathcal{M}}(\bar{b}), \dots, t_n^{\mathcal{M}}(\bar{b})) \\ &= f^{\mathcal{N}}(t_1^{\mathcal{M}}(\bar{b}), \dots, t_n^{\mathcal{M}}(\bar{b})) \\ &= f^{\mathcal{N}}(t_1^{\mathcal{N}}(\bar{b}), \dots, t_n^{\mathcal{N}}(\bar{b})) \\ &= t^{\mathcal{N}}(\bar{b}). \end{aligned}$$

We now prove the proposition by induction on formulas.

If ϕ is $t_1 = t_2$, then

$$\mathcal{M} \models \phi(\bar{a}) \Leftrightarrow t_1^{\mathcal{M}}(\bar{a}) = t_2^{\mathcal{M}}(\bar{a}) \Leftrightarrow t_1^{\mathcal{N}}(\bar{a}) = t_2^{\mathcal{N}}(\bar{a}) \Leftrightarrow \mathcal{N} \models \phi(\bar{a}).$$

If ϕ is $R(t_1, \dots, t_n)$, where R is an n -ary relation symbol, then

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow (t_1^{\mathcal{M}}(\bar{a}), \dots, t_n^{\mathcal{M}}(\bar{a})) \in R^{\mathcal{M}} \\ &\Leftrightarrow (t_1^{\mathcal{M}}(\bar{a}), \dots, t_n^{\mathcal{M}}(\bar{a})) \in R^{\mathcal{N}} \\ &\Leftrightarrow (t_1^{\mathcal{N}}(\bar{a}), \dots, t_n^{\mathcal{N}}(\bar{a})) \in R^{\mathcal{N}} \\ &\Leftrightarrow \mathcal{N} \models \phi(\bar{a}). \end{aligned}$$

Thus, the proposition is true for all atomic formulas.

Suppose that the proposition is true for ψ and that ϕ is $\neg\psi$. Then,

$$\mathcal{M} \models \phi(\bar{a}) \Leftrightarrow \mathcal{M} \not\models \psi(\bar{a}) \Leftrightarrow \mathcal{N} \not\models \psi(\bar{a}) \Leftrightarrow \mathcal{N} \models \phi(\bar{a}).$$

Finally, suppose that the proposition is true for ψ_0 and ψ_1 and that ϕ is $\psi_0 \wedge \psi_1$. Then,

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow \mathcal{M} \models \psi_0(\bar{a}) \text{ and } \mathcal{M} \models \psi_1(\bar{a}) \\ &\Leftrightarrow \mathcal{N} \models \psi_0(\bar{a}) \text{ and } \mathcal{M} \models \psi_1(\bar{a}) \\ &\Leftrightarrow \mathcal{N} \models \phi(\bar{a}). \end{aligned}$$

We have shown that the proposition holds for all atomic formulas and that if it holds for ϕ and ψ , then it also holds for $\neg\phi$ and $\phi \wedge \psi$. Because the set of quantifier-free formulas is the smallest set of formulas containing the atomic formulas and closed under negation and conjunction, the proposition is true for all quantifier-free formulas.

Elementary Equivalence and Isomorphism

We next consider structures that satisfy the same sentences.

Definition 1.9 We say that two \mathcal{L} -structures \mathcal{M} and \mathcal{N} are *elementarily equivalent* and write $\mathcal{M} \equiv \mathcal{N}$ if

$$\mathcal{M} \models \phi \text{ if and only if } \mathcal{N} \models \phi$$

for all \mathcal{L} -sentences ϕ .

We let $\text{Th}(\mathcal{M})$, the *full theory of \mathcal{M}* , be the set of \mathcal{L} -sentences ϕ such that $\mathcal{M} \models \phi$. It is easy to see that $\mathcal{M} \equiv \mathcal{N}$ if and only if $\text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$. Our next result shows that $\text{Th}(\mathcal{M})$ is an isomorphism invariant of \mathcal{M} . The proof uses the important technique of “induction on formulas.”

Theorem 1.10 *Suppose that $j : \mathcal{M} \rightarrow \mathcal{N}$ is an isomorphism. Then, $\mathcal{M} \equiv \mathcal{N}$.*

Proof We show by induction on formulas that $\mathcal{M} \models \phi(a_1, \dots, a_n)$ if and only if $\mathcal{N} \models \phi(j(a_1), \dots, j(a_n))$ for all formulas ϕ .

We first must show that terms behave well.

Claim Suppose that t is a term and the free variables in t are from $\bar{v} = (v_1, \dots, v_n)$. For $\bar{a} = (a_1, \dots, a_n) \in M$, we let $j(\bar{a})$ denote $(j(a_1), \dots, j(a_n))$. Then $j(t^{\mathcal{M}}(\bar{a})) = t^{\mathcal{N}}(j(\bar{a}))$.

We prove this by induction on terms.

- i) If $t = c$, then $j(t^{\mathcal{M}}(\bar{a})) = j(c^{\mathcal{M}}) = c^{\mathcal{N}} = t^{\mathcal{N}}(j(\bar{a}))$.
- ii) If $t = v_i$, then $j(t^{\mathcal{M}}(\bar{a})) = j(a_i) = t^{\mathcal{N}}(j(\bar{a}))$.
- iii) If $t = f(t_1, \dots, t_m)$, then

$$\begin{aligned} j(t^{\mathcal{M}}(\bar{a})) &= j(f^{\mathcal{M}}(t_1^{\mathcal{M}}(\bar{a}), \dots, t_m^{\mathcal{M}}(\bar{a}))) \\ &= f^{\mathcal{N}}(j(t_1^{\mathcal{M}}(\bar{a})), \dots, j(t_m^{\mathcal{M}}(\bar{a}))) \\ &= f^{\mathcal{N}}(t_1^{\mathcal{N}}(j(\bar{a})), \dots, t_m^{\mathcal{N}}(j(\bar{a}))) \\ &= t^{\mathcal{N}}(j(\bar{a})). \end{aligned}$$

We proceed by induction on formulas.

i) If $\phi(\bar{v})$ is $t_1 = t_2$, then

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow t_1^{\mathcal{M}}(\bar{a}) = t_2^{\mathcal{M}}(\bar{a}) \\ &\Leftrightarrow j(t_1^{\mathcal{M}}(\bar{a})) = j(t_2^{\mathcal{M}}(\bar{a})) \text{ because } j \text{ is injective} \\ &\Leftrightarrow t_1^{\mathcal{N}}(j(\bar{a})) = t_2^{\mathcal{N}}(j(\bar{a})) \\ &\Leftrightarrow \mathcal{N} \models \phi(j(\bar{a})). \end{aligned}$$

ii) If $\phi(\bar{v})$ is $R(t_1, \dots, t_n)$, then

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow (t_1^{\mathcal{M}}(\bar{a}), \dots, t_n^{\mathcal{M}}(\bar{a})) \in R^{\mathcal{M}} \\ &\Leftrightarrow (j(t_1^{\mathcal{M}}(\bar{a})), \dots, j(t_n^{\mathcal{M}}(\bar{a}))) \in R^{\mathcal{N}} \\ &\Leftrightarrow (t_1^{\mathcal{N}}(j(\bar{a})), \dots, t_n^{\mathcal{N}}(j(\bar{a}))) \in R^{\mathcal{N}} \\ &\Leftrightarrow \mathcal{N} \models \phi(j(\bar{a})). \end{aligned}$$

iii) If ϕ is $\neg\psi$, then by induction

$$\mathcal{M} \models \phi(\bar{a}) \Leftrightarrow \mathcal{M} \not\models \psi(\bar{a}) \Leftrightarrow \mathcal{N} \not\models \psi(j(\bar{a})) \Leftrightarrow \mathcal{N} \models \phi(j(\bar{a})).$$

iv) If ϕ is $\psi \wedge \theta$, then

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow \mathcal{M} \models \psi(\bar{a}) \text{ and } \mathcal{M} \models \theta(\bar{a}) \\ &\Leftrightarrow \mathcal{N} \models \psi(j(\bar{a})) \text{ and } \mathcal{N} \models \theta(j(\bar{a})) \Leftrightarrow \mathcal{N} \models \phi(j(\bar{a})). \end{aligned}$$

v) If $\phi(\bar{v})$ is $\exists w \psi(\bar{v}, w)$, then

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow \mathcal{M} \models \psi(\bar{a}, b) \text{ for some } b \in M \\ &\Leftrightarrow \mathcal{N} \models \psi(j(\bar{a}), c) \text{ for some } c \in N \text{ because } j \text{ is onto} \\ &\Leftrightarrow \mathcal{N} \models \phi(j(\bar{a})). \end{aligned}$$

Theories

Let \mathcal{L} be a language. An \mathcal{L} -theory T is simply a set of \mathcal{L} -sentences. We say that \mathcal{M} is a *model* of T and write $\mathcal{M} \models T$ if $\mathcal{M} \models \phi$ for all sentences $\phi \in T$.

The set $T = \{\forall x x = 0, \exists x x \neq 0\}$ is a theory. Because the two sentences in T are contradictory, there are no models of T . We say that a theory is *satisfiable* if it has a model.

We say that a class of \mathcal{L} -structures \mathcal{K} is an *elementary class* if there is an \mathcal{L} -theory T such that $\mathcal{K} = \{\mathcal{M} : \mathcal{M} \models T\}$.

One way to get a theory is to take $\text{Th}(\mathcal{M})$, the full theory of an \mathcal{L} -structure \mathcal{M} . In this case, the elementary class of models of $\text{Th}(\mathcal{M})$ is exactly the class of \mathcal{L} -structures elementarily equivalent to \mathcal{M} . More typically, we have a class of structures in mind and try to write a set of properties T describing these structures. We call these sentences *axioms* for the elementary class.

We give a few basic examples of theories and elementary classes that we will return to frequently.

Example 1.11 *Infinite Sets*

Let $\mathcal{L} = \emptyset$.

Consider the \mathcal{L} -theory where we have, for each n , the sentence ϕ_n given by

$$\exists x_1 \exists x_2 \dots \exists x_n \bigwedge_{i < j \leq n} x_i \neq x_j.$$

The sentence ϕ_n asserts that there are at least n distinct elements, and an \mathcal{L} -structure \mathcal{M} with universe M is a model of T if and only if M is infinite.

Example 1.12 *Linear Orders*

Let $\mathcal{L} = \{<\}$, where $<$ is a binary relation symbol. The class of linear orders is axiomatized by the \mathcal{L} -sentences

$$\begin{aligned} &\forall x \neg(x < x), \\ &\forall x \forall y \forall z ((x < y \wedge y < z) \rightarrow x < z), \\ &\forall x \forall y (x < y \vee x = y \vee y < x). \end{aligned}$$

There are a number of interesting extensions of the theory of linear orders. For example, we could add the sentence

$$\forall x \forall y (x < y \rightarrow \exists z (x < z \wedge z < y))$$

to get the theory of dense linear orders, or we could instead add the sentence

$$\forall x \exists y (x < y \wedge \forall z (x < z \rightarrow (z = y \vee y < z)))$$

to get the theory of linear orders where every element has a unique successor. We could also add sentences that either assert or deny the existence of top or bottom elements.

Example 1.13 *Equivalence Relations*

Let $\mathcal{L} = \{E\}$, where E is a binary relation symbol. The theory of equivalence relations is given by the sentences

$$\begin{aligned} &\forall x E(x, x), \\ &\forall x \forall y (E(x, y) \rightarrow E(y, x)), \\ &\forall x \forall y \forall z ((E(x, y) \wedge E(y, z)) \rightarrow E(x, z)). \end{aligned}$$

If we added the sentence

$$\forall x \exists y (x \neq y \wedge E(x, y) \wedge \forall z (E(x, z) \rightarrow (z = x \vee z = y)))$$

we would have the theory of equivalence relations where every equivalence class has exactly two elements. If instead we added the sentence

$$\exists x \exists y (\neg E(x, y) \wedge \forall z (E(x, z) \vee E(y, z)))$$

and the infinitely many sentences

$$\forall x \exists x_1 \exists x_2 \dots \exists x_n \left(\bigwedge_{i < j \leq n} x_i \neq x_j \wedge \bigwedge_{i=1}^n E(x, x_i) \right)$$

we would axiomatize the class of equivalence relations with exactly two classes, both of which are infinite.

Example 1.14 *Graphs*

Let $\mathcal{L} = \{R\}$ where R is a binary relation. We restrict our attention to irreflexive graphs. These are axiomatized by the two sentences

$$\begin{aligned} \forall x \neg R(x, x), \\ \forall x \forall y (R(x, y) \rightarrow R(y, x)). \end{aligned}$$

Example 1.15 *Groups*

Let $\mathcal{L} = \{\cdot, e\}$, where \cdot is a binary function symbol and e is a constant symbol. We will write $x \cdot y$ rather than $\cdot(x, y)$. The class of groups is axiomatized by

$$\begin{aligned} \forall x e \cdot x = x \cdot e = x, \\ \forall x \forall y \forall z x \cdot (y \cdot z) = (x \cdot y) \cdot z, \\ \forall x \exists y x \cdot y = y \cdot x = e. \end{aligned}$$

We could also axiomatize the class of Abelian groups by adding $\forall x \forall y x \cdot y = y \cdot x$.

Let $\phi_n(x)$ be the \mathcal{L} -formula

$$\underbrace{x \cdot x \cdots x}_{n\text{-times}} = e;$$

which asserts that $nx = e$.

We could axiomatize the class of torsion-free groups by adding $\{\forall x (x = e \vee \neg \phi_n(x)) : n \geq 2\}$ to the axioms for groups. Alternatively, we could axiomatize the class of groups where every element has order at most N by adding to the axioms for groups the sentence

$$\forall x \bigvee_{n \leq N} \phi_n(x).$$

Note that the same idea will not work to axiomatize the class of torsion groups because the corresponding sentence would be infinitely long. In the next chapter, we will see that the class of torsion groups is not elementary.

Let $\psi_n(x, y)$ be the formula

$$\underbrace{x \cdot x \cdots x}_{n\text{-times}} = y;$$

which asserts that $x^n = y$. We can axiomatize the class of divisible groups by adding the axioms $\{\forall y \exists x \psi_n(x, y) : n \geq 2\}$.

It will often be useful to deal with additive groups instead of multiplicative groups. The class of additive groups is the collection structures in the language $\mathcal{L} = \{+, 0\}$, axiomatized as above replacing \cdot by $+$ and e by 0 .

Example 1.16 *Ordered Abelian Groups*

Let $\mathcal{L} = \{+, <, 0\}$, where $+$ is a binary function symbol, $<$ is a binary relation symbol, and 0 is a constant symbol. The axioms for ordered groups are

- the axioms for additive groups,
- the axioms for linear orders, and
- $\forall x \forall y \forall z (x < y \rightarrow x + z < y + z)$.

Example 1.17 *Left R -modules*

Let R be a ring with multiplicative identity 1 . Let $\mathcal{L} = \{+, 0\} \cup \{r : r \in R\}$ where $+$ is a binary function symbol, 0 is a constant, and r is a unary function symbol for $r \in R$. In an R -module, we will interpret r as scalar multiplication by R . The axioms for left R -modules are

- the axioms for additive commutative groups,
- $\forall x \ r(x + y) = r(x) + r(y)$ for each $r \in R$,
- $\forall x \ (r + s)(x) = r(x) + s(x)$ for each $r, s \in R$,
- $\forall x \ r(s(x)) = rs(x)$ for $r, s \in R$,
- $\forall x \ 1(x) = x$.

Example 1.18 *Rings and Fields*

Let \mathcal{L}_r be the language of rings $\{+, -, \cdot, 0, 1\}$, where $+$, $-$, and \cdot are binary function symbols and 0 and 1 are constants. The axioms for rings are given by

- the axioms for additive commutative groups,
- $\forall x \forall y \forall z (x - y = z \leftrightarrow x = y + z)$,
- $\forall x \ x \cdot 0 = 0$,
- $\forall x \forall y \forall z (x \cdot (y \cdot z) = (x \cdot y) \cdot z)$,
- $\forall x \ x \cdot 1 = 1 \cdot x = x$,
- $\forall x \forall y \forall z \ x \cdot (y + z) = (x \cdot y) + (x \cdot z)$,
- $\forall x \forall y \forall z (x + y) \cdot z = (x \cdot z) + (y \cdot z)$.

The second axiom is only necessary because we include $-$ in the language (this will be useful later). We axiomatize the class of fields by adding the axioms

- $\forall x \forall y \ x \cdot y = y \cdot x$,
- $\forall x (x \neq 0 \rightarrow \exists y \ x \cdot y = 1)$.

We axiomatize the class of algebraically closed fields by adding to the field axioms the sentences

$$\forall a_0 \dots \forall a_{n-1} \exists x \ x^n + \sum_{i=0}^{n-1} a_i x^i = 0$$

for $n = 1, 2, \dots$. Let ACF be the axioms for algebraically closed fields.

Let ψ_p be the \mathcal{L}_r -sentence $\forall x \underbrace{x + \dots + x}_{p\text{-times}} = 0$, which asserts that a field has characteristic p . For $p > 0$ a prime, let $\text{ACF}_p = \text{ACF} \cup \{\psi_p\}$ and $\text{ACF}_0 = \text{ACF} \cup \{\neg \psi_p : p > 0\}$, be the theories of algebraically closed fields of characteristic p and characteristic zero, respectively.

Example 1.19 *Ordered Fields*

Let $\mathcal{L}_{\text{or}} = \mathcal{L}_r \cup \{<\}$. The class of ordered fields is axiomatized by the axioms for fields,

$$\begin{aligned} & \text{the axioms for linear orders,} \\ & \forall x \forall y \forall z (x < y \rightarrow x + z < y + z), \\ & \forall x \forall y \forall z ((x < y \wedge z > 0) \rightarrow x \cdot z < y \cdot z). \end{aligned}$$

Example 1.20 *Differential Fields*

Let $\mathcal{L} = \mathcal{L}_r \cup \{\delta\}$, where δ is a unary function symbol. The class of differential fields is axiomatized by

$$\begin{aligned} & \text{the axioms of fields,} \\ & \forall x \forall y \delta(x + y) = \delta(x) + \delta(y), \\ & \forall x \forall y \delta(x \cdot y) = x \cdot \delta(y) + y \cdot \delta(x). \end{aligned}$$

Example 1.21 *Peano Arithmetic*

Let $\mathcal{L} = \{+, \cdot, s, 0\}$, where $+$ and \cdot are binary functions, s is a unary function, and 0 is a constant. We think of s as the successor function $x \mapsto x + 1$. The Peano axioms for arithmetic are the sentences

$$\begin{aligned} & \forall x s(x) \neq 0, \\ & \forall x (x \neq 0 \rightarrow \exists y s(y) = x), \\ & \forall x x + 0 = x, \\ & \forall x \forall y x + (s(y)) = s(x + y), \\ & \forall x x \cdot 0 = 0, \\ & \forall x \forall y x \cdot s(y) = (x \cdot y) + x, \end{aligned}$$

and the axioms $\text{Ind}(\phi)$ for each formula $\phi(v, \bar{w})$, where $\text{Ind}(\phi)$ is the sentence $\forall \bar{w} [(\phi(0, \bar{w}) \wedge \forall v (\phi(v, \bar{w}) \rightarrow \phi(s(v), \bar{w}))) \rightarrow \forall x \phi(x, \bar{w})]$.

The axiom $\text{Ind}(\phi)$ formalizes an instance of induction. It asserts that if $\bar{a} \in M$, $X = \{m \in M : \mathcal{M} \models \phi(m, \bar{a})\}$, $0 \in X$, and $s(m) \in X$ whenever $m \in X$, then $X = M$.

Logical Consequence

Definition 1.22 Let T be an \mathcal{L} -theory and ϕ an \mathcal{L} -sentence. We say that ϕ is a *logical consequence* of T and write $T \models \phi$ if $\mathcal{M} \models \phi$ whenever $\mathcal{M} \models T$.

We give two examples.

Proposition 1.23 a) Let $\mathcal{L} = \{+, <, 0\}$ and let T be the theory of ordered Abelian groups. Then, $\forall x (x \neq 0 \rightarrow x + x \neq 0)$ is a logical consequence of T .

b) Let T be the theory of groups where every element has order 2. Then, $T \not\models \exists x_1 \exists x_2 \exists x_3 (x_1 \neq x_2 \wedge x_2 \neq x_3 \wedge x_1 \neq x_3)$.

Proof

a) Suppose that $\mathcal{M} = (M, +, <, 0)$ is an ordered Abelian group. Let $a \in M \setminus \{0\}$. We must show that $a + a \neq 0$. Because $(M, <)$ is a linear order $a < 0$

or $0 < a$. If $a < 0$, then $a + a < 0 + a = a < 0$. Because $\neg(0 < 0)$, $a + a \neq 0$. If $0 < a$, then $0 < a = 0 + a < a + a$ and again $a + a \neq 0$.

b) Clearly, $\mathbb{Z}/2\mathbb{Z} \models T \wedge \neg\exists x_1\exists x_2\exists x_3(x_1 \neq x_2 \wedge x_2 \neq x_3 \wedge x_1 \neq x_3)$.

In general, to show that $T \models \phi$, we give an informal mathematical proof as above that $\mathcal{M} \models \phi$ whenever $\mathcal{M} \models T$. To show that $T \not\models \phi$, we usually construct a counterexample.

Definable Sets

Definition 1.24 Let $\mathcal{M} = (M, \dots)$ be an \mathcal{L} -structure. We say that $X \subseteq M^n$ is *definable* if and only if there is an \mathcal{L} -formula $\phi(v_1, \dots, v_n, w_1, \dots, w_m)$ and $\bar{b} \in M^m$ such that $X = \{\bar{a} \in M^n : \mathcal{M} \models \phi(\bar{a}, \bar{b})\}$. We say that $\phi(\bar{v}, \bar{b})$ *defines* X . We say that X is *A-definable* or *definable over A* if there is a formula $\psi(\bar{v}, w_1, \dots, w_l)$ and $\bar{b} \in A^l$ such that $\psi(\bar{v}, \bar{b})$ defines X .

We give a number of examples using \mathcal{L}_r , the language of rings.

• Let $\mathcal{M} = (R, +, -, \cdot, 0, 1)$ be a ring. Let $p(X) \in R[X]$. Then, $Y = \{x \in R : p(x) = 0\}$ is definable. Suppose that $p(X) = \sum_{i=0}^m a_i X^i$. Let $\phi(v, w_0, \dots, w_n)$ be the formula

$$w_n \cdot \underbrace{v \cdots v}_{n\text{-times}} + \dots + w_1 \cdot v + w_0 = 0$$

(in the future, when no confusion arises, we will abbreviate such a formula as “ $w_n v^n + \dots + w_1 v + w_0 = 0$ ”). Then, $\phi(v, a_0, \dots, a_n)$ defines Y . Indeed, Y is A -definable for any $A \supseteq \{a_0, \dots, a_n\}$.

• Let $\mathcal{M} = (\mathbb{R}, +, -, \cdot, 0, 1)$ be the field of real numbers. Let $\phi(x, y)$ be the formula

$$\exists z(z \neq 0 \wedge y = x + z^2).$$

Because $a < b$ if and only if $\mathcal{M} \models \phi(a, b)$, the ordering is \emptyset -definable.

• Let $\mathcal{M} = (\mathbb{Z}, +, -, \cdot, 0, 1)$ be the ring of integers. Let $X = \{(m, n) \in \mathbb{Z}^2 : m < n\}$. Then, X is definable (indeed \emptyset -definable). By Lagrange’s Theorem, every nonnegative integer is the sum of four squares. Thus, if we let $\phi(x, y)$ be the formula

$$\exists z_1\exists z_2\exists z_3\exists z_4(z_1 \neq 0 \wedge y = x + z_1^2 + z_2^2 + z_3^2 + z_4^2),$$

then $X = \{(m, n) \in \mathbb{Z}^2 : \mathcal{M} \models \phi(m, n)\}$.

• Let F be a field and $\mathcal{M} = (F[X], +, -, \cdot, 0, 1)$ be the ring of polynomials over F . Then F is definable in \mathcal{M} . Indeed, F is the set of units of $F[X]$ and is defined by the formula $x = 0 \vee \exists y xy = 1$.

• Let $\mathcal{M} = (\mathbb{C}(X), +, -, \cdot, 0, 1)$ be the field of complex rational functions in one variable. We claim that \mathbb{C} is defined in $\mathbb{C}(X)$ by the formula

$$\exists x\exists y y^2 = v \wedge x^3 + 1 = v.$$

For any $z \in \mathbb{C}$ we can find x and y such that $y^2 = x^3 + 1 = z$. Suppose that h is a nonconstant rational function and that there are nonconstant rational functions f and g such that $h = g^2 = f^3 + 1$. Then $t \mapsto (f(t), g(t))$ is a nonconstant rational function from an open subset of \mathbb{C} into the curve E given by the equation $y^2 = x^3 + 1$. But E is an elliptic curve and it is known (see for example [?]) that there are no such functions.

A similar argument shows that \mathbb{C} is the set of rational functions f such that f and $f + 1$ are both fourth powers. These ideas generalize to show that \mathbb{C} is definable in any finite algebraic extension of $\mathbb{C}(X)$.

- Let $\mathcal{M} = (\mathbb{Q}_p, +, -, \cdot, 0, 1)$ be the field of p -adic numbers. Then \mathbb{Z}_p the ring of p -adic integers is definable. Suppose $p \neq 2$ (we leave \mathbb{Q}_2 for Exercise ??) and $\phi(x)$ is the formula $\exists y y^2 = px^2 + 1$. We claim that $\phi(x)$ defines \mathbb{Z}_p .

First, suppose that $y^2 = pa^2 + 1$. Let v denote the p -adic valuation. Because $v(pa^2) = 2v(a) + 1$, if $v(a) < 0$, then $v(pa^2)$ is an odd negative integer and $v(y^2) = v(pa^2 + 1) = v(pa^2)$. On the other hand, $v(y^2) = 2v(y)$, an even integer. Thus, if $\mathcal{M} \models \phi(a)$, then $v(a) \geq 0$ so $a \in \mathbb{Z}_p$.

On the other hand, suppose that $a \in \mathbb{Z}_p$. Let $F(X) = X^2 - (pa^2 + 1)$. Let \bar{F} be the reduction of $F \bmod p$. Because $v(a) \geq 0$, $v(pa) > 0$ and $\bar{F}(X) = X^2 - 1$ and $\bar{F}' = 2X$. Thus, $\bar{F}(1) = 0$ and $\bar{F}'(1) \neq 0$ so, by Hensel's Lemma, there is $b \in \mathbb{Z}_p$ such that $F(b) = 0$. Hence $\mathcal{M} \models \phi(a)$.

- Let $\mathcal{M} = (\mathbb{Q}, +, -, \cdot, 0, 1)$ be the field of rational numbers. Let $\phi(x, y, z)$ be the formula

$$\exists a \exists b \exists c xyz^2 + 2 = a^2 + xy^2 - yc^2$$

and let $\psi(x)$ be the formula

$$\forall y \forall z ([\phi(y, z, 0) \wedge (\forall w (\phi(y, z, w) \rightarrow \phi(y, z, w + 1)))] \rightarrow \phi(y, z, x)).$$

A remarkable result of Julia Robinson (see [?]) shows that $\psi(x)$ defines the integers in \mathbb{Q} .

- Consider the natural numbers \mathbb{N} as an $\mathcal{L} = \{+, \cdot, 0, 1\}$ structure. The definable sets are quite complex. For example, there is an \mathcal{L} -formula $T(e, x, s)$ such that $\mathbb{N} \models T(e, x, s)$ if and only if the Turing machine with program coded by e halts on input x in at most s steps (see, for example, [?]). Thus, the Turing machine with program e halts on input x if and only if $\mathbb{N} \models \exists s T(e, x, s)$, so the set of halting computations is definable. It is well known that this set is not computable (see, for example, [?]). This leads to an interesting conclusion.

Proposition 1.25 *The full \mathcal{L} -theory of the natural numbers is undecidable (i.e., there is no algorithm that when given an \mathcal{L} -sentence ψ as input will always halt answering “yes” if $\mathbb{N} \models \psi$ and “no” if $\mathbb{N} \models \neg\psi$).*

Proof For each e and x , let $\phi_{e,x}$ be the \mathcal{L} -sentence

$$\exists s T(\underbrace{1 + \dots + 1}_{e\text{-times}}, \underbrace{1 + \dots + 1}_{x\text{-times}}, s).$$

If there were such an algorithm we could decide whether the program coded by e halts on input x by asking whether $\mathbb{N} \models \phi_{e,x}$.

Recursively enumerable sets have simple mathematical definitions. By the Matijasevič–Robinson–Davis–Putnam solution to Hilbert’s 10th Problem (see [?]) for any recursively enumerable set $A \subseteq \mathbb{N}^n$ there is a polynomial

$$p(X_1, \dots, X_n, Y_1, \dots, Y_m) \in \mathbb{Z}[\overline{X}, \overline{Y}]$$

such that

$$A = \{\overline{x} \in \mathbb{N}^n : \mathbb{N} \models \exists y_1 \dots \exists y_m p(\overline{x}, \overline{y}) = 0\}.$$

The following example will be useful later.

Lemma 1.26 *Let \mathcal{L}_r be the language of ordered rings and $(\mathbb{R}, +, -, \cdot, <, 0, 1)$ be the ordered field of real numbers. Suppose that $X \subseteq \mathbb{R}^n$ is A -definable. Then, the topological closure of X is also A -definable.*

Proof Let $\phi(v_1, \dots, v_n, \overline{a})$ define X . Let $\psi(v_1, \dots, v_n, \overline{w})$ be the formula

$$\forall \epsilon \left[\epsilon > 0 \rightarrow \exists y_1, \dots, y_n (\phi(\overline{y}, \overline{w}) \wedge \sum_{i=1}^n (v_i - y_i)^2 < \epsilon) \right].$$

Then, \overline{b} is in the closure of X if and only if $\mathcal{M} \models \psi(\overline{b}, \overline{a})$.

How do we show that $X \subset M^n$ is not definable? The following proposition will often be useful.

Proposition 1.27 *Let \mathcal{M} be an \mathcal{L} -structure. If $X \subset M^n$ is A -definable, then every \mathcal{L} -automorphism of \mathcal{M} that fixes A pointwise fixes X setwise (that is, if σ is an automorphism of M and $\sigma(a) = a$ for all $a \in A$, then $\sigma(X) = X$).*

Proof Let $\psi(\overline{v}, \overline{a})$ be the \mathcal{L} -formula defining X where $\overline{a} \in A$. Let σ be an automorphism of \mathcal{M} with $\sigma(\overline{a}) = \overline{a}$, and let $\overline{b} \in M^n$.

In the proof of Theorem 1.10, we showed that if $j : \mathcal{M} \rightarrow \mathcal{N}$ is an isomorphism, then $\mathcal{M} \models \phi(\overline{a})$ if and only if $\mathcal{N} \models \phi(j(\overline{a}))$. Thus

$$\mathcal{M} \models \psi(\overline{b}, \overline{a}) \leftrightarrow \mathcal{M} \models \psi(\sigma(\overline{b}), \sigma(\overline{a})) \leftrightarrow \mathcal{M} \models \psi(\sigma(\overline{b}), \overline{a}).$$

In other words, $\overline{b} \in X$ if and only if $\sigma(\overline{b}) \in X$ as desired.

We give a sample application.

Corollary 1.28 *The set of real numbers is not definable in the field of complex numbers.*

Proof If \mathbb{R} were definable, then it would be definable over a finite $A \subset \mathbb{C}$. Let $r, s \in \mathbb{C}$ be algebraically independent over A with $r \in \mathbb{R}$ and $s \notin \mathbb{R}$. There is an automorphism σ of \mathbb{C} such that $\sigma|_A$ is the identity and $\sigma(r) = s$. Thus, $\sigma(\mathbb{R}) \neq \mathbb{R}$ and \mathbb{R} is not definable over A .

This proof worked because \mathbb{C} has many automorphisms. The situation is much different for \mathbb{R} . Any automorphism of the real field must fix the rational numbers. Because the ordering is definable it must be preserved by any automorphism. Because the rationals are dense in \mathbb{R} , the only automorphism of the real field is the identity. Most subsets of \mathbb{R} are undefinable (there are $2^{2^{\aleph_0}}$ subsets of \mathbb{R} and only 2^{\aleph_0} possible definitions), but we cannot use Proposition 1.27 to show any particular set is undefinable. In fact, the converse to Proposition 1.27 holds for sufficiently rich models.

2 The Compactness Theorem

Let T be an \mathcal{L} -theory and ϕ an \mathcal{L} -sentence. To show that $T \models \phi$, we must show that ϕ holds in every model of T . Checking all models of T sounds like a daunting task, but in practice we usually show that $T \models \phi$ by giving an informal mathematical proof that ϕ is true in every model of T . One of the first great achievements of mathematical logic was giving a rigorous definition of “proof” that completely captures the notion of “logical consequence.”

A proof of ϕ from T is a finite sequence of \mathcal{L} -formulas ψ_1, \dots, ψ_m such that $\psi_m = \phi$ and $\psi_i \in T$ or ψ_i follows from $\psi_1, \dots, \psi_{i-1}$ by a simple logical rule for each i . We write $T \vdash \phi$ if there is a proof of ϕ from T . Examples of “simple” logical rules are:

- “from ϕ and ψ conclude $\phi \wedge \psi$,” or
- “from $\phi \wedge \psi$ conclude ϕ .”

It will not be important for our purposes to go into the details of the proof system, but we stress the following points. (See [?], for example, for complete details of one possible proof system.)

- Proofs are finite.
- (Soundness) If $T \vdash \phi$, then $T \models \phi$.
- If T is a finite set of sentences, then there is an algorithm that, when given a sequence of \mathcal{L} -formulas σ and an \mathcal{L} -sentence ϕ , will decide whether σ is a proof of ϕ from T .

Note that the last point does not say that there is an algorithm that will decide if $T \vdash \phi$. It only says that there is an algorithm that can check each purported proof.

We say that a language \mathcal{L} is *recursive* if there is an algorithm that decides whether a sequence of symbols is an \mathcal{L} -formula. We say that an \mathcal{L} -theory T is recursive if there is an algorithm that, when given an \mathcal{L} -sentence ϕ as input, decides whether $\phi \in T$.

Proposition 2.1 *If \mathcal{L} is a recursive language and T is a recursive \mathcal{L} -theory, then $\{\phi : T \vdash \phi\}$ is recursively enumerable; that is, there is an algorithm, that when given ϕ as input will halt accepting if $T \vdash \phi$ and not halt if $T \not\vdash \phi$.*

Proof There is $\sigma_0, \sigma_1, \sigma_2, \dots$, a computable listing of all finite sequences of \mathcal{L} -formulas. At stage i of our algorithm, we check to see whether σ_i is a proof of ψ from T . This involves checking that each formula either is in T (which we can check because T is recursive) or follows by a logical rule from earlier formulas in the sequence σ_i and that the last formula is ϕ . If σ_i is a proof of ϕ from T , then we halt accepting; otherwise we go on to stage $i + 1$.

Remarkably, the finitistic syntactic notion of “proof” completely captures the semantic notion of “logical consequence.”

Theorem 2.2 (Gödel’s Completeness Theorem) *Let T be an \mathcal{L} -theory and ϕ an \mathcal{L} -sentence, then $T \models \phi$ if and only if $T \vdash \phi$.*

The Completeness Theorem gives a criterion for testing whether an \mathcal{L} -theory is satisfiable. We say that an \mathcal{L} -theory T is *inconsistent* if $T \vdash (\phi \wedge \neg\phi)$ for some sentence ϕ ; otherwise we say that T is *consistent*. Because our proof system is sound, any satisfiable theory is consistent. The Completeness Theorem implies that the converse is true.

Corollary 2.3 *T is consistent if and only if T is satisfiable.*

Proof Suppose that T is not satisfiable. Because there are no models of T , every model of T is a model of $(\phi \wedge \neg\phi)$. Thus, $T \models (\phi \wedge \neg\phi)$ and by the Completeness Theorem $T \vdash (\phi \wedge \neg\phi)$.

This has a deceptively simple consequence.

Theorem 2.4 (Compactness Theorem) *T is satisfiable if and only if every finite subset of T is satisfiable.*

Proof Clearly, if T is satisfiable, then every subset of T is satisfiable. On the other hand, if T is not satisfiable, then T is inconsistent. Let σ be a proof of a contradiction from T . Because σ is finite, only finitely many assumptions from T are used in the proof. Thus, there is a finite $T_0 \subseteq T$ such that σ is a proof of a contradiction from T_0 . But then T_0 is a finite unsatisfiable subset of T .

Although it is a simple consequence of the Completeness Theorem and the finite nature of proof, the Compactness Theorem is the cornerstone of model theory. Because it will not be useful for us to understand the exact nature of our proof system, we will not prove the Completeness Theorem. Instead, in the next section, we will give a second proof of the Compactness Theorem that does not appeal directly to the Completeness Theorem.

Basic Applications of Compactness

We conclude this section with several standard applications of the Compactness Theorem.

Corollary 2.5 *Suppose T has arbitrarily large finite models, then T has an infinite model.*

Proof Let ϕ_n be the sentence:

$$\exists v_1 \dots \exists v_n \bigwedge_{i < j \leq n} v_i \neq v_j.$$

Let $T^* = T \cup \{\phi_n : n = 1, 2, \dots\}$. Clearly any model of T^* is an infinite model of T . If $\Delta \subset T^*$ is finite, then for some N , $\Delta \subset T \cup \{\phi_1, \dots, \phi_N\}$. There is $\mathcal{A} \models T$ with $|\mathcal{A}| \geq N$, thus $\mathcal{A} \models \Delta$. By the Compactness Theorem, T^* has a model.

Proposition 2.6 Let $\mathcal{L} = \{\cdot, +, <, 0, 1\}$ and let $\text{Th}(\mathbb{N})$ be the full \mathcal{L} -theory of the natural numbers. There is $\mathcal{M} \models \text{Th}(\mathbb{N})$ and $a \in M$ such that a is larger than every natural number.

Proof Let $\mathcal{L}^* = \mathcal{L} \cup \{c\}$ where c is a new constant symbol and let

$$T = \text{Th}(\mathbb{N}) \cup \left\{ \underbrace{1 + 1 + \dots + 1}_{n\text{-times}} < c : \text{for } n = 1, 2, \dots \right\}.$$

If Δ is a finite subset of T , we can make \mathbb{N} a model of Δ by interpreting c as a suitably large natural number. Thus, T is finitely satisfiable and there is $\mathcal{M} \models T$. If $a \in M$ is the interpretation of c , then a is larger than every natural number.

Proposition 2.7 Let \mathcal{L} be a language containing $\{\cdot, e\}$, the language of groups, let T be an \mathcal{L} -theory extending the theory of groups, and let $\phi(v)$ be an \mathcal{L} -formula. Suppose that for all n there is $G_n \models T$ and $g_n \in G_n$ with finite order greater than n such that $G_n \models \phi(g_n)$. Then, there is $G \models T$ and $g \in G$ such that $G \models \phi(g)$ and g has infinite order. In particular, there is no formula that defines the torsion points in all models of T .

Proof Let $\mathcal{L}^* = \mathcal{L} \cup \{c\}$, where c is a new constant symbol. Let T^* be the \mathcal{L} -theory

$$T \cup \{\phi(c)\} \cup \left\{ \underbrace{c \cdot c \cdots c}_{n\text{-times}} \neq e : n = 1, 2, \dots \right\}.$$

If G is a model of T^* and g is the interpretation of c in G then $G \models \phi(g)$ and g has infinite order. Hence, it suffices to show that T^* is satisfiable.

Let $\Delta \subseteq T^*$ be finite. Then

$$\Delta \subseteq T \cup \{\phi(c)\} \cup \left\{ \underbrace{c \cdot c \cdots c}_{n\text{-times}} \neq e : n = 1, 2, \dots, m \right\}$$

for some m . View G_m as an \mathcal{L}^* structure by interpreting c as the element g_m . Because $G_m \models T \cup \{\phi(g_m)\}$ and g_m has order greater than m , $G_m \models \Delta$. Thus, T^* is finitely satisfiable and hence, by the Compactness Theorem, satisfiable.

Example 2.8 *Four Coloring Graphs*

Let $G = (V, E)$ be a graph such that every finite subgraph can be four colored.² We claim that G can be four colored. Let $\mathcal{L} = \{R, B, Y, G\} \cup \{c_v : v \in V\}$. Let Γ be the \mathcal{L} -theory with axioms:

- i) $\forall x [(R(x) \wedge \neg B(x) \wedge \neg Y(x) \wedge \neg G(x)) \vee \dots \vee (\neg R(x) \wedge \neg B(x) \wedge \neg Y(x) \wedge G(x))]$
- ii) if $(v, w) \in E$ add the axiom: $\neg(R(c_v) \wedge R(c_w)) \wedge \dots \wedge \neg(G(c_v) \wedge G(c_w))$.

²That is, we can color the vertices with four colors so that no adjacent vertices have the same color. For example, the Four Color Theorem says that every finite planar graph can be four colored.

If Δ is a finite subset of Γ , let V_Δ be the vertices such that c_v is used in Δ . Since the restriction of G to V_Δ is four colorable, Δ is consistent. Thus Γ is consistent. Let $\mathcal{A} \models \Gamma$.

Color G by coloring v as \mathcal{A} colors c_v .

Theorem 2.9 [*Upward Löwenheim–Skolem Theorem*] *Suppose Γ is an \mathcal{L} -theory. If Γ has an infinite model, then it has a model of cardinality κ for every $\kappa \geq \max(|\mathcal{L}|, \aleph_0)$.*

Proof Let I be a set of cardinality κ . Let $\mathcal{L}^* = \mathcal{L} \cup \{c_\alpha : \alpha \in I\}$. Let

$$\Gamma^* = \Gamma \cup \{c_\alpha \neq c_\beta : \alpha \neq \beta\}.$$

If Δ is a finite subset of Γ^* , then in any infinite model \mathcal{A} of Γ we can interpret the constants such that $\mathcal{A} \models \Delta$. Thus Γ has a model of size at most κ . But certainly any model of Γ^* has size at least κ (the map $\alpha \mapsto \hat{c}_\alpha$ is one to one).

The next lemma is an easy consequence of the Completeness Theorem, but it also can be deduced from the Compactness Theorem.

Lemma 2.10 *If $T \models \phi$, then $\Delta \models \phi$ for some finite $\Delta \subseteq T$.*

Proof Suppose not. Let $\Delta \subseteq T$ be finite. Because $\Delta \not\models \phi$, $\Delta \cup \{\neg\phi\}$ is satisfiable. Thus, $T \cup \{\neg\phi\}$ is finitely satisfiable and, by the Compactness Theorem, $T \not\models \phi$.

3 Ultraproducts and Compactness

In this section we will give an alternative proof of the Compactness Theorem using ultraproducts, an algebraic method of *averaging* structures.

Let I be an infinite set. We let

$$\mathcal{P}(I) = \{A : A \subseteq I\}$$

be the *power set* of I .

Definition 3.1 We say that $\mathcal{F} \subseteq \mathcal{P}(I)$ is a *filter* if

- i) $I \in \mathcal{F}$, $\emptyset \notin \mathcal{F}$;
- ii) If $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$;
- iii) If $A, B \in \mathcal{F}$ then $A \cap B \in \mathcal{F}$.

We say that \mathcal{F} is an *ultrafilter* if in addition,

- iv) for all $A \subseteq I$ either $A \in \mathcal{F}$ or $I \setminus A \in \mathcal{F}$.

Example 3.2 $\text{Cof} = \{A \subseteq I : I \setminus A \text{ is finite}\}$ is a filter.

Example 3.3 Let $I = \mathbb{R}$ then $\mathcal{F} = \{A : \mathbb{R} \setminus A \text{ has Lebesgue measure zero}\}$, is a filter.

If \mathcal{F} is a filter, we think of elements of \mathcal{F} as *llarge*, so if $A \in \mathcal{F}$ we think of A as large and that $i \in A$ for *almost all* $i \in I$.

We can think of an ultrafilter \mathbb{F} as finitely additive two valued measures $\mu : \mathcal{P}(I) \rightarrow \{0, 1\}$, where $\mu(A) = 1$ if and only if $A \in \mathbb{F}$.

Lemma 3.4 If $\mathcal{F} \subseteq \mathcal{P}(I)$ is a filter, $A \subseteq I$ and $I \setminus A \notin \mathcal{F}$, then

$$\mathcal{F}' = \{C : \text{there is } B \in \mathcal{F}, C \supseteq A \cap B\}$$

is an ultrafilter and $A \in \mathcal{F}'$.

Proof Since $I \supseteq I \cap A$, $I \in \mathcal{F}'$.

If $\emptyset \in \mathcal{F}'$, then there is $B \in \mathcal{F}$ such that $A \cap B = \emptyset$. But then $B \subseteq I \setminus A$ and $I \setminus A \in \mathcal{F}$, a contradiction.

It is easy to see that \mathcal{F}' is closed under superset.

If $C_1, C_2 \in \mathcal{F}'$ there are $B_1, B_2 \in \mathcal{F}$ such that $C_i \supseteq B_i \cap A$. Then $C_1 \cap C_2 \supseteq B_1 \cap B_2 \cap A$, so $C_1 \cap C_2 \in \mathcal{F}'$.

Corollary 3.5 If $\mathcal{F} \subseteq \mathcal{P}(I)$ is a filter, then there is an ultrafilter $\mathcal{U} \supseteq \mathcal{F}$.

Proof Let $\mathcal{I} = \{\mathcal{F}' : \mathcal{F} \subseteq \mathcal{F}' \subseteq \mathcal{P}(I) \text{ is a filter}\}$.

If $(X, <)$ is a linearly ordered set, $\mathcal{F}_x \in \mathcal{I}$ for $x \in X$ and $\mathcal{F}_x \subseteq \mathcal{F}_y$ for $x < y$, then $\mathcal{F}^* = \bigcup_{x \in X} \mathcal{F}_x$ is a filter. Thus we can apply Zorn's Lemma to find $\mathcal{U} \in \mathcal{I}$ maximal. Suppose $A \subseteq I$. If $I \setminus A \notin \mathcal{U}$, then, by the Lemma and the maximality of \mathcal{U} , $A \in \mathcal{U}$.

Corollary 3.6 *There are non-principal ultrafilters.*

Proof Let $\mathcal{U} \supseteq \text{Cof}$ be an ultrafilter. Then \mathcal{U} contains no finite sets.

Our proof of the existence of non-principal ultrafilters is non-constructive as it depends heavily on the Axiom of Choice. Unfortunately, some use of choice is unavoidable.

We will use ultrafilters to give a new construction of models. Let \mathcal{L} be a first order language. Suppose that \mathcal{M}_i is an \mathcal{L} -structure for all $i \in I$ with universe M_i . Let $\mathcal{U} \subseteq \mathcal{P}(I)$ be an ultrafilter.

We define \sim on $\prod_{i \in I} M_i$ by

$$f \sim g \Leftrightarrow \{i \in I : f(i) = g(i)\} \in \mathcal{U}.$$

Lemma 3.7 *\sim is an equivalence relation*

Proof Let $f, g, h \in \prod_{i \in I} M_i$. Clearly $f \sim f$ and if $f \sim g$, then $g \sim f$. Suppose $f \sim g$ and $g \sim h$. Since

$$\{i : f(i) = h(i)\} \supseteq \{i : f(i) = g(i)\} \cap \{i : g(i) = h(i)\} \in \mathcal{U},$$

$f \sim h$.

For $f \in \prod_{i \in I} M_i$, let $[f]$ be the \sim -equivalence class of f and let

$$M = \left\{ [f] : f \in \prod_{i \in I} M_i \right\}.$$

We will interpret the symbols of \mathcal{L} in M to construct an \mathcal{L} -structure \mathcal{M} , which we also denote $\prod M_i / \mathcal{U}$.

If c is a constant symbol of \mathcal{L} , let $f \in \prod M_i$ be the function $f(i) = c^{\mathcal{M}_i}$ and let $c^{\mathcal{M}} = [f]$.

Let R be an n -ary relation symbol of \mathcal{L} .

Lemma 3.8 *$f_1, \dots, f_n, g_1, \dots, g_n \in \prod M_i$ such that $f_j \sim g_j$ for all $j = 1, \dots, n$. Then*

$$\{i \in I : (f_1(i), \dots, f_n(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{U} \Leftrightarrow \{i \in I : (g_1(i), \dots, g_n(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{U}.$$

Proof Suppose $\{i \in I : (f_1(i), \dots, f_n(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{U}$. Then $\{i \in I : (g_1(i), \dots, g_n(i)) \in R^{\mathcal{M}_i}\}$ contains

$$\{i \in I : (f_1(i), \dots, f_n(i)) \in R^{\mathcal{M}_i}\} \cap \{i \in I : g_1(i) = f_1(i)\} \cap \dots \cap \{i \in I : g_n(i) = f_n(i)\}.$$

Since \mathcal{U} is a filter this later set is in \mathcal{U} .

The other direction is symmetric.

We define

$$R^{\mathcal{M}} = \{([f_1], \dots, [f_n]) : \{i \in I : (f_1(i), \dots, f_n(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{U}\}.$$

By the Lemma, this is well-defined and does not depend on the choice of representatives for the equivalence classes.

Let F be an n -ary function symbol of \mathcal{L} . Let $f_1, \dots, f_n, g_1, \dots, g_n \in \prod M_i$ with $f_j \sim g_j$ for $j = 1, \dots, n$. Define $f_{n+1}, g_{n+1} \in \prod M_i$ by

$$f_{n+1}(i) = F(f_1(i), \dots, f_n(i)) \text{ and } g_{n+1}(i) = F(g_1(i), \dots, g_n(i)).$$

Exercise 3.9 Argue as in Lemma 3.8 that $f_{n+1} \sim g_{n+1}$.

We define $F^{\mathcal{M}} : M^n \rightarrow M$ by

$$F([f_1], \dots, [f_n]) = [g]$$

where $g(i) = F(f_1(i), \dots, f_n(i))$. By Exercise 3.9 this is well defined and does not depend on choice of representatives.

We have now completely defined the structure $\mathcal{M} = \prod M_i / U$. We call \mathcal{M} an *ultraproduct* of $(\mathcal{M}_i : i \in I)$

The following exercise is an easy induction on terms.

Exercise 3.10 If t is an \mathcal{L} -term, then $t^{\mathcal{M}}([f_1], \dots, [f_n]) = [g]$ where $g(i) = t^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))$.

We can now state the Fundamental Theorem of Ultraproducts.

Theorem 3.11 (Łos's Theorem) *Let $\phi(v_1, \dots, v_n)$ be any \mathcal{L} -formula. Then*

$$\mathcal{M} \models \phi([f_1], \dots, [f_n]) \Leftrightarrow \{i : \mathcal{M}_i \models \phi(f_1(i), \dots, f_n(i))\} \in \mathcal{U}.$$

Proof We prove this by induction on complexity of formulas

1) Suppose ϕ is $t_1 = t_2$ where t_1 and t_2 are terms.

Define $g_j(i) = t_j^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))$. Then

$$\begin{aligned} \mathcal{M} \models t_1([f_1], \dots, [f_n]) = t_2([f_1], \dots, [f_n]) &\Leftrightarrow [g_1] = [g_2] \\ &\Leftrightarrow \{i : t_1^{\mathcal{M}_i}(f_1(i), \dots, f_n(i)) = t_2^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))\} \in \mathcal{U} \end{aligned}$$

as desired.

2) Suppose ϕ is $R(t_1, \dots, t_m)$.

For $j = 1, \dots, m$ let $g_j(i) = t_j^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))$. Then

$$\begin{aligned} \mathcal{M} \models \phi([f_1], \dots, [f_n]) &\Leftrightarrow \{i : (g_1(i), \dots, g_m(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{U} \\ &\Leftrightarrow \{i : \mathcal{M}_i \models \phi(f_1(i), \dots, f_n(i))\} \in \mathcal{U} \end{aligned}$$

3) Suppose the theorem is true for θ and ψ , and ϕ is $\theta \wedge \psi$. (We suppress the parameters $[f_1], \dots, [f_n]$)

Then

$$\begin{aligned} \mathcal{M} \models \phi &\Leftrightarrow \mathcal{M} \models \psi \text{ and } \mathcal{M} \models \theta \\ &\Leftrightarrow \{i : \mathcal{M}_i \models \psi\} \in \mathcal{U} \text{ and } \{i : \mathcal{M}_i \models \theta\} \in \mathcal{U} \end{aligned}$$

$$\Leftrightarrow \{i : \mathcal{M}_i \models \psi \wedge \theta\} \in \mathcal{U}$$

4) Suppose the theorem is true for ψ and ϕ is $\neg\psi$ Then

$$\begin{aligned} \mathcal{M} \models \phi &\Leftrightarrow \mathcal{M} \not\models \psi \\ &\Leftrightarrow \{i : \mathcal{M}_i \models \psi\} \notin \mathcal{U} \\ &\Leftrightarrow \{i : \mathcal{M}_i \models \neg\psi\} \in \mathcal{U} \end{aligned}$$

5) Suppose the theorem is true for $\psi(v)$ and ϕ is $\exists v \psi(v)$.

If $\mathcal{M} \models \exists v \psi(v)$, then there is g such that $\mathcal{M} \models \psi([g])$. But then

$$\{i : \mathcal{M}_i \models \exists v \psi(v)\} \supseteq \{i : \mathcal{M}_i \models \psi(g(i))\} \in \mathcal{U}$$

On the other hand if $A = \{i : \mathcal{M}_i \models \exists v \psi(v)\} \in \mathcal{U}$ define $g \in \prod M_i$ such that $\mathcal{M}_i \models \psi(g(i))$ for all $i \in A$. Then $\mathcal{M} \models \psi([g])$, so $\mathcal{M} \models \phi$.

Note that step 4) is the only place in the construction that we used that \mathcal{U} is an ultrafilter rather than just a filter.

Exercise 3.12 Let \mathcal{U} be a non-principal ultrafilter on the set of prime numbers. For each prime p , let $\mathbb{F}_p^{\text{alg}}$ be the algebraic closure of \mathbb{F}_p the field with p elements. Prove that $\prod \mathbb{F}_p^{\text{alg}}/\mathcal{U}$ is an algebraically closed field of characteristic 0.

Another Proof of Compactness

We can use Los's Theorem to give a proof of the Compactness Theorem that avoids the Completeness Theorem.

Let Γ be an \mathcal{L} -theory such that every finite $\Delta \subseteq \Gamma$ has a model. Let I be the collection of finite subsets of Γ .

For $\phi \in \Gamma$ let

$$X_\phi = \{\Delta \in I : \Delta \models \phi\}$$

and let

$$\mathcal{F} = \{Y \subseteq I : X_\phi \subseteq Y \text{ for some } \phi \in \Gamma\}.$$

We claim that \mathcal{F} is a filter. It is easy to see that $I \in \mathcal{F}$, $\emptyset \notin \mathcal{F}$ and \mathcal{F} is closed under superset. Also if $Y_1, Y_2 \in \mathcal{F}$ there are ϕ_1, ϕ_2 such that $X_{\phi_i} \subseteq Y_i$. Then $X_{\phi_1 \wedge \phi_2} = X_{\phi_1} \cap X_{\phi_2}$, so

$$X_{\phi_1 \wedge \phi_2} \subseteq Y_1 \cap Y_2$$

and $Y_1 \cap Y_2 \in \mathcal{F}$

Let $\mathcal{U} \supseteq \mathcal{F}$ be an ultrafilter. For $\Delta \in I$, let $\mathcal{M}_\Delta \models \Delta$ and let $\mathcal{M} = \prod \mathcal{M}_\Delta/\mathcal{U}$. Since $X_\phi \in \mathcal{U}$ for all $\phi \in \Gamma$, by Los's Theorem $\mathcal{M} \models \Gamma$.

Ultrapowers and Elementary Extensions

Fix \mathcal{M} and \mathcal{L} structure and let \mathcal{U} be an ultrafilter on an infinite set I . An interesting special case of the ultraproduct construction is when we take all of the $\mathcal{M}_i = \mathcal{M}$. In this case we let $\mathcal{M}^* = \mathcal{M}^I/\mathcal{U}$.

Exercise 3.13 Prove that if \mathcal{M} is finite or \mathcal{U} is principal, then $\mathcal{M} \cong \mathcal{M}^*$.

For each $a \in M$, let $f_a : I \rightarrow M$ be the constant function $f_a(i) = a$. If $a \neq b$, then $[f_a] \neq [f_b]$. By Los's Theorem if $a_1, \dots, a_n \in M$ and ϕ is an \mathcal{L} -formula, then

$$\mathcal{M} \models \phi(a_1, \dots, a_n) \Leftrightarrow \mathcal{M}^* \models \phi([f_{a_1}], \dots, [f_{a_n}])$$

Identifying \mathcal{M} and its image under the embedding $a \mapsto [f_a]$ we can think of \mathcal{M} as substructure of \mathcal{M}^* . Then for $a_1, \dots, a_n \in M$.

$$\mathcal{M} \models \phi(a_1, \dots, a_n) \Leftrightarrow \mathcal{M}^* \models \phi(a_1, \dots, a_n).$$

Definition 3.14 If $\mathcal{M} \subseteq \mathcal{N}$ we say that \mathcal{N} is an *elementary extension* of \mathcal{M} and write $\mathcal{M} \prec \mathcal{N}$ if

$$\mathcal{M} \models \phi(\bar{a}) \Leftrightarrow \mathcal{N} \models \phi(\bar{a})$$

for all $\bar{a} \in M$.

We have argued that \mathcal{M}^* is an elementary extension of \mathcal{M} . This is only interesting if we can also prove \mathcal{M}^* properly extends \mathcal{M} .

Proposition 3.15 If $|I| \leq |\mathcal{M}|$ and \mathcal{U} is a non-principal ultrafilter, then \mathcal{M}^* is a proper extension of \mathcal{M} .

Proof Let $f : I \rightarrow M$ be injective. Then for all $a \in M$, $|\{i : f(i) = a\}| \leq 1$. Since \mathcal{U} is non-principal, $f \not\sim f_a$. Thus $[f] \in \mathcal{M}^* \setminus M$.

Cardinalities of Ultraproducts

Suppose we have $(\mathcal{M}_i : i \in I)$ and an ultrafilter $\mathcal{U} \subseteq \mathcal{P}(I)$.

Exercise 3.16 Suppose $\{i \in I : |\mathcal{M}_i| = n\} \in \mathcal{U}$, then $|\prod \mathcal{M}_i/\mathcal{U}| = n$

Exercise 3.17 If we also have $(\mathcal{N}_i : i \in I)$ and $\{i : |\mathcal{M}_i| = |\mathcal{N}_i|\} \in \mathcal{U}$, then $|\prod \mathcal{M}_i/\mathcal{U}| = |\prod \mathcal{N}_i/\mathcal{U}|$.

Exercise 3.18 If $\lambda \leq |\mathcal{M}_i| \leq \kappa$ for all $i \in I$, then

$$\lambda \leq \prod \mathcal{M}_i/\mathcal{U} \leq \kappa^{|I|}.$$

For the rest of these Exercises we will assume $I = \mathbb{N}$.

Exercise 3.19 Suppose that for all $n \in \mathbb{N}$, $\{i : |\mathcal{M}_i| = n\} \notin \mathcal{U}$ and \mathcal{U} is non-principal.

a) Show there is a family X of functions $f : \mathbb{N} \rightarrow \mathbb{N}$ such that:

i) $|X| = 2^{\aleph_0}$

ii) for each $f \in X$ $f(n) < 2^n$

iii) $f \neq g \in X$, then $\{n : f(n) = g(n)\}$ is finite.

[Hint: For $\alpha : \mathbb{N} \rightarrow \{0, 1\}$ let $f_\alpha(n) = \sum_{i=0}^{n-1} \alpha(i)2^i$.

b) Show there is a partition $I = \bigcup_{n=0}^{\infty} A_n$ such that

i) each $A_n \notin \mathcal{U}$

ii) if $i \in A_n$, then $|\mathcal{M}_i| \geq 2^i$.

[Hint: Let $A_n = \{i : 2^n \leq |\mathcal{M}_i| < 2^{n+1} \text{ or } i = n \text{ and } |\mathcal{M}_i| \geq \aleph_0\}$.]

For $i \in I$ let $n(i)$ be unique such that $i \in A_{n(i)}$. For $i \in I$ choose $(m_{i,j} : 0 \leq j < 2^{n(i)})$ distinct elements of M_i . For $f \in X$, let $\alpha_f \in \prod M_i$ such that $\alpha_f(i) = m_{i,f(n(i))}$.

c) Prove that if $f \neq g \in X$, then $\alpha_f \not\sim \alpha_g$. Conclude that $|\prod \mathcal{M}_i/U| \geq 2^{\aleph_0}$.

Corollary 3.20 *Suppose that \mathcal{U} is a non-principal ultrafilter on \mathbb{N} , $|\mathcal{M}_n| \leq \aleph_0$ for all n , and $\{n : |\mathcal{M}_n| = m\} \notin \mathcal{U}$ for any m . Then $|\prod \mathcal{M}_i/U| = 2^{\aleph_0}$.*

Exercise 3.21 Let \mathcal{U} be a non-principal ultrafilter on the set of primes. Prove $\prod \mathbb{F}_p^{\text{alg}}/\mathcal{U}$ is isomorphic to \mathbb{C} the field of complex numbers.

4 Complete Theories

Definition 4.1 A satisfiable theory T is *complete* if $T \models \phi$ or $T \models \neg\phi$ for all \mathcal{L} -sentences ϕ .

It is easy to see that T is complete if and only if $\mathcal{M} \equiv \mathcal{N}$ for any $\mathcal{M}, \mathcal{N} \models T$. If \mathcal{M} is an \mathcal{L} -structure, then $\text{Th}(\mathcal{M})$ is a complete theory, but it may be difficult to figure out if $\phi \in \text{Th}(\mathcal{M})$.

When we are trying to understand $\text{Th}(\mathcal{M})$ for a particular structure \mathcal{M} we will often do this by looking for easy to understand theory T such that $\mathcal{M} \models T$ and T is complete. If $T \models \phi$, then $\mathcal{M} \models \phi$. On the other hand, if $T \not\models \phi$, then, since T is complete, $T \models \neg\phi$ and, as before, $\mathcal{M} \models \neg\phi$ so $\mathcal{M} \not\models \phi$. Thus we would have

$$\mathcal{M} \models \phi \Leftrightarrow T \models \phi$$

In this section, will give one useful test to decide if a theory is complete.

Categoricity

We know from the Löwenheim-Skolem Theorem (Theorem 2.9) that if a theory has an infinite model it has arbitrarily large models. Thus the theory of an infinite structure can not capture the structure up to isomorphism. Sometimes though knowing the theory and the cardinality determines the structure.

Definition 4.2 T is κ -categorical if and only if any two models of T of cardinality κ are isomorphic.

- Let \mathcal{L} be the empty language. Then the theory of an infinite set is κ -categorical for all cardinals κ .

- Let $\mathcal{L} = \{E\}$, where E is a binary relation, and let T be the theory of an equivalence relation with exactly two classes, both of which are infinite. It is easy to see that any two countable models of T are isomorphic. On the other hand, T is not κ -categorical for $\kappa > \aleph_0$. To see this, let \mathcal{M}_0 be a model where both classes have cardinality κ , and let \mathcal{M}_1 be a model where one class has cardinality κ and the other has cardinality \aleph_0 . Clearly, \mathcal{M}_0 and \mathcal{M}_1 are not isomorphic.

Let $\mathcal{L} = \{+, 0\}$ be the language of additive groups and let T be the \mathcal{L} -theory of nontrivial torsion-free divisible Abelian groups. The axioms of T are the axioms for Abelian groups together with the axioms

$$\begin{aligned} & \exists x \ x \neq 0, \\ & \forall x (x \neq 0 \rightarrow \underbrace{x + \dots + x}_{n\text{-times}} \neq 0) \end{aligned}$$

and

$$\forall y \exists x \ \underbrace{x + \dots + x}_{n\text{-times}} = y$$

for $n = 1, 2, \dots$

Proposition 4.3 *The theory of torsion-free divisible Abelian groups is κ -categorical for all $\kappa > \aleph_0$.*

Proof We first argue that models of T are essentially vector spaces over the field of rational numbers \mathbb{Q} . Clearly, if V is any vector space over \mathbb{Q} , then the underlying additive group of V is a model of T . On the other hand, if $G \models T$, $g \in G$, and $n \in \mathbb{N}$ with $n > 0$, we can find $h \in G$ such that $nh = g$. If $nk = g$, then $n(h - k) = 0$. Because G is torsion-free there is a unique $h \in G$ such that $nh = g$. We call this element g/n . We can view G as a \mathbb{Q} -vector space under the action $\frac{m}{n}g = m(g/n)$.

Two \mathbb{Q} -vector spaces are isomorphic if and only if they have the same dimension. Thus, models of T are determined up to isomorphism by their dimension. If G has dimension λ , then $|G| = \lambda + \aleph_0$. If κ is uncountable and G has cardinality κ , then G has dimension κ . Thus, for $\kappa > \aleph_0$ any two models of T of cardinality κ are isomorphic.

Note that T is not \aleph_0 -categorical. Indeed, there are \aleph_0 nonisomorphic models corresponding to vector spaces of dimension $1, 2, 3, \dots$ and \aleph_0 .

A similar argument applies to the theory of algebraically closed fields. Let ACF_p be the theory of algebraically closed fields of characteristic p , where p is either 0 or a prime number.

Proposition 4.4 *ACF_p is κ -categorical for all uncountable cardinals κ .*

Proof Two algebraically closed fields are isomorphic if and only if they have the same characteristic and transcendence degree (see, for example Lang's *Algebra* X §1). An algebraically closed field of transcendence degree λ has cardinality $\lambda + \aleph_0$. If $\kappa > \aleph_0$, an algebraically closed field of cardinality κ also has transcendence degree κ . Thus, any two algebraically closed fields of the same characteristic and same uncountable cardinality are isomorphic.

Vaught's Test

Categoricity give a very simple test for completeness.

Theorem 4.5 (Vaught's Test) *Suppose every model of T is infinite, $\kappa \geq \max(|\mathcal{L}|, \aleph_0)$ and T is κ -categorical. Then T is complete.*

Proof Suppose not. Let ϕ be an \mathcal{L} -sentence such that $T \not\models \phi$ and $T \not\models \neg\phi$. Let $T_0 = T \cup \{\phi\}$ and $T_1 = T \cup \{\neg\phi\}$. Each T_i has a model, thus since T has only infinite models, each T_i has an infinite model. By the Löwenheim-Skolem theorem there is $\mathcal{A}_i \models T_i$ where \mathcal{A}_i has cardinality κ . Since T is κ -categorical, $\mathcal{A}_0 \cong \mathcal{A}_1$ and hence by 1.10, $\mathcal{A}_0 \equiv \mathcal{A}_1$. But $\mathcal{A}_0 \models \phi$ and $\mathcal{A}_1 \models \neg\phi$, a contradiction.

The assumption that T has no finite models is necessary. Suppose that T is the $\{+, 0\}$ -theory of Abelian groups, where every element has order 2.

Exercise 4.6 Show that T is κ -categorical for all $\kappa \geq \aleph_0$. [Hint: Models of T

are essentially vector spaces over \mathbb{F}_2 .]

However, T is not complete. The sentence $\exists x \exists y \exists z (x \neq y \wedge y \neq z \wedge z \neq x)$ is false in the two-element group but true in every other model of T .

Vaught's Test implies that all of the categorical theories discussed above are complete. In particular, the theory of algebraically closed fields of a fixed characteristic is complete. This result of Tarski has several immediate interesting consequences.

Definition 4.7 We say that an \mathcal{L} -theory T is *decidable* if there is an algorithm that when given an \mathcal{L} -sentence ϕ as input decides whether $T \models \phi$.

Lemma 4.8 *Let T be a recursive complete satisfiable theory in a recursive language \mathcal{L} . Then T is decidable.*

Proof Start enumerating all finite sequence of strings of \mathcal{L} -symbols. For each one, check to see if it is a derivation of $\Delta \vdash \phi$ or $\Delta \vdash \neg\phi$. If it is then check to see if all of the sentences in Δ are in T . If so output yes if $\Delta \vdash \phi$ and no if $\Delta \vdash \neg\phi$. If not, go on to the next string. Since T is complete, the Completeness Theorem implies there is a finite $\Delta \subseteq T$ such that $\Delta \vdash \phi$ or $\Delta \vdash \neg\phi$. Thus our search will halt at some stage.

Informally, to decide whether ϕ is a logical consequence of a complete satisfiable recursive theory T , we begin searching through possible proofs from T until we find either a proof of ϕ or a proof of $\neg\phi$. Because T is satisfiable, we will not find proofs of both. Because T is complete, we will eventually find a proof of one or the other.

Corollary 4.9 *For $p = 0$ or p prime, ACF_p is decidable. In particular, $\text{Th}(\mathbb{C})$, the first-order theory of the field of complex numbers, is decidable.*

The completeness of ACF_p can also be thought of as a first-order version of the Lefschetz Principle from algebraic geometry.

Corollary 4.10 *Let ϕ be a sentence in the language of rings. The following are equivalent.*

- i) ϕ is true in the complex numbers.
- ii) ϕ is true in every algebraically closed field of characteristic zero.
- iii) ϕ is true in some algebraically closed field of characteristic zero.
- iv) There are arbitrarily large primes p such that ϕ is true in some algebraically closed field of characteristic p .
- v) There is an m such that for all $p > m$, ϕ is true in all algebraically closed fields of characteristic p .

Proof The equivalence of i)–iii) is just the completeness of ACF_0 and v) \Rightarrow iv) is obvious.

For ii) \Rightarrow v) suppose that $\text{ACF}_0 \models \phi$. There is a finite $\Delta \subset \text{ACF}_0$ such that $\Delta \vdash \phi$. Thus, if we choose p large enough, then $\text{ACF}_p \models \Delta$. Thus, $\text{ACF}_p \models \phi$ for all sufficiently large primes p .

For iv) \Rightarrow ii) suppose $\text{ACF}_0 \not\models \phi$. Because ACF_0 is complete, $\text{ACF}_0 \models \neg\phi$. By the argument above, $\text{ACF}_p \models \neg\phi$ for sufficiently large p ; thus, iv) fails.

Ax found the following striking application of Corollary 4.10.

Theorem 4.11 (Ax) *Every injective polynomial map from \mathbb{C}^n to \mathbb{C}^n is surjective.*

Proof Remarkably, the key to the proof is the simple observation that if k is a finite field, then every injective function $f : k^n \rightarrow k^n$ is surjective. From this observation it is easy to show that the same is true for $\mathbb{F}_p^{\text{alg}}$, the algebraic closure of the p -element field.

Claim Every injective polynomial map $f : (\mathbb{F}_p^{\text{alg}})^n \rightarrow (\mathbb{F}_p^{\text{alg}})^n$ is surjective.

Suppose not. Let $\bar{a} \in \mathbb{F}_p^{\text{alg}}$ be the coefficients of f and let $\bar{b} \in (\mathbb{F}_p^{\text{alg}})^n$ such that \bar{b} is not in the range of f . Let k be the subfield of $\mathbb{F}_p^{\text{alg}}$ generated by \bar{a}, \bar{b} . Then $f|_{k^n}$ is an injective but not surjective polynomial map from k^n into itself. But $\mathbb{F}_p^{\text{alg}} = \bigcup_{n=1}^{\infty} \mathbb{F}_{p^n}$ is a locally finite field. Thus k is finite, a contradiction.

Suppose that the theorem is false. Let $X = (X_1, \dots, X_n)$. Let

$$f(X) = (f_1(X), \dots, f_n(X))$$

be a counterexample where each $f_i \in \mathbb{C}[X]$ has degree at most d . There is an \mathcal{L} -sentence $\Phi_{n,d}$ such that for K a field, $K \models \Phi_{n,d}$ if and only if every injective polynomial map from K^n to K^n where each coordinate function has degree at most d is surjective. We can quantify over polynomials of degree at most d by quantifying over the coefficients. For example, $\Phi_{2,2}$ is the sentence

$$\forall a_{0,0} \forall a_{0,1} \forall a_{0,2} \forall a_{1,0} \forall a_{1,1} \forall a_{2,0} \forall b_{0,0} \forall b_{0,1} \forall b_{0,2} \forall b_{1,0} \forall b_{1,1} \forall b_{2,0} \\ \left[(\forall x_1 \forall y_1 \forall x_2 \forall y_2 ((\sum a_{i,j} x_1^i y_1^j = \sum a_{i,j} x_2^i y_2^j \wedge \sum b_{i,j} x_1^i y_1^j = \sum b_{i,j} x_2^i y_2^j) \rightarrow \right. \\ \left. (x_1 = x_2 \wedge y_1 = y_2))) \rightarrow \forall u \forall v \exists x \exists y \sum a_{i,j} x^i y^j = u \wedge \sum b_{i,j} x^i y^j = v \right].$$

By the claim $\mathbb{F}_p^{\text{alg}} \models \Phi_{n,d}$ for all primes p . By Corollary 4.10, $\mathbb{C} \models \Phi_{n,d}$, a contradiction.

We will return to the model theory of algebraically closed fields in §6.

There are other interesting applications of Vaught's Test. Let $\mathbf{L} = \{<\}$ and let DLO be the theory says we have a dense linear order with no top or bottom element. Then $\mathbb{Q} \models \text{DLO}$ and $\mathbb{R} \models \text{DLO}$.

Theorem 4.12 (Cantor) *Any two countable models of DLO are isomorphic. Thus DLO is \aleph_0 -categorical. Since DLO has no finite models it is complete.*

It follows the $(\mathbb{R}, <) \equiv (\mathbb{Q}, <)$. Thus we can not express the fact that \mathbb{R} is complete. DLO is not κ -categorical for any uncountable cardinal κ . Indeed, if κ is uncountable there are 2^κ non-isomorphic models of cardinality κ .

5 Quantifier Elimination

In model theory we try to understand structures by studying their definable sets. Recall that if \mathcal{M} is an \mathcal{L} -structure, then $X \subseteq M^n$ is *definable* if there is an \mathcal{L} -formula $\phi(v_1, \dots, v_n, w_1, \dots, w_m)$ and $b_1, \dots, b_m \in M$ such that

$$X = \{\bar{a} \in M^n : \mathcal{M} \models \phi(\bar{a}, \bar{b})\}.$$

The study of definable sets is often complicated by quantifiers. For example, in the structure $(\mathbb{N}, +, \cdot, <, 0, 1)$ the quantifier-free definable sets are defined by polynomial equations and inequalities. Even if we use only existential quantifiers the definable sets become complicated. By the Matijasevič–Robinson–Davis–Putnam solution to Hilbert’s 10th problem every recursively enumerable subset of \mathbb{N} is defined by a formula

$$\exists v_1 \dots \exists v_n p(x, v_1, \dots, v_n) = 0$$

for some polynomial $p \in \mathbb{N}[X, Y_1, \dots, Y_n]$. As we allow more alternations of quantifiers, we get even more complicated definable sets.

Not surprisingly, it will be easiest to study definable sets that are defined by quantifier-free formulas. Sometimes formulas with quantifiers can be shown to be equivalent to formulas without quantifiers. Here are two well-known examples. Let $\phi(a, b, c)$ be the formula

$$\exists x ax^2 + bx + c = 0.$$

By the quadratic formula,

$$\mathbb{R} \models \phi(a, b, c) \leftrightarrow [(a \neq 0 \wedge b^2 - 4ac \geq 0) \vee (a = 0 \wedge (b \neq 0 \vee c = 0))],$$

whereas in the complex numbers

$$\mathbb{C} \models \phi(a, b, c) \leftrightarrow (a \neq 0 \vee b \neq 0 \vee c = 0).$$

In either case, ϕ is equivalent to a quantifier-free formula. However, ϕ is not equivalent to a quantifier-free formula over the rational numbers \mathbb{Q} .

For a second example, let $\phi(a, b, c, d)$ be the formula

$$\exists x \exists y \exists u \exists v (xa + yc = 1 \wedge xb + yd = 0 \wedge ua + vc = 0 \wedge ub + vd = 1).$$

The formula $\phi(a, b, c, d)$ asserts that the matrix

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

is invertible. By the determinant test,

$$F \models \phi(a, b, c, d) \leftrightarrow ad - bc \neq 0$$

for any field F .

Definition 5.1 We say that a theory T has *quantifier elimination* if for every formula ϕ there is a quantifier-free formula ψ such that

$$T \models \phi \leftrightarrow \psi.$$

Our goal in this section is to give a very useful model theoretic test for elimination of quantifiers. In the next section we will show that this method can be applied to the theory of algebraically closed fields and develop some rich consequences. We begin by introducing some preliminary tools.

Diagrams

We begin by giving a way to construct \mathcal{L} -embeddings.

Definition 5.2 Suppose that \mathcal{M} is an \mathcal{L} -structure. Let \mathcal{L}_M be the language where we add to \mathcal{L} constant symbols m for each element of M . The *atomic diagram* of \mathcal{M} is $\{\phi(m_1, \dots, m_n) : \phi \text{ is either an atomic } \mathcal{L}\text{-formula or the negation of an atomic } \mathcal{L}\text{-formula and } \mathcal{M} \models \phi(m_1, \dots, m_n)\}$. We let $\text{Diag}(\mathcal{M})$ denote the atomic diagram of \mathcal{M}

Lemma 5.3 *Suppose that \mathcal{N} is an \mathcal{L}_M -structure and $\mathcal{N} \models \text{Diag}(\mathcal{M})$; then, viewing \mathcal{N} as an \mathcal{L} -structure, there is an \mathcal{L} -embedding of \mathcal{M} into \mathcal{N} .*

Proof Let $j : M \rightarrow N$ be defined by $j(m) = m^{\mathcal{N}}$; that is, $j(m)$ is the interpretation of this constant symbol m in \mathcal{N} . If m_1, m_2 are distinct elements of M , then $m_1 \neq m_2 \in \text{Diag}(\mathcal{M})$; thus, $j(m_1) \neq j(m_2)$ so j is an embedding. If f is a function symbol of \mathcal{L} and $f^{\mathcal{M}}(m_1, \dots, m_n) = m_{n+1}$, then $f(m_1, \dots, m_n) = m_{n+1}$ is a formula in $\text{Diag}(\mathcal{M})$ and $f^{\mathcal{N}}(j(m_1), \dots, j(m_n)) = j(m_{n+1})$. If R is a relation symbol and $\bar{m} \in R^{\mathcal{M}}$, then $R(m_1, \dots, m_n) \in \text{Diag}(\mathcal{M})$ and $(j(m_1), \dots, j(m_n)) \in R^{\mathcal{N}}$. Hence, j is an \mathcal{L} -embedding.

Quantifier Elimination Tests

Theorem 5.4 *Suppose that \mathcal{L} contains a constant symbol c , T is an \mathcal{L} -theory, and $\phi(\bar{v})$ is an \mathcal{L} -formula. The following are equivalent:*

- i) There is a quantifier-free \mathcal{L} -formula $\psi(\bar{v})$ such that $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \psi(\bar{v}))$.*
- ii) If \mathcal{M} and \mathcal{N} are models of T , \mathcal{A} is an \mathcal{L} -structure, $\mathcal{A} \subseteq \mathcal{M}$, and $\mathcal{A} \subseteq \mathcal{N}$, then $\mathcal{M} \models \phi(\bar{a})$ if and only if $\mathcal{N} \models \phi(\bar{a})$ for all $\bar{a} \in \mathcal{A}$.*

Proof i) \Rightarrow ii) Suppose that $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \psi(\bar{v}))$, where ψ is quantifier-free. Let $\bar{a} \in \mathcal{A}$, where \mathcal{A} is a common substructure of \mathcal{M} and \mathcal{N} and the latter two structures are models of T . In Proposition 1.8, we saw that quantifier-free formulas are preserved under substructure and extension. Thus

$$\begin{aligned} \mathcal{M} \models \phi(\bar{a}) &\Leftrightarrow \mathcal{M} \models \psi(\bar{a}) \\ &\Leftrightarrow \mathcal{A} \models \psi(\bar{a}) \quad (\text{because } \mathcal{A} \subseteq \mathcal{M}) \\ &\Leftrightarrow \mathcal{N} \models \psi(\bar{a}) \quad (\text{because } \mathcal{A} \subseteq \mathcal{N}) \\ &\Leftrightarrow \mathcal{N} \models \phi(\bar{a}). \end{aligned}$$

ii) \Rightarrow i) First, if $T \models \forall \bar{v} \phi(\bar{v})$, then $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow c = c)$. Second, if $T \models \forall \bar{v} \neg \phi(\bar{v})$, then $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow c \neq c)$.

Thus, we may assume that both $T \cup \{\phi(\bar{v})\}$ and $T \cup \{\neg \phi(\bar{v})\}$ are satisfiable.

Let $\Gamma(\bar{v}) = \{\psi(\bar{v}) : \psi \text{ is quantifier-free and } T \models \forall \bar{v} (\phi(\bar{v}) \rightarrow \psi(\bar{v}))\}$. Let d_1, \dots, d_m be new constant symbols. We will show that $T \cup \Gamma(\bar{d}) \models \phi(\bar{d})$. Then, by compactness, there are $\psi_1, \dots, \psi_n \in \Gamma$ such that

$$T \models \forall \bar{v} \left(\bigwedge_{i=1}^n \psi_i(\bar{v}) \rightarrow \phi(\bar{v}) \right).$$

Thus

$$T \models \forall \bar{v} \left(\bigwedge_{i=1}^n \psi_i(\bar{v}) \leftrightarrow \phi(\bar{v}) \right)$$

and $\bigwedge_{i=1}^n \psi_i(\bar{v})$ is quantifier-free. We need only prove the following claim.

Claim $T \cup \Gamma(\bar{d}) \models \phi(\bar{d})$.

Suppose not. Let $\mathcal{M} \models T \cup \Gamma(\bar{d}) \cup \{\neg \phi(\bar{d})\}$. Let \mathcal{A} be the substructure of \mathcal{M} generated by \bar{d} .

Let $\Sigma = T \cup \text{Diag}(\mathcal{A}) \cup \phi(\bar{d})$. If Σ is unsatisfiable, then there are quantifier-free formulas $\psi_1(\bar{d}), \dots, \psi_n(\bar{d}) \in \text{Diag}(\mathcal{A})$ such that

$$T \models \forall \bar{v} \left(\bigwedge_{i=1}^n \psi_i(\bar{v}) \rightarrow \neg \phi(\bar{v}) \right).$$

But then

$$T \models \forall \bar{v} \left(\phi(\bar{v}) \rightarrow \bigvee_{i=1}^n \neg \psi_i(\bar{v}) \right),$$

so $\bigvee_{i=1}^n \neg \psi_i(\bar{v}) \in \Gamma$ and $\mathcal{A} \models \bigvee_{i=1}^n \neg \psi_i(\bar{d})$, a contradiction. Thus, Σ is satisfiable.

Let $\mathcal{N} \models \Sigma$. Then $\mathcal{N} \models \phi(\bar{d})$. Because $\Sigma \supseteq \text{Diag}(\mathcal{A})$, $\mathcal{A} \subseteq \mathcal{N}$, by Lemma 5.3 i). But $\mathcal{M} \models \neg \phi(\bar{d})$; thus, by ii), $\mathcal{N} \models \neg \phi(\bar{d})$, a contradiction.

The proof above can easily be adapted to the case where \mathcal{L} contains no constant symbols. In this case, there are no quantifier-free sentences, but for each sentence we can find a quantifier-free formula $\psi(v_1)$ such that $T \models \phi \leftrightarrow \psi(v_1)$.

The next lemma shows that we can prove quantifier elimination by getting rid of one existential quantifier at a time.

Lemma 5.5 *Let T be an \mathcal{L} -theory. Suppose that for every quantifier-free \mathcal{L} -formula $\theta(\bar{v}, w)$ there is a quantifier-free formula $\psi(\bar{v})$ such that $T \models \exists w \theta(\bar{v}, w) \leftrightarrow \psi(\bar{v})$. Then, T has quantifier elimination.*

Proof Let $\phi(\bar{v})$ be an \mathcal{L} -formula. We wish to show that $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \psi(\bar{v}))$ for some quantifier-free formula $\psi(\bar{v})$. We prove this by induction on the complexity of $\phi(\bar{v})$.

If ϕ is quantifier-free, there is nothing to prove. Suppose that for $i = 0, 1$, $T \models \forall \bar{v} (\theta_i(\bar{v}) \leftrightarrow \psi_i(\bar{v}))$, where ψ_i is quantifier free.

If $\phi(\bar{v}) = \neg\theta_0(\bar{v})$, then $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \neg\psi_0(\bar{v}))$.

If $\phi(\bar{v}) = \theta_0(\bar{v}) \wedge \theta_1(\bar{v})$, then $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow (\psi_0(\bar{v}) \wedge \psi_1(\bar{v})))$.

In either case, ϕ is equivalent to a quantifier-free formula.

Suppose that $T \models \forall \bar{v} (\theta(\bar{v}, w) \leftrightarrow \psi_0(\bar{v}, w))$, where ψ_0 is quantifier-free and $\phi(\bar{v}) = \exists w \theta(\bar{v}, w)$. Then $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \exists w \psi_0(\bar{v}, w))$. By our assumptions, there is a quantifier-free $\psi(\bar{v})$ such that $T \models \forall \bar{v} (\exists w \psi_0(\bar{v}, w) \leftrightarrow \psi(\bar{v}))$. But then $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \psi(\bar{v}))$.

Combining Theorem 5.4 and Lemma 5.5 gives us the following simple, yet useful, test for quantifier elimination.

Corollary 5.6 *Let T be an \mathcal{L} -theory. Suppose that for all quantifier-free formulas $\phi(\bar{v}, w)$, if $\mathcal{M}, \mathcal{N} \models T$, \mathcal{A} is a common substructure of \mathcal{M} and \mathcal{N} , $\bar{a} \in \mathcal{A}$, and there is $b \in M$ such that $\mathcal{M} \models \phi(\bar{a}, b)$, then there is $c \in N$ such that $\mathcal{N} \models \phi(\bar{a}, c)$. Then, T has quantifier elimination.*

Theories with Quantifier Elimination

We conclude with several observations about theories with quantifier elimination.

Definition 5.7 An \mathcal{L} -theory T is *model-complete* $\mathcal{M} \prec \mathcal{N}$ whenever $\mathcal{M} \subseteq \mathcal{N}$ and $\mathcal{M}, \mathcal{N} \models T$.

Stated in terms of embeddings: T is model-complete if and only if all embeddings are elementary.

Proposition 5.8 *If T has quantifier elimination, then T is model-complete.*

Proof Suppose that $\mathcal{M} \subseteq \mathcal{N}$ are models of T . We must show that \mathcal{M} is an elementary submodel. Let $\phi(\bar{v})$ be an \mathcal{L} -formula, and let $\bar{a} \in M$. There is a quantifier-free formula $\psi(\bar{v})$ such that $\mathcal{M} \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \psi(\bar{v}))$. Because quantifier-free formulas are preserved under substructures and extensions, $\mathcal{M} \models \psi(\bar{a})$ if and only if $\mathcal{N} \models \psi(\bar{a})$. Thus

$$\mathcal{M} \models \phi(\bar{a}) \Leftrightarrow \mathcal{M} \models \psi(\bar{a}) \Leftrightarrow \mathcal{N} \models \psi(\bar{a}) \Leftrightarrow \mathcal{N} \models \phi(\bar{a}).$$

There are model-complete theories that do not have quantifier elimination, but model completeness implies that we can eliminate all but the last existential quantifiers.

Proposition 5.9 *If T is model complete, then for any formula $\phi(\bar{v})$, there is a quantifier free formula $\psi(\bar{v}, \bar{w})$ such that*

$$T \models \forall \bar{v} [\phi(\bar{v}) \leftrightarrow \exists \bar{w} \psi(\bar{v}, \bar{w})].$$

Let us just point out the following test for completeness of model-complete theories.

Proposition 5.10 *Let T be a model-complete theory. Suppose that there is $\mathcal{M}_0 \models T$ such that \mathcal{M}_0 embeds into every model of T . Then, T is complete.*

Proof If $\mathcal{M} \models T$, the embedding of \mathcal{M}_0 into \mathcal{M} is elementary. In particular $\mathcal{M}_0 \equiv \mathcal{M}$. Thus, any two models of T are elementarily equivalent.

We will use Proposition 5.10 below in cases where Vaught's test does not apply.

We have provided a number of proofs of quantifier elimination without explicitly explaining how to take an arbitrary formula and produce a quantifier free one. In all of these cases, one can give explicit effective procedures. After the fact, the following lemma tells us that there is an algorithm to eliminate quantifiers.

Proposition 5.11 *Suppose that T is a decidable theory with quantifier elimination. Then, there is an algorithm which when given a formula ϕ as input will output a quantifier-free formula ψ such that $T \models \phi \leftrightarrow \psi$.*

Proof Given input $\phi(\bar{v})$ we search for a quantifier-free formula $\psi(\bar{v})$ such that $T \models \forall \bar{v} (\phi(\bar{v}) \leftrightarrow \psi(\bar{v}))$. Because T is decidable this is an effective search. Because T has quantifier elimination, we will eventually find ψ .

6 Algebraically Closed Fields

We now return to the theory of algebraically closed fields. In Proposition 4.4, we proved that the theory of algebraically closed fields of a fixed characteristic is complete. We begin this section by showing that algebraically closed fields have quantifier elimination. For convenience we will formulate ACF in the language $\mathcal{L} = \{+, -, \cdot, 0, 1\}$. We add $-$ to the language, so that substructures are integral domains. Without $-$ we would have structures that are a bit more cumbersome to deal with.

Theorem 6.1 *ACF has quantifier elimination.*

Proof

Suppose K and L are algebraically closed fields and A is an integral domain with $A \subseteq K \cap L$. By Corollary 5.6, we need to show that if $\phi(v, \bar{w})$ is a quantifier free formula, $\bar{a} \in A$, $b \in K$ and $K \models \phi(b, \bar{a})$, then there is $c \in L$ such that $L \models \phi(c, \bar{a})$.

Let F be the algebraic closure of the fraction field of A . We, may without loss of generality, assume that $F \subseteq K \cap L$. It will be enough to show that, $\bar{a} \in F$, and $K \models \phi(b, \bar{a})$ for some $b \in K$, then there is $c \in F$ such that $F \models \phi(c, \bar{a})$, for then, by Proposition 1.8, $L \models \phi(c, \bar{a})$.

We first note that ϕ can be put in disjunctive normal form, namely there are atomic or negated atomic formulas $\theta_{i,j}(\bar{v}, w)$ such that:

$$\phi(\bar{v}, w) \leftrightarrow \bigvee_{i=1}^n \bigwedge_{j=1}^m \theta_{i,j}(\bar{v}, w).$$

Because $K \models \phi(\bar{a}, b)$, $K \models \bigwedge_{j=1}^m \theta_{i,j}(\bar{a}, b)$ for some i . Thus, without loss of generality, we may assume that ϕ is a conjunction of atomic and negated atomic formulas. In our language atomic formulas $\theta(v_1, \dots, v_n)$ are of the form $p(\bar{v}) = 0$, where $p \in \mathbb{Z}[X_1, \dots, X_n]$. If $p(X, \bar{Y}) \in \mathbb{Z}[X, \bar{Y}]$, we can view $p(X, \bar{a})$ as a polynomial in $F[X]$. Thus, there are polynomials $p_1, \dots, p_n, q_1, \dots, q_m \in F[X]$ such that $\phi(v, \bar{a})$ is equivalent to

$$\bigwedge_{i=1}^n p_i(v) = 0 \wedge \bigwedge_{i=1}^m q_i(v) \neq 0.$$

If any of the polynomials p_i are nonzero, then b is algebraic over F . In this case, because F is algebraically closed, $b \in F$. Thus, we may assume that $\phi(v, \bar{a})$ is equivalent to

$$\bigwedge_{i=1}^m q_i(v) \neq 0.$$

But $q_i(X) = 0$ has only finitely many solutions for each $i \leq m$. Thus, there are only finitely many elements of F that do not satisfy F . Because algebraically closed fields are infinite, there is a $c \in F$ such that $F \models \phi(c, \bar{a})$.

Corollary 6.2 *ACF is model-complete and ACF_p is complete where $p = 0$ or p is prime.*

Proof Model-completeness is an immediate consequence of quantifier elimination.

The completeness of ACF_p was proved in Proposition 4.4, but it also follows from quantifier elimination. Suppose that $K, L \models \text{ACF}_p$. Let ϕ be any sentence in the language of rings. By quantifier elimination, there is a quantifier-free sentence ψ such that

$$\text{ACF} \models \phi \leftrightarrow \psi.$$

Because quantifier-free sentences are preserved under extension and substructure,

$$K \models \psi \leftrightarrow \mathbb{F}_p \models \psi \leftrightarrow L \models \psi,$$

where \mathbb{F}_p is the p -element field if $p > 0$ and the rationals if $p = 0$. Thus,

$$K \models \phi \leftrightarrow K \models \psi \leftrightarrow L \models \psi \leftrightarrow L \models \phi.$$

Thus $K \equiv L$ and ACF_p is complete.

Definable Sets and Constructible Sets

Quantifier elimination has a geometric interpretation. We begin by looking at the sets defined by quantifier free formulas.

Lemma 6.3 *Let K be a field. The subsets of K^n defined by atomic formulas are exactly those of the form $V(p) = \{x \text{ for some } p \in K[\bar{X}]\}$. A subset of K^n is quantifier-free definable if and only if it is a Boolean combination of Zariski closed subsets.*

Proof If $\phi(\bar{x}, \bar{y})$ is an atomic \mathcal{L}_r -formula, then there is $q(\bar{X}, \bar{Y}) \in \mathbb{Z}[\bar{X}, \bar{Y}]$ such that $\phi(\bar{x}, \bar{y})$ is equivalent to $q(\bar{x}, \bar{y}) = 0$. If $X = \{\bar{x} : \phi(\bar{x}, \bar{a})\}$, then $X = V(q(\bar{X}, \bar{a}))$ and $q(\bar{X}, \bar{a}) \in K[\bar{X}]$. On the other hand, if $p \in K[\bar{X}]$, there is $q \in \mathbb{Z}[\bar{X}, \bar{Y}]$ and $\bar{a} \in K^m$ such that $p(\bar{X}) = q(\bar{X}, \bar{a})$. Then, $V(p)$ is defined by the quantifier-free formula $q(\bar{X}, \bar{a}) = 0$.

If $X \subseteq K^n$ is a finite Boolean combination of Zariski closed sets we call X *constructible*. If K is algebraically closed, the constructible sets have much stronger closure properties.

Corollary 6.4 *Let K be an algebraically closed field.*

- i) $X \subseteq K^n$ is constructible if and only if it is definable.*
- ii) (Chevalley's Theorem) The image of a constructible set under a polynomial map is constructible.*

Proof i) By Lemma 6.3, the constructible sets are exactly the quantifier-free definable sets, but by quantifier elimination every definable set is quantifier-free definable.

ii) Let $X \subseteq K^n$ be constructible and $p : K^n \rightarrow K^m$ be a polynomial map. Then, the image of $X = \{y \in K^m : \exists x \in K^n p(x) = y\}$. This set is definable and hence constructible.

Quantifier elimination has very strong consequences for definable subsets of K .

Corollary 6.5 *If K is an algebraically closed field and $X \subseteq K$ is definable, then either X or $K \setminus X$ is finite.*

Proof By quantifier elimination X is a finite Boolean combination of sets of the form $V(p)$, where $p \in K[X]$. But $V(p)$ is either finite or (if $p = 0$) all of K .

We say that a theory T is *strongly minimal* if for any $\mathcal{M} \models T$ and any definable $X \subseteq M$ either X or $M \setminus X$ is finite. This is a very powerful assumption. For example, it can be shown that any strongly minimal theory in a countable language is κ -categorical for every uncountable κ .

The model-completeness of algebraically closed fields can be used to give a proof of the Nullstellensatz.

Theorem 6.6 (Hilbert's Nullstellensatz) *Let K be an algebraically closed field. Suppose that I and J are radical ideals in $K[X_1, \dots, X_n]$ and $I \subset J$. Then $V(J) \subset V(I)$. Thus $X \mapsto I(X)$ is a bijective correspondence between Zariski closed sets and radical ideals.*

Proof Let $p \in J \setminus I$. By Primary Decomposition, there is a prime ideal $P \supseteq I$ such that $p \notin P$. We will show that there is $x \in V(P) \subseteq V(I)$ such that $p(x) \neq 0$. Thus $V(I) \neq V(J)$. Because P is prime, $K[\bar{X}]/P$ is a domain and we can take F , the algebraic closure of its fraction field.

Let $q_1, \dots, q_m \in K[X_1, \dots, X_n]$ generate I . Let a_i be the element X_i/P in F . Because each $q_i \in P$ and $p \notin P$,

$$F \models \bigwedge_{i=1}^m q_i(\bar{a}) = 0 \wedge p(\bar{a}) \neq 0.$$

Thus

$$F \models \exists \bar{w} \bigwedge_{i=1}^m q_i(\bar{w}) = 0 \wedge p(\bar{w}) \neq 0$$

and by model-completeness

$$K \models \exists \bar{w} \bigwedge_{i=1}^m q_i(\bar{w}) = 0 \wedge p(\bar{w}) \neq 0.$$

Thus there is $\bar{b} \in K^n$ such that $q_1(\bar{b}) = \dots = q_m(\bar{b}) = 0$ and $p(\bar{b}) \neq 0$. But then $\bar{b} \in V(P) \setminus V(J)$.

The next corollary is a simple consequence of model completeness.

Corollary 6.7 *Suppose $K \subseteq L$ are algebraically closed fields, V and W are varieties defined over K and $f : V \rightarrow W$ is a polynomial isomorphism defined over L . Then there is an isomorphism defined over K .*

Proof Suppose $f : V \rightarrow W$ is a polynomial isomorphism defined over L and f and f^{-1} both have degree at most d . As in the proof of Ax's Theorem we can write down an \mathcal{L} -formula Ψ with parameters from K saying that for some choice of coefficients there is a polynomial bijection from V between V and W where the polynomials have degree at most d . Since $L \models \Psi$, by model completeness, $K \models \Psi$. Thus we can choose an isomorphism defined over K .

Quantifier elimination gives us a powerful tool for analyzing definability in algebraically closed fields. For example, we will give the following characterization of definable functions.

Definition 6.8 Let $X \subseteq K^n$. We say that $f : X \rightarrow K$ is *quasirational* if either
i) K has characteristic zero and for some rational function $q(\bar{X}) \in K(X_1, \dots, X_n)$, $f(\bar{x}) = q(\bar{x})$ on X , or
ii) K has characteristic $p > 0$ and for some rational function $q(\bar{X}) \in K(\bar{X})$, $f(\bar{x}) = q(\bar{x})^{\frac{1}{p^n}}$.

Rational functions are easily seen to be definable. In algebraically closed fields of characteristic p , the formula $x = y^p$ defines the function $x \mapsto x^{\frac{1}{p}}$, because every element has a unique p^{th} -root. Thus, every quasirational function is definable.

Proposition 6.9 *If $X \subseteq K^n$ is constructible and $f : X \rightarrow K$ is definable, then there are constructible sets X_1, \dots, X_m and quasirational functions ρ_1, \dots, ρ_m such that $\bigcup X_i = X$ and $f|_{X_i} = \rho_i|_{X_i}$.*

Proof Let $\Gamma(v_1, \dots, v_n) = \{f(\bar{v}) \neq \rho(\bar{v}) : \rho \text{ a quasirational function}\} \cup \{\bar{v} \in X\} \cup \text{ACF} \cup \text{Diag}(K)$.

Claim Γ is not satisfiable.

Suppose that Γ is consistent. Let $L \models \text{ACF} + \text{Diag}(K)$ with $b_1, \dots, b_n \in L$ such that for all $\gamma(\bar{v}) \in \Gamma$, $L \models \gamma(\bar{b})$.

Let K_0 be the subfield of L generated by K and \bar{b} . Then, K_0 is the closure of $B = \{b_1, \dots, b_n\}$ under the rational functions of K . Let K_1 be the closure of B under all quasirational functions. If K has characteristic 0, then $K_0 = K_1$. If K has characteristic $p > 0$, $K_1 = \bigcup K_0^{\frac{1}{p^n}}$, the perfect closure of K_0 .

By model-completeness, $K \prec L$, thus f^L , the interpretation of f in L , is a function from X^L to L , extending f . Because $L \models \Gamma(\bar{b})$, $f(\bar{b})$ is not in K_1 . Because K_1 is perfect there is an automorphism α of L fixing K_1 pointwise such that $\alpha(f^L(\bar{b})) \neq f^L(\bar{b})$. But f^L is definable with parameters from K ; thus, any automorphism of L which fixes K and fixes \bar{a} must fix $f(\bar{a})$, a contradiction. Thus Γ is unsatisfiable.

Thus, by compactness, there are quasirational functions ρ_1, \dots, ρ_m such that

$$K \models \forall x \in X \bigwedge f(\bar{x}) = \rho_i(\bar{x}).$$

Let $X_i = \{\bar{x} \in X : f(\bar{x}) = \rho_i(\bar{x})\}$. Each X_i is definable.

We end by stating two more far reaching definability results for algebraically closed. They are a bit more involved—and ideally best understood using the model theoretic tool of ω -stability that we will not discuss in these lectures.

Let K be algebraically closed.

Theorem 6.10 (Elimination of Imaginaries) *Suppose $X \subseteq K^n$ is definable and E is a definable equivalence relation on X . There is a definable $f : X \rightarrow K^m$ for some m such that xEy if and only if $f(x) = f(y)$.*

This is related to the existence of fields of definitions. It is a useful tool for viewing projective, quasiprojective or abstract varieties (at least in the style of Weil) as constructible objects.

Theorem 6.11 *Let $G \subseteq K^n$ be a definable group. Then G is definably isomorphic to an algebraic group.*

Combining these we could conclude that if G is an algebraic group and H is a normal algebraic subgroup, then G/H is an algebraic group.

7 Real Closed Fields and o-minimality

In this section, we will concentrate on the field of real numbers. Unlike algebraically closed fields, the theory of the real numbers does not have quantifier elimination in $\mathcal{L}_r = \{+, \cdot, 0, 1\}$, the language of rings. The proof of Corollary 6.5 shows that any field with quantifier elimination is strongly minimal, whereas in \mathbb{R} , if $\phi(x)$ is the formula $\exists z z^2 = x$, then ϕ defines an infinite coinfinite definable set. In fact, algebraically closed fields are the only infinite fields with quantifier elimination.

In fact, the ordering is the only obstruction to quantifier elimination. We will eventually analyze the real numbers in the language $\mathcal{L}_{or} = \{+, -, \cdot, <, 0, 1\}$ and show that we have quantifier elimination in this language. Because the ordering $x < y$ is definable in the real field by the formula

$$\exists z (z \neq 0 \wedge x + z^2 = y),$$

any subset of \mathbb{R}^n definable using an \mathcal{L}_{or} -formula is already definable using an \mathcal{L}_r -formula). We will see that quantifier elimination in \mathcal{L}_{or} leads us to a good geometric understanding of the definable sets.

We begin by reviewing some of the necessary algebraic background on ordered fields. All of the algebraic results stated in this chapter are due to Artin and Schreier. These results are all proved in the appendix

Definition 7.1 We say that a field F is *orderable* if there is a linear order $<$ of F making $(F, <)$ an ordered field.

Although there are unique orderings of the fields \mathbb{R} and \mathbb{Q} , orderable fields may have many possible orderings. The field of rational functions $\mathbb{Q}(X)$ has 2^{\aleph_0} distinct orderings. To see this, let x be any real number transcendental over \mathbb{Q} . The evaluation map $f(X) \mapsto f(x)$ is a field isomorphism between $\mathbb{Q}(X)$ and $\mathbb{Q}(x)$, the subfield of \mathbb{R} generated by x . We can lift the ordering of the reals to an ordering $\mathbb{Q}(X)$ by $f(X) < g(X)$ if and only if $f(x) < g(x)$. Because $X < q$ if and only if $x < q$, choosing a different transcendental real would yield a different ordering. These are not the only orderings. We can also order $\mathbb{Q}(X)$ by making X infinite or infinitesimally close to a rational.

There is a purely algebraic characterization of the orderable fields.

Definition 7.2 We say that F is *formally real* if -1 is not a sum of squares.

In any ordered field all squares are nonnegative. Thus, every orderable field is formally real. The following result shows that the converse is also true.

Theorem 7.3 *If F is a formally real field, then F is orderable. Indeed, if $a \in F$ and $-a$ is not a sum of squares of elements of F , then there is an ordering of F where a is positive.*

Because the field of complex numbers is the only proper algebraic extension of the real field, the real numbers have no proper formally real algebraic extensions. Fields with this property will play a key role.

Definition 7.4 A field F is *real closed* if it is formally real with no proper formally real algebraic extensions.

Although it is not obvious at first that real closed fields form an elementary class, the next theorem allows us to axiomatize the real closed fields.

Theorem 7.5 *Let F be a formally real field. The following are equivalent.*

- i) F is real closed.
- ii) $F(i)$ is algebraically closed (where $i^2 = -1$).
- iii) For any $a \in F$, either a or $-a$ is a square and every polynomial of odd degree has a root.

Corollary 7.6 *The class of real closed fields is an elementary class of \mathcal{L}_r -structures.*

Proof We can axiomatize real closed fields by:

- i) axioms for fields
- ii) for each $n \geq 1$, the axiom

$$\forall x_1 \dots \forall x_n \ x_1^2 + \dots + x_n^2 + 1 \neq 0$$

- iii) $\forall x \exists y \ (y^2 = x \vee y^2 + x = 0)$
- iv) for each $n \geq 0$, the axiom

$$\forall x_0 \dots \forall x_{2n} \exists y \ y^{2n+1} + \sum_{i=0}^{2n} x_i y^i = 0.$$

Although we can axiomatize real closed fields in the language of rings, we already noticed that we do not have quantifier elimination in this language. Instead, we will study real closed fields in \mathcal{L}_{or} , the language of ordered rings. If F is a real closed field and $0 \neq a \in F$, then exactly one of a and $-a$ is a square. This allows us to order F by

$$x < y \text{ if and only if } y - x \text{ is a nonzero square.}$$

It is easy to check that this is an ordering and it is the only possible ordering of F .

Definition 7.7 We let RCF be the \mathcal{L}_{or} -theory axiomatized by the axioms above for real closed fields and the axioms for ordered fields.

The models of RCF are exactly real closed fields with their canonical ordering. Because the ordering is defined by the \mathcal{L}_r -formula

$$\exists z \ (z \neq 0 \wedge x + z^2 = y),$$

the next result tells us that using the ordering does not change the definable sets.

Proposition 7.8 *If F is a real closed field and $X \subseteq F^n$ is definable by an \mathcal{L}_{or} -formula, then X is definable by an \mathcal{L}_r -formula.*

Proof Replace all instances of $t_i < t_j$ by $\exists v (v \neq 0 \wedge v^2 + t_i = t_j)$, where t_i and t_j are terms occurring in the definition of X .

The next result suggests another possible axiomatization of RCF.

Theorem 7.9 *An ordered field F is real closed if and only if whenever $p(X) \in F[X]$, $a, b \in F$, $a < b$, and $p(a)p(b) < 0$, there is $c \in F$ such that $a < c < b$ and $p(c) = 0$.*

Definition 7.10 *If F is a formally real field, a real closure of F is a real closed algebraic extension of F .*

By Zorn's Lemma, every formally real field F has a maximal formally real algebraic extension. This maximal extension is a real closure of F .

The real closure of a formally real field may not be unique. Let $F = \mathbb{Q}(X)$, $F_0 = F(\sqrt{X})$, and $F_1 = F(\sqrt{-X})$. By Theorem 7.3, F_0 and F_1 are formally real. Let R_i be a real closure of F_i . There is no isomorphism between R_0 and R_1 fixing F because X is a square in R_0 but not in R_1 . Thus, some work needs to be done to show that any ordered field $(F, <)$ has a real closure where the canonical order extends the ordering of F .

Lemma 7.11 *If $(F, <)$ is an ordered field, $0 < x \in F$, and x is not a square in F , then we can extend the ordering of F to $F(\sqrt{x})$.*

Proof We can extend the ordering to $F(\sqrt{x})$ by $0 < a + b\sqrt{x}$ if and only if

- i) $b = 0$ and $a > 0$, or
- ii) $b > 0$ and $(a > 0 \text{ or } x > \frac{a^2}{b^2})$, or
- iii) $b < 0$ and $(a < 0 \text{ and } x < \frac{a^2}{b^2})$.

Corollary 7.12 *If $(F, <)$ is an ordered field, there is a real closure R of F such that the canonical ordering of R extends the ordering on F .*

Proof

By successive applications of Lemma 7.11, we can find an ordered field $(L, <)$ extending $(F, <)$ such that every positive element of F has a square root in L . We now apply Zorn's Lemma to find a maximal formally real algebraic extension R of L . Because every positive element of F is a square in R , the canonical ordering of R extends the ordering of F .

Although a formally real field may have nonisomorphic real closures, if $(F, <)$ is an ordered field there will be a unique real closure compatible with the ordering of F .

Theorem 7.13 *If $(F, <)$ is an ordered field, and R_1 and R_2 are real closures of F where the canonical ordering extends the ordering of F , then there is a unique field isomorphism $\phi : R_1 \rightarrow R_2$ that is the identity on F .*

Note that because the ordering of a real closed field is definable in \mathcal{L}_r , ϕ also preserves the ordering. We often say that any ordered field $(F, <)$ has a unique real closure. By this we mean that there is a unique real closure that extends the given ordering.

Quantifier Elimination for Real Closed Fields

We are now ready to prove quantifier elimination.

Theorem 7.14 *The theory RCF admits elimination of quantifiers in \mathcal{L}_{or} .*

Proof We use the quantifier elimination tests of §5. Suppose K and L are real closed ordered fields and A is a common substructure. Then A is an ordered integral domain. We extend the ordering on A to its fraction field to obtain an ordered subfield $F_0 \subseteq K \cap L$. Let F be the real closure of F_0 . By uniqueness of real closures, F is isomorphic, as an ordered field, to the algebraic closure of F_0 inside K and L . Without loss of generality we may assume $F \subseteq K \cap L$.

It suffices then to show that if $\phi(v, \bar{w})$ is a quantifier-free formula, $\bar{a} \in F$, $b \in K$ and $K \models \phi(b, \bar{a})$, then there is $b' \in F$ such that $F \models \phi(b', \bar{a})$.

Note that

$$p(X) \neq 0 \leftrightarrow (p(\bar{X}) > 0 \vee -p(\bar{X}) > 0)$$

and

$$p(\bar{X}) \not> 0 \leftrightarrow (p(\bar{X}) = 0 \vee -p(\bar{X}) > 0).$$

With this in mind, we may assume that ϕ is a disjunction of conjunctions of formulas of the form $p(v, \bar{w}) = 0$ or $p(v, \bar{w}) > 0$. As in Theorem 6.1, we may assume that there are polynomials p_1, \dots, p_n and $q_1, \dots, q_m \in F[X]$ such that

$$\phi(v, \bar{a}) \leftrightarrow \bigwedge_{i=1}^n p_i(v) = 0 \wedge \bigwedge_{i=1}^m q_i(v) > 0.$$

If any of the polynomials $p_i(X)$ is nonzero, then b is algebraic over F . Because F has no proper formally real algebraic extensions, in this case $b \in F$. Thus, we may assume that

$$\phi(v, \bar{a}) \leftrightarrow \bigwedge_{i=1}^m q_i(v) > 0.$$

The polynomial $q_i(X)$ can only change signs at zeros of q_i and all zeros of q_i are in F . Thus, we can find $c_i, d_i \in F$ such that $c_i < b < d_i$ and $q_i(x) > 0$ for all $x \in (c_i, d_i)$. Let $c = \max(c_1, \dots, c_m)$ and $d = \min(d_1, \dots, d_m)$. Then, $c < d$ and $\bigwedge_{i=1}^m q_i(x) > 0$ whenever $c < x < d$. Thus, we can find $b' \in F$ such that $F \models \phi(b', \bar{a})$.

Corollary 7.15 *RCF is complete, model complete and decidable. Thus RCF is the theory of $(\mathbb{R}, +, \cdot, <)$ and RCF is decidable.*

Proof By quantifier elimination, RCF is model complete.

Every real closed field has characteristic zero; thus, the rational numbers are embedded in every real closed field. Therefore, \mathbb{R}_{alg} , the field of real algebraic numbers (i.e., the real closure of the rational numbers) is a subfield of any real closed field. Thus, for any real closed field R , $\mathbb{R}_{\text{alg}} \prec R$, so $R \equiv \mathbb{R}_{\text{alg}}$.

In particular, $R \equiv \mathbb{R}_{\text{alg}} \equiv \mathbb{R}$.

Because RCF is complete and recursively axiomatized, it is decidable.

Semialgebraic Sets

Quantifier elimination for real closed fields has a geometric interpretation.

Definition 7.16 Let F be an ordered field. We say that $X \subseteq F^n$ is *semialgebraic* if it is a Boolean combination of sets of the form $\{\bar{x} : p(\bar{x}) > 0\}$, where $p(\bar{X}) \in F[X_1, \dots, X_n]$.

By quantifier elimination, the semialgebraic sets are exactly the definable sets. The next corollary is a geometric restatement of quantifier elimination. It is analogous to Chevalley's Theorem (6.4) for algebraically closed fields.

Corollary 7.17 (Tarski–Seidenberg Theorem) *The semialgebraic sets are closed under projection.*

The next corollary is a typical application of quantifier elimination.

Corollary 7.18 *If $F \models \text{RCF}$ and $A \subseteq F^n$ is semialgebraic, then the closure (in the Euclidean topology) of A is semialgebraic.*

Proof We repeat the main idea of Lemma 1.26. Let d be the definable function

$$d(x_1, \dots, x_n, y_1, \dots, y_n) = z \text{ if and only if } z \geq 0 \wedge z^2 = \sum_{i=1}^n (x_i - y_i)^2.$$

The closure of A is

$$\{\bar{x} : \forall \epsilon > 0 \exists \bar{y} \in A \ d(\bar{x}, \bar{y}) < \epsilon\}.$$

Because this set is definable, it is semialgebraic.

We say that a function is semialgebraic if its graph is semialgebraic. The next result shows how we can use the completeness of RCF to transfer results from \mathbb{R} to other real closed fields.

Corollary 7.19 *Let F be a real closed field. If $X \subseteq F^n$ is semialgebraic, closed and bounded, and f is a continuous semialgebraic function, then $f(X)$ is closed and bounded.*

Proof If $F = \mathbb{R}$, then X is closed and bounded if and only if X is compact. Because the continuous image of a compact set is compact, the continuous image of a closed and bounded set is closed and bounded.

In general, there are $\bar{a}, \bar{b} \in F$ and formulas ϕ and ψ such that $\phi(\bar{x}, \bar{a})$ defines X and $\psi(\bar{x}, y, \bar{b})$ defines $f(\bar{x}) = y$. There is a sentence Φ asserting:

$\forall \bar{u}, \bar{w}$ [if $\psi(\bar{x}, y, \bar{w})$ defines a continuous function with domain $\phi(\bar{x}, \bar{u})$ and $\phi(\bar{x}, \bar{u})$ is a closed and bounded set, then the range of the function is closed and bounded].

By the remarks above, $\mathbb{R} \models \Phi$. Therefore, by the completeness of RCF, $F \models \Phi$ and the range of f is closed and bounded.

Model-completeness has several important applications. A typical application is Abraham Robinson's simple proof of Artin's positive solution to Hilbert's 17th problem.

Definition 7.20 Let F be a real closed field and $f(\bar{X}) \in F(X_1, \dots, X_n)$ be a rational function. We say that f is *positive semidefinite* if $f(\bar{a}) \geq 0$ for all $\bar{a} \in F^n$.

Theorem 7.21 (Hilbert's 17th Problem) *If f is a positive semidefinite rational function over a real closed field F , then f is a sum of squares of rational functions.*

Proof Suppose that $f(X_1, \dots, X_n)$ is a positive semidefinite rational function over F that is not a sum of squares. By Theorem 7.3, there is an ordering of $F(\bar{X})$ so that f is negative. Let R be the real closure of $F(\bar{X})$ extending this order. Then

$$R \models \exists \bar{v} f(\bar{v}) < 0$$

because $f(\bar{X}) < 0$ in R . By model-completeness

$$F \models \exists \bar{v} f(\bar{v}) < 0,$$

contradicting the fact that f is positive semidefinite.

We will show that quantifier elimination gives us a powerful tool for understanding the definable subsets of a real closed field.

Definition 7.22 Let $\mathcal{L} \supseteq \{<\}$. Let T be an \mathcal{L} -theory extending the theory of linear orders. We say that T is *o-minimal* if for all $\mathcal{M} \models T$ if $X \subseteq M$ is definable, then X is a finite union of points and intervals with endpoints in $M \cup \{\pm\infty\}$.

We can think of o-minimality as an analog of strong minimality for ordered structures. Strong minimality says that the only definable subsets in dimension one can be defined using only equality—i.e., the ones that can be defined in any structure. O-minimality says the only sets that can be defined in one dimension are the ones definable in any ordered structure.

Corollary 7.23 *RCF is an o-minimal theory.*

Proof Let $R \models \text{RCF}$. We need to show that every definable subset of R is a finite union of points and intervals with endpoints in $R \cup \{\pm\infty\}$. By quantifier elimination, every definable subset of R is a finite Boolean combination of sets of the form

$$\{x \in R : p(x) = 0\}$$

and

$$\{x \in R : q(x) > 0\}.$$

Solution sets to nontrivial equations are finite, whereas sets of the second form are finite unions of intervals. Thus, any definable set is a finite union of points and intervals.

Next we will show that definable functions in one variable are piecewise continuous. The first step is to prove a lemma about \mathbb{R} that we will transfer to all real closed fields.

Lemma 7.24 *If $f : \mathbb{R} \rightarrow \mathbb{R}$ is semialgebraic, then for any open interval $U \subseteq \mathbb{R}$ there is a point $x \in U$ such that f is continuous at x .*

Proof

case 1: There is an open set $V \subseteq U$ such that f has finite range on V .

Pick an element b in the range of f such that $\{x \in V : f(x) = b\}$ is infinite. By o-minimality, there is an open set $V_0 \subseteq V$ such that f is constantly b on V .

case 2: Otherwise.

We build a chain $U = V_0 \supset V_1 \supset V_2 \dots$ of open subsets of U such that the closure $\overline{V_{n+1}}$ of V_{n+1} is contained in V_n . Given V_n , let X be the range of f on V_n . Because X is infinite, by o-minimality, X contains an interval (a, b) of length at most $\frac{1}{n}$. The set $Y = \{x \in V_n : f(x) \in (a, b)\}$ contains a suitable open interval V_{n+1} . Because \mathbb{R} is locally compact,

$$\bigcap_{i=1}^{\infty} V_i = \bigcap_{i=1}^{\infty} \overline{V_i} \neq \emptyset.$$

If $x \in \bigcap_{i=1}^{\infty} V_i$, then f is continuous at x .

The proof above makes essential use of the completeness of the ordering of the reals. However, because the statement is first order, it is true for all real closed fields, by the completeness of RCF.

Corollary 7.25 *Let F be a real closed field and $f : F \rightarrow F$ is a semialgebraic function. Then, we can partition F into $I_1 \cup \dots \cup I_m \cup X$, where X is finite and the I_j are pairwise disjoint open intervals with endpoints in $F \cup \{\pm\infty\}$ such that f is continuous on each I_j .*

Proof Let

$$D = \{x : F \models \exists \epsilon > 0 \forall \delta > 0 \exists y |x - y| < \delta \wedge |f(x) - f(y)| > \epsilon\}$$

be the set of points where f is discontinuous. Because D is definable, by o-minimality D is either finite or has a nonempty interior. By Corollary 7.23, D must be finite. Thus, $F \setminus D$ is a finite union of intervals on which F is continuous.

If F is real closed, then o-minimality tells us what the definable subsets of F look like. Definable subsets of F^n are also relatively simple.

Definition 7.26 We inductively define the collection of *cells* as follows.

- $X \subseteq F^n$ is a 0-cell if it is a single point.
- $X \subseteq F$ is a 1-cell if it is an interval (a, b) , where $a \in F \cup \{-\infty\}$, $b \in F \cup \{+\infty\}$, and $a < b$.
- If $X \subseteq F^n$ is an n -cell and $f : X \rightarrow F$ is a continuous definable function, then $Y = \{(\bar{x}, f(\bar{x})) : \bar{x} \in X\}$ is an n -cell.
- Let $X \subseteq F^n$ be an n -cell. Suppose that f is either a continuous definable function from X to F or identically $-\infty$ and g is either a continuous definable function from X to F such that $f(\bar{x}) < g(\bar{x})$ for all $\bar{x} \in X$ or g is identically $+\infty$; then

$$Y = \{(\bar{x}, y) : \bar{x} \in X \wedge f(\bar{x}) < y < g(\bar{x})\}$$

is an $n + 1$ -cell.

In a real closed field, every nonempty definable set is a finite disjoint union of cells. The proof relies on the following lemma.

Lemma 7.27 (Uniform Bounding) *Let $X \subseteq F^{n+1}$ be semialgebraic. There is a natural number N such that if $\bar{a} \in F^n$ and $X_{\bar{a}} = \{y : (\bar{a}, y) \in X\}$ is finite, then $|X_{\bar{a}}| < N$.*

Proof First, note that $X_{\bar{a}}$ is infinite if and only if there is an interval (c, d) such that $(c, d) \subseteq X_{\bar{a}}$. Thus $\{(\bar{a}, b) \in X : X_{\bar{a}} \text{ is finite}\}$ is definable. Without loss of generality, we may assume that for all $\bar{a} \in F^n$, $X_{\bar{a}}$ is finite. In particular, we may assume that

$$F \models \forall \bar{x} \forall c \forall d \neg [c < d \wedge \forall y (c < y < d \rightarrow y \in X_{\bar{x}})].$$

Consider the following set of sentences in the language of fields with constants added for each element of F and new constants c_1, \dots, c_n . Let Γ be

$$\text{RCF} + \text{Diag}(F) + \left\{ \exists y_1, \dots, y_m \left[\bigwedge_{i < j} y_i \neq y_j \wedge \bigwedge_{i=1}^m y_i \in X_{\bar{c}} \right] : m \in \omega \right\}$$

Suppose that Γ is satisfiable. Then, there is a real closed field $K \supseteq F$ and elements $\bar{c} \in K^n$ such that $X_{\bar{c}}$ is infinite. By model-completeness, $F \prec K$. Therefore

$$K \models \forall \bar{x} \forall c, d \neg [c < d \wedge \forall y (c < y < d \rightarrow y \in X_{\bar{x}})].$$

This contradicts the o-minimality of K . Thus, Γ is unsatisfiable and there is an N such that

$$\text{RCF} + \text{Diag}(F) \models \forall \bar{x} \neg \left(\exists y_1, \dots, y_N \left[\bigwedge_{i < j} y_i \neq y_j \wedge \bigwedge_{i=1}^N y_i \in X_{\bar{x}} \right] \right).$$

In particular, for all $\bar{a} \in F^n$, $|X_{\bar{a}}| < N$.

We now state the Cell Decomposition Theorem and give the proof for subsets of F^2 . In the exercises, we will outline the results needed for the general case.

Theorem 7.28 (Cell Decomposition) *Let $X \subseteq F^m$ be semialgebraic. There are finitely many pairwise disjoint cells C_1, \dots, C_n such that $X = C_1 \cup \dots \cup C_n$.*

Proof (for $m = 2$) For each $a \in F$, let

$$C_a = \{x : \forall \epsilon > 0 \exists y, z \in (x - \epsilon, x + \epsilon) [(a, y) \in X \wedge (a, z) \notin X]\}.$$

We call C_a the *critical values* above a . By o-minimality, there are only finitely many critical values above a . By uniform bounding, there is a natural number N such that for all $a \in F$, $|C_a| \leq N$. We partition F into A_0, A_1, \dots, A_N , where $A_n = \{a : |C_a| = n\}$.

For each $n \leq N$, we have a definable function $f_n : A_1 \cup \dots \cup A_n \rightarrow F$ by $f_n(a) = n$ th element of C_a . As above, $X_a = \{y : (a, y) \in X\}$.

For $n \leq N$ and $a \in A_n$, we define $P_a \in 2^{2n+1}$, the *pattern* of X above a , as follows.

If $n = 0$, then $P_a(0) = 1$ if and only if $X_a = F$. Suppose that $n > 0$.

$P_a(0) = 1$ if and only if $x \in X_a$ for all $x < f_1(a)$.

$P_a(2i - 1) = 1$ if and only if $f_i(a) \in X$.

For $i < n$, $P_a(2i) = 1$ if and only if $x \in X_a$ for all $x \in (f_i(a), f_{i+1}(a))$.

$P_a(2n) = 1$ if and only if $x \in X_a$ for all $x > f_n(a)$.

For each possible pattern $\sigma \in 2^{2n+1}$, let $A_{n,\sigma} = \{a \in A_n : P_a = \sigma\}$. Each $A_{n,\sigma}$ is semialgebraic. For each $A_{n,\sigma}$, we will give a decomposition of $\{(x, y) \in X : x \in A_{n,\sigma}\}$ into disjoint cells. Because the $A_{n,\sigma}$ partition F , this will suffice.

Fix one $A_{n,\sigma}$. By Corollary 7.25, we can partition $A_{n,\sigma} = C_1 \cup \dots \cup C_l$, where each C_j is either an interval or a singleton and f_i is continuous on C_j for $i \leq n, j \leq l$. We can now give a decomposition of $\{(x, y) : x \in A_{n,\sigma}\}$ into cells such that each cell is either contained in X or disjoint from X .

For $j \leq l$, let $D_{j,0} = \{(x, y) : x \in C_j \text{ and } y < f_1(x)\}$.

For $j \leq l$ and $1 \leq i \leq n$, let $D_{j,2i-1} = \{(x, f_i(x)) : x \in C_j\}$.

For $j \leq l$ and $1 \leq i < n$, let $D_{j,2i} = \{(x, y) : x \in C_j, f_i(x) < y < f_{i+1}(x)\}$.

For $j \leq l$, let $D_{j,2n} = \{(x, y) : x \in C_j, y > f_n(x)\}$.

Clearly, each $D_{j,i}$ is a cell, $\bigcup D_{j,i} = \{(x, y) : x \in A_{n,\sigma}\}$, and each $D_{j,i}$ is either contained in X or disjoint from X . Thus, taking the $D_{j,i}$ that are contained in X , we get a partition of $\{(x, y) \in X : x \in A_{n,\sigma}\}$ into disjoint cells.

o-minimal Expansions of \mathbb{R}

The proofs above used very little about semialgebraic sets beyond o-minimality. Indeed, they would work in any o-minimal expansion of the real field. Indeed, there is a rich theory of definable sets in o-minimal expansions of the reals. We

will survey some of the results in this section. For full details, see van den Dries book *Tame topology and o-minimal structures*.

Let $\mathcal{R} = (\mathbb{R}, +, \cdot, <, \dots)$ be an o-minimal expansion of the reals, i.e., a structure obtained by adding extra structure to the reals such that $\text{Th}(\mathcal{R})$ is o-minimal. Below by “definable” we will mean definable in \mathcal{R} .

Theorem 7.29 *Assume \mathcal{R} is an o-minimal expansion of \mathbb{R} .*

- i) Every definable subset of \mathbb{R}^n is a finite union of cells.*
- ii) If $f : X \rightarrow \mathbb{R}^n$ is definable, there is a finite partition of X into cells X_1, \dots, X_n such that $f|_{X_i}$ is continuous for each i . Indeed, for any $r \geq 0$, we can choose the partition such that $f|_{X_i}$ is C^r for each i .*

An easy consequence of ii) is that definable sets have only finitely many connected components. Much more is true, for example:

- Definable bounded sets can be definably triangulated.
- Suppose $X \subseteq \mathbb{R}^{n+m}$ is definable. For $a \in \mathbb{R}^m$ let

$$X_a = \{\bar{x} = (x_1, \dots, x_n) \in \mathbb{R}^n : (\bar{x}, a) \in X\}.$$

There are only finitely many definable homeomorphism types for the sets X_a .

- (Curve selection) If $X \subseteq \mathbb{R}^n$ is definable and a is in the closure of X , then there is a continuous definable $f : (0, 1) \rightarrow X$ such that

$$\lim_{x \rightarrow 1} f(x) = a.$$

- If G is a definable group, then G is definably isomorphic to a Lie group.
- If we assume in addition that all definable functions are majorized by polynomials, then many of the metric properties of semialgebraic sets and asymptotic properties of semialgebraic functions also generalize.

Of course, this leads to the question: are there interesting o-minimal expansions of \mathbb{R} ?

\mathbb{R}_{an} and subanalytic sets

Most of the results on o-minimal structures mentioned above were proved before we knew of any interesting o-minimal structures other than the real field. The first new example of an o-minimal theory was given by van den Dries.

Let $\mathcal{L}_{\text{an}} = \mathcal{L} \cup \{\hat{f} : \text{for some open } U \supset [0, 1]^n, f : U \rightarrow \mathbb{R} \text{ is analytic}\}$.

We define $\hat{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$\hat{f}(x) = \begin{cases} f(x) & x \in [0, 1]^n \\ 0 & \text{otherwise.} \end{cases}$$

We let \mathbb{R}_{an} be the resulting \mathcal{L}_{an} -structure. Denef and van den Dries proved that \mathbb{R}_{an} is o-minimal and that \mathbb{R}_{an} has quantifier elimination if we add a function

$$D(x, y) = \begin{cases} x/y & \text{if } 0 \leq |x| \leq |y| \\ 0 & \text{otherwise} \end{cases}$$

to the language. Quantifier elimination is proven by using the Weierstrass preparation theorem to replace arbitrary analytic functions of several variables by analytic functions that are polynomial in one of the variables. Tarski's elimination procedure is then used to eliminate this variable.

Denef and van den Dries also showed that if $f : \mathbb{R} \rightarrow \mathbb{R}$ is definable in \mathbb{R}_{an} , then f is asymptotic to a rational function. In particular, although we can define the restriction of the exponential function to bounded intervals, we cannot define the exponential function globally. It is also impossible to define the sine function globally; for its zero set would violate o-minimality.

Although \mathbb{R}_{an} may seem unnatural, the definable sets form an interesting class.

We say that $X \subseteq \mathbb{R}^n$ is *semi-analytic* if for all x in \mathbb{R}^n there is an open neighborhood U of x such that $X \cap U$ is a finite Boolean combination of sets $\{\bar{x} \in U : f(\bar{x}) = 0\}$ and $\{\bar{x} \in U : g(\bar{x}) > 0\}$ where $f, g : U \rightarrow \mathbb{R}$ are analytic. We say that $X \subseteq \mathbb{R}^n$ is *subanalytic* if for all x in \mathbb{R}^n there is an open U and $Y \subset \mathbb{R}^{n+m}$ a bounded semianalytic set such that $X \cap U$ is the projection of Y into U . It is well known that subanalytic sets share many of the nice properties of semialgebraic sets.

If $X \subset \mathbb{R}^n$ is bounded, then X is definable in \mathbb{R}_{an} if and only if X is subanalytic. Indeed $Y \subseteq \mathbb{R}^n$ is definable in \mathbb{R}_{an} if and only if it is the image of a bounded subanalytic set under a semialgebraic map. Most of the known properties of subanalytic sets generalize to sets defined in any polynomial bounded o-minimal theory.

Exponentiation

The big breakthrough in the subject came in 1991. While quantifier elimination for \mathbb{R}_{exp} is impossible, Wilkie proved the next best thing.

Theorem 7.30 (Wilkie) *Let $\phi(x_1, \dots, x_m)$ be an \mathcal{L}_{exp} formula. Then there is $n \geq m$ and $f_1, \dots, f_s \in \mathbb{Z}[x_1, \dots, x_n, e^{x_1}, \dots, e^{x_n}]$ such that $\phi(x_1, \dots, x_n)$ is equivalent to*

$$\exists x_{m+1} \dots \exists x_n f_1(x_1, \dots, x_n, e^{x_1}, \dots, e^{x_n}) = \dots = f_s(x_1, \dots, x_n, e^{x_1}, \dots, e^{x_n}) = 0.$$

Thus every formula is equivalent to an existential formula (this property is equivalent to model completeness) and every definable set is the projection of an exponential variety.

Wilkie's proof depends heavily on the following special case of a theorem of Khovanski. Before Wilkie's theorem, Khovanski's result was the best evidence that \mathbb{R}_{exp} is o-minimal; indeed Khovanski's theorem is also the crucial tool needed to deduce o-minimality from model completeness.

Theorem 7.31 (Khovanski) *If $f_1, \dots, f_m : \mathbb{R}^n \rightarrow \mathbb{R}$ are exponential polynomials, then $\{x \in \mathbb{R}^n : f_1(x) = \dots = f_m(x) = 0\}$ has finitely many connected components.*

If $X \subseteq \mathbb{R}$ is definable in \mathbb{R}_{exp} then by Wilkie's Theorem there is an exponential variety $V \subseteq \mathbb{R}^n$ such that X is the projection of V . By Khovanski's Theorem V has finitely many connected components and X is a finite union of points and intervals. Thus \mathbb{R}_{exp} is o-minimal.

Using the o-minimality of \mathbb{R}_{exp} one can improve some of Khovanski's results on "fewnomials". From algebraic geometry we know that we can bound the number of connected components of a hypersurface in \mathbb{R}^n uniformly in the degree of the defining polynomial. Khovanski showed that it is also possible to bound the number of connected component uniformly in the number of monomials in the defining polynomial. We will sketch the simplest case of this. Let $\mathcal{F}_{n,m}$ be the collection of polynomials in $\mathbb{R}[X_1, \dots, X_n]$ with at most m monomials. For $p \in \mathcal{F}_{n,m}$ let

$$V^+(p) = \{\bar{x} = (x_1, \dots, x_n) \in \mathbb{R}^n : \bigwedge_{i=1}^n x_i \geq 0 \wedge p(\bar{x}) = 0\}.$$

We claim that there are only finitely many homeomorphism types of $V^+(p)$ for $p \in \mathcal{F}_{n,m}$. Let $\Phi_{m,n}(x_1, \dots, x_n, r_{1,1}, \dots, r_{1,n}, \dots, r_{m,1}, \dots, r_{m,n}, a_1, \dots, a_m)$ be the formula

$$\exists w_1, \dots, w_m \left(\bigwedge_{i=1}^m e^{w_i} = x_i \wedge \sum_{i=1}^m a_i \prod_{j=1}^n e^{w_i r_{i,j}} = 0 \right).$$

We see that Φ expresses

$$\sum_{i=1}^m a_i \prod_{j=1}^n x_j^{r_{i,j}} = 0.$$

Let $X_{\bar{r}, \bar{a}}$ denote the set of $\bar{x} \in \mathbb{R}^n$ such that $\Phi(\bar{x}, \bar{r}, \bar{a})$ holds. By o-minimality, $\{X_{\bar{r}, \bar{a}} : \bar{r} \in \mathbb{R}^{mn}, \bar{a} \in \mathbb{R}^m\}$ represents only finitely many homeomorphism types.

In addition to answering the question of o-minimality, some headway has been made on the problem of decidability. Making heavy use of Wilkie's methods and Khovanski's theorem, Macintyre and Wilkie have shown that if Schanuel's Conjecture is true then the first order theory of \mathbb{R}_{exp} is decidable. Where Schanuel's Conjecture is the assertion that if $\lambda_1, \dots, \lambda_n$ are complex numbers linearly independent over \mathbb{Q} , then the transcendence degree of the field

$$\mathbb{Q}(\lambda_1, \dots, \lambda_n, e^{\lambda_1}, \dots, e^{\lambda_n})$$

is at least n .

Miller provided an interesting counterpoint to Wilkie's theorem. Using ideas of Rosenlicht he showed that if \mathcal{R} is any o-minimal expansion of the real field that contains a function that is not majorized by a polynomial, then exponentiation is definable in \mathcal{R} .

Let $\mathcal{L}_{\text{an,exp}}$ be $\mathcal{L}_{\text{an}} \cup \{e^x\}$ and let $\mathbb{R}_{\text{an,exp}}$ be the real numbers with both exponentiation and restricted analytic functions. Using the Denef-van den Dries

quantifier elimination for \mathbb{R}_{an} and a mixture of model-theoretic and valuation theoretic ideas, van den Dries, Macintyre, and I were able to show that $\mathbb{R}_{\text{an,exp}}$ has quantifier elimination if we add \log to the language. Using quantifier elimination and Hardy field style arguments (but avoiding the geometric type of arguments used by Khovanski) we were able to show that $\mathbb{R}_{\text{an,exp}}$ is o-minimal.

Since the language $\mathcal{L}_{\text{an,exp}}$ has size 2^{\aleph_0} , one would not expect to give a simple axiomatization of the first order theory of $\mathbb{R}_{\text{an,exp}}$. Ressayre noticed that the model-theoretic analysis of $\mathbb{R}_{\text{an,exp}}$ uses very little global information about exponentiation. This observation leads to a “relative” axiomatization. The theory $\text{Th}(\mathbb{R}_{\text{an,exp}})$ is axiomatized by the theory of \mathbb{R}_{an} and axioms asserting that exponentiation is an increasing homomorphism from the additive group onto the multiplicative group of positive elements that majorizes every polynomial.

Using this axiomatization and quantifier elimination one can show that any definable function is piecewise given by a composition of polynomials, \exp , \log , and restricted analytic functions on $[0, 1]^n$. For example, the definable function $f(x) = e^{e^x} - e^{x^2} - 3x$ is eventually increasing and unbounded. Thus for some large enough $r \in \mathbb{R}$ there is a function $g : (r, +\infty) \rightarrow \mathbb{R}$ such that $f(g(x)) = x$ for $x > r$. The graph of g is the definable set $\{(x, y) : x > r \text{ and } e^{e^y} - e^{y^2} - 3y = x\}$. Thus g is a definable function and there is some way to express g explicitly as a composition of rational functions, \exp , \log , and restricted analytic functions. In most cases it is in no way clear how to get these explicit representations of an implicitly defined function. One important corollary is that every definable function is majorized by an iterated exponential.

8 The Pila-Zannier proof of the Manin-Mumford Conjecture

My goal in these lecture notes is to describe a variant of the Pila-Zannier proof of the Manin-Mumford Conjecture.³

Theorem 8.1 *Let A be an abelian variety defined over a number field. Suppose $V \subseteq A$ is an irreducible subvariety. There are finitely many cosets of algebraic subgroups $b_1 + B_1, \dots, b_n + B_n$ such that each $b_i + B_i \subseteq V$ and $V \cap \text{Tor}(A) \subseteq b_1 + \text{Tor}(B_1) \cup \dots \cup b_n + \text{Tor}(B_n)$. In particular, if V contains no cosets of infinite algebraic subgroups, then $V \cap \text{Tor}(A)$ is finite.*

The novelty of the Pila-Zannier proof is that it relies on a result of Pila and Wilkie on the asymptotics of rational points on sets definable in o-minimal structures. As such it is the only proof of Manin-Mumford that relies on real- rather than p -adic-methods.

Let $x \in \mathbb{Q}$, $x = \frac{a}{b}$ where $a, b \in \mathbb{Z}$ and $\gcd(a, b) = 1$, we define $h(x)$ the height of x to be the maximum of $|a|$ and $|b|$. If $x = (x_1, \dots, x_n) \in \mathbb{Q}^n$ we let $h(x)$ be the maximum of $h(x_1), \dots, h(x_n)$.

For $X \subseteq \mathbb{R}^n$ and $r \in \mathbb{R}$ we let $N(X, r)$ be the number of points in $X \cap \mathbb{Q}^n$ of height at most r .

For $X \subseteq \mathbb{R}^n$ we let X^{alg} be the union of all connected infinite semialgebraic subsets of X .

Theorem 8.2 (Pila-Wilkie) *Suppose $X \subseteq \mathbb{R}^n$ is definable in an o-minimal expansion of \mathbb{R} . Then for any $\epsilon > 0$ there is a constant c such that*

$$N(X \setminus X^{\text{alg}}, r) < cr^\epsilon$$

for all $r \geq 1$.

The case of Tori

As an instructive example we will first prove the theorem where we work with \mathbb{G}_m^d , a power of the multiplicative group rather than an Abelian variety.

Step 1 Move to an o-minimal setting.

Let $g : [0, 1]^d \rightarrow \mathbb{C}^d$ be defined by the function

$$g(x_1, \dots, x_d) = (2\pi i x_1, \dots, 2\pi i x_d)$$

and let $\exp : \mathbb{C}^d \rightarrow \mathbb{G}_m^d$ be the function

$$\exp(y_1, \dots, y_d) = (e^{y_1}, \dots, e^{y_d}).$$

³These lectures are based on notes of Anand Pillay.

Let $f = \exp \circ g$.

If $a \in \mathbb{G}_m$ has order n , then $a = e^{2\pi i \frac{m}{n}}$ where $0 < m < n$ and m and n are relatively prime. Thus $\text{Tor}(\mathbb{G}_m^d)$ is contained in the image of f on $[0, 1]^d \cap \mathbb{Q}^d$. Let $X = f^{-1}(V)$.

If we identify \mathbb{C}^d with \mathbb{R}^{2d} in the usual way, then

$$f(x_1, \dots, x_n) = (\cos(2\pi x_1), \sin(2\pi(x_1)), \dots, \cos(2\pi x_n), \sin(2\pi(x_n))).$$

In particular, then f and X are definable in the o-minimal structure \mathbb{R}_{an} .

Step 2 Understand X^{alg}

If $x \in X^{\text{alg}}$, then there is a connected one-dimensional semialgebraic set C such that $x \in C$. By quantifier elimination it is easy to see the C is a piece of a real algebraic curve.

Our analysis will use Ax's differential field version of Schanuel's Conjecture.

Theorem 8.3 (Ax) *Let (K, δ) be a differential field with constants k . Suppose $y_1, \dots, y_n, z_1, \dots, z_n \in K$ such that $\delta(y_i) = \frac{\delta(z_i)}{z_i}$ for $i = 1, \dots, n$. Suppose the transcendence degree of $k(y_1, \dots, y_n, z_1, \dots, z_n)$ over k is at most n , then there are integers m_1, \dots, m_n such that:*

i)

$$\prod_{i=1}^n z_i^{m_i} \in k,$$

and

ii)

$$\sum_{i=1}^n m_i y_i \in k.$$

If B is an infinite irreducible algebraic subgroup of \mathbb{G}_m^d . There is a $k \times d$ integer matrix $M = (a_{i,j})$ such that

$$z \in B \Leftrightarrow \prod_{j=1}^d z_j^{a_{i,j}} = 1, \text{ for } i = 1, \dots, k.$$

Define $LB \subseteq \mathbb{C}^d$, $LB = \{y : My = 0\}$.

Suppose $C \subseteq X$ is a connected one-dimensional semialgebraic set and x is a generic point of C , in the sense of the o-minimal structure \mathbb{R}_{an} . As above, let $y = g(x)$ and $z = \exp(y)$.

Lemma 8.4 *Let B be a minimal irreducible algebraic subgroup of \mathbb{G}_m^d such that $y \in b + LB$ for some $b \in \mathbb{C}^d$. Then the transcendence degree of $\mathbb{C}(z)$ over \mathbb{C} is the dimension of B and $\exp(b) + B$ is contained in V .⁴*

⁴To make the transition to the case of Abelian varieties smoother, we abuse notation and write cosets in \mathbb{G}_m^d additively.

Proof Let l be the dimension of LB . Since x is a generic point we must have $l > 0$. We may, without loss of generality, assume that y_1, \dots, y_l satisfies no equation $\sum m_i y_i = c$ where $m_i \in \mathbb{Z}$ and $c \in \mathbb{C}$. We claim that z_1, \dots, z_l are algebraically independent over \mathbb{C} . Suppose not. Since $x \in C$, $\text{td}(y/\mathbb{C}) = \text{td}(x/\mathbb{C}) = 1$. Then $\text{td}(y_1, \dots, y_l, z_1, \dots, z_l/\mathbb{C})$ is at most n . Thus by Ax, there are m_1, \dots, m_l such that $\sum m_i y_i \in \mathbb{C}$, a contradiction.

Thus z is an (algebraic) generic point of $\exp(b)+B$. Since $z \in V$, $\exp(b)+B \subseteq V$.

Corollary 8.5 *If $a \in X^{\text{alg}} \cap \mathbb{Q}^d$, then $f(a) \in b + B$ where b is a torsion point of A , B is an infinite algebraic subgroup of A and $b + B \subseteq V$.*

Step 3 Finiteness of $\mathbb{Q}^d \cap X \setminus X^{\text{alg}}$

We may, without loss of generality, assume that $V \cap \text{Tor}(\mathbb{G}_m^d)$ is Zariski dense in V . If not, then we can proceed by first proving the result for each irreducible component of the Zariski closure of $V \cap \text{Tor}(\mathbb{G}_m^d)$. Then, any automorphism of \mathbb{C} that fixes the roots of unity will fix V . Thus V is defined over a number field k . We may assume that k is a Galois extension of \mathbb{Q} of degree l .

Suppose $a = (a_1, \dots, a_d) \in V \cap \text{Tor}(\mathbb{G}_m^d)$. If σ is an automorphism of \mathbb{C} , fixing k , then $\sigma(a) \in V \cap \text{Tor}(\mathbb{G}_m^d)$.

Let a_1 have order exactly n_i then a has order n where n is the least common multiple of n_1, \dots, n_m . If b is a primitive n^{th} -root of unity, then $\mathbb{Q}(b) = \mathbb{Q}(a_1, \dots, a_d)$. Thus the degree of $k(a)/k$ is at most $\phi(n)$ and at least $\phi(n)/l$, where $\phi(n)$ is Euler's function, i.e.,

$$\phi(n) = |\{m : 1 \leq x < n, x \text{ relatively prime to } n\}|.$$

The asymptotics of $\phi(n)$ are well understood. In particular, for any $0 < \epsilon < 1$,

$$n^\epsilon < \phi(n)$$

for large enough n .⁵ In particular there is M such that if $a \in \text{Tor}(\mathbb{G}_m^d)$ is a torsion point of order $n > M$, then a has at least $\frac{n^{1/2}}{l}$ conjugates over k .

Suppose $a \in V \cap \text{Tor}(\mathbb{G}_m^d)$ is not in an infinite coset $b + B$ where $b + B \subseteq V$. Then the same is true of all conjugates of a over k . If a has order n , there is $x \in (X \setminus X^{\text{alg}}) \cap \mathbb{Q}^d$ such that $f(x) = a$ and $h(x) = n$. Thus if $(X \setminus X^{\text{alg}}) \cap \mathbb{Q}^d$ is infinite, then

$$N(X \setminus X^{\text{alg}}, n) \geq \frac{n^{1/2}}{l}$$

for infinitely many n . But this contradicts the Pila-Wilkie Theorem.

Corollary 8.6 *There is a finite set F such that every element of $V \cap \text{Tor}(\mathbb{G}_m^d)$ is either contained in F or contained in $b + B$ where $b \in \text{Tor}(\mathbb{G}_m^d)$ and B is an infinite irreducible algebraic subgroup of \mathbb{G}_m^d .*

⁵Better bounds can be found using the Prime Number Theorem.

We say that an infinite irreducible coset $b + B$ is *maximal* $b + B \subseteq V$ and there is no irreducible algebraic subgroup $C \supset B$ with $b + C \subseteq V$. We need to show there are only finitely many maximal cosets $b + B \subseteq V$.

Step 4 Finitely many choices for B .

Bombieri and Zannier proved, in the Abelian variety case, that there are only finitely many B such that $b + B$ is a maximal coset in V . This follows from the next two lemmas.

Lemma 8.7 *For any M there are only finitely many subgroups of \mathbb{G}_m^d of degree M .*

Proof For any dimension $m < d$, there is a definable family $(W_y : y \in Y)$ of all subvarieties of \mathbb{G}_m^d of dimension m and degree M . There is a definable $Y_0 \subseteq Y$ such that $y \in Y_0$ if and only if W_y is a subgroup. Since semi-abelian varieties have no infinite definable families of subgroups Y_0 is finite.⁶

Lemma 8.8 *There is a number M , depending on the dimension and the degree of V such that if $b + B$ is a maximal coset then the degree of B is at most M .*

Proof Suppose $b + B$ is a maximal coset. We build a sequence of subvarieties $V = V_1 \supset V_2 \supset \dots \supset V_m$ as follows. Given V_i if there is $g \in B$ such that

$$\dim(V_i \cap V_i + g) < \dim V_i,$$

then choose some such g and let V_{i+1} be an irreducible component of the intersection containing $b + B$. If there is no such $g \in B$, we let $m = i$. Let $W = V_i$. Since $m < \dim V$, we can bound $\deg W$ in terms of the dimension and degree of W .

Note that $b \in W$ and $B + W = W$.

We next build a sequence $W = W_1 \supset W_2 \supset \dots \supset W_m$ such that $b + B \subseteq W_i$ and $B + W_i = W + i$ for all i . Start with W_i . If there is $x \in W_i$ such that

$$\dim(W_i \cap (b - x) + W_i) < \dim W_i$$

then we choose some such x . Let Y_1, \dots, Y_m be the irreducible components of $W_i \cap (b - x) + W_i$. Note that $b + B \subseteq (b - x) + W_i$. Thus $b + B$ is contained in one irreducible component, say Y_1 . Let $B_0 = \{b \in B : b + Y_1 = Y_1\}$. Then B_0 is a finite index subgroup of B . But B is irreducible, so $B = B_0$ and $B + Y_1 = Y_1$. Let $Y_1 = W_{i+1}$. If there is no such x , we let $m = i$ and stop.

Let $Z = W_m$. Once again, we can bound the degree of Z in terms of the dimension and degree of V . We also have that Z is irreducible, $b + B \subseteq Z$, $Z + B = B$ and $b - z + Z = Z$ for all $z \in Z$.

Let $C = \{a \in \mathbb{G}_m^d : a + Z = Z\}$. Then C is an algebraic subgroup of A and $B \subseteq C$. Since $C + Z = Z$ and $b \in Z$, $b + C \subseteq Z \subseteq V$, thus, by the maximality of V , C is a finite union of B cosets.

⁶This is really much easier for \mathbb{G}_m^d where we can easily describe the subgroups.

On the other hand, $B \subseteq b - Z$, while, by construction of Z , $b - Z \subseteq C$. Since Z is irreducible we must have $B = b - Z$. Thus we can bound the degree of B .

We can now finish the proof. We claim that for any infinite irreducible subgroup B , there are only finitely many maximal cosets $b + B \subseteq V$ where b is a torsion point of A .

Suppose for contradiction that there are infinitely many maximal torsion cosets $b + B \subseteq V$. Consider the projection map $\pi : A \rightarrow A/B$. Let $W = \{a : a + B \subseteq V\}$. Let W' be the projection of W . If V contains infinitely many maximal torsion cosets $b + B$, then W' contains infinitely many torsion points. By the arguments above we can find $b \in W$ such that $b + B$ is a maximal coset and $\pi(b)$ is contained in an infinite torsion coset $\pi(b) + C$ of W' . But that $b + \pi^{-1}(C)$ is a coset in V with $\pi^{-1}(C) \supset B$, contradicting the maximality of B .

Abelian Varieties

We outline the changes that need to be made to adapt the argument for Abelian varieties rather than the multiplicative group.

Step 1

We let $\exp_A : \mathbb{C}^d \rightarrow A$ be the usual exponential map. Let $\Lambda = \bigoplus_{i=1}^{2d} \mathbb{Z}\lambda_i$ be the kernel of \exp_A . Let $g : [0, 1]^{2d} \rightarrow \mathbb{C}^d$ be the map

$$g(x_1, \dots, x_{2n}) = \sum_{i=1}^n x_i \lambda_i$$

and let f be the composition $\exp_A \circ g$. Once again, we can view f as a function definable in \mathbb{R}_{an} and $\text{Tor}(A)$ is contained in the image of $[0, 1]^{2d} \cap \mathbb{Q}^{2d}$.

Step 2

We need the extension of Ax's theorem for abelian, or semiabelian varieties defined over the constants due independently to Bertrand and Kirby.

Theorem 8.9 (Bertrand/Kirby) *Let K be a differential field with constants k . Suppose A is a semiabelian variety defined over k with Lie algebra LA . Let $l_A : A \rightarrow LA$ and $l_{LA} : LA \rightarrow LA$ be the logarithmic derivatives. Suppose $(y, z) \in LA \times A$ with $l_{LA}y = l_Az$. If $\text{td}(y, z/k) \leq \dim A + 1$, then there is a proper algebraic $B \leq A$ such that:*

- i) $z \in B + b$, for some $b \in A(k)$;*
- ii) $y \in LB + c$, for some $c \in LA(k)$.*

Step 3

We need the following Theorem of Masser.

Theorem 8.10 (Masser) *Suppose A is an abelian variety defined over a number field k . There is $l > 0$, $c > 0$ and $N > 0$ such that if a is a torsion point of A of order $n \geq N$, then the degree of a over k is at least $cn^{1/l}$.*

The remainder of the proof is as above.

Semiabelian Varieties

Let's consider the case where G is a semiabelian variety defined over a number field. Suppose \mathbb{G}_m^d is a subgroup of G and the projection map $\pi : G \rightarrow A$ has kernel \mathbb{G}_m^d and A is an abelian variety. We suppose that G , A and π are all defined over a number field k .

The n -torsion subgroup of G is of the form $B \oplus C$ where B is the n -torsion of A and π maps C isomorphically onto the n -torsion of A . If $g \in G$ has order n , then $g = b + c$ where $b \in B$ has order n_1 , $c \in C$ has order n_2 and n is the least common multiple of n_1 and n_2 . At least one of n_1 and n_2 is at least \sqrt{n} . Suppose $n_2 \geq \sqrt{n}$. For n large enough and c and l as in Theorem 8.10, $\pi(c)$ has at least $cn^{\frac{1}{2l}}$ conjugates over k . then the same is true of c and g . The argument is similar if $n_1 \geq \sqrt{n}$.

Questions

- Raynaud showed, using specialization arguments, that one could deduce the general version of Manin-Mumford, from the number field version. Masser's Theorem is the one place we used the number field assumption. Are there extensions of Masser's Theorem that would allow us to deduce the general result by these methods?

A Real Algebra

We prove some of the algebraic facts needed in Section 7. All of these results are due to Artin and Schreier. See Lang's *Algebra* §XI for more details.

All fields are assumed to be of characteristic 0.

Definition A.1 A field K is *real* if -1 can not be expressed as a sum of squares of elements of K . In general, we let $\sum K^2$ be the sums of squares from K .

If F is orderable, then F is real because squares are nonnegative with respect to any ordering.

Lemma A.2 *Suppose that F is real and $a \in F \setminus \{0\}$. Then, at most one of a and $-a$ is a sum of squares.*

Proof If a and b are both sums of squares, then $\frac{a}{b} = \frac{a}{b^2}b$ is a sum of squares. Thus, if F is real, at least one of a and $-a$ is not in $\sum F^2$.

Lemma A.3 *If F is real and $-a \in F \setminus \sum F^2$, then $F(\sqrt{a})$ is real. Thus, if F is real and $a \in F$, then $F(\sqrt{a})$ is real or $F(\sqrt{-a})$ is real.*

Proof We may assume that $\sqrt{a} \notin F$. If $F(\sqrt{a})$ is not real, then there are $b_i, c_i \in F$ such that

$$-1 = \sum (b_i + c_i\sqrt{a})^2 = \sum (b_i^2 + 2c_i b_i\sqrt{a} + c_i^2 a).$$

Because \sqrt{a} and 1 are a vector space basis for $F(\sqrt{a})$ over F ,

$$-1 = \sum b_i^2 + a \sum c_i^2.$$

Thus

$$-a = \frac{1 + \sum b_i^2}{\sum c_i^2} = \frac{(\sum b_i^2)(\sum c_i^2) + (\sum c_i^2)}{(\sum c_i^2)^2}$$

and $-a \in \sum F^2$, a contradiction.

Lemma A.4 *If F is real, $f(X) \in F[X]$ is irreducible of odd degree n , and $f(\alpha) = 0$, then $F(\alpha)$ is real.*

Proof We proceed by induction on n . If $n = 1$, this is clear. Suppose, for purposes of contradiction, that $n > 1$ is odd, $f(X) \in F[X]$ is irreducible of degree n , $f(\alpha) = 0$, and $F(\alpha)$ is not real. There are polynomials g_i of degree at most $n-1$ such that $-1 = \sum g_i(\alpha)^2$. Because F is real, some g_i is nonconstant. Because $F(\alpha) \cong F[X]/(f)$, there is a polynomial $q(X) \in F[X]$ such that

$$1 = \sum g_i^2(X) + q(X)f(X).$$

The polynomial $\sum g_i^2(X)$ has a positive even degree at most $2n - 2$. Thus, q has odd degree at most $n - 2$. Let β be the root of an irreducible factor of q . By induction, $F(\beta)$ is real, but $-1 = \sum g_i^2(\beta)$, a contradiction.

Definition A.5 We say that a field R is *real closed* if and only if R is real and has no proper real algebraic extensions.

If R is real closed and $a \in R$, then, by Lemmas A.2 and A.3, either $a \in R^2$ or $-a \in R^2$. Thus, we can define an order on R by

$$a \geq 0 \Leftrightarrow a \in R^2.$$

Moreover, this is the only way to define an order on R because the squares must be nonnegative. Also, if R is real closed, every polynomial of odd degree has a root in R .

Lemma A.6 *Let F be a real field. There is $R \supseteq F$ a real closed algebraic extension. We call R a real closure of F .*

Proof Let $I = \{K \supseteq F : K \text{ real, } K/F \text{ algebraic}\}$. The union of any chain of real fields is real; thus, by Zorn's Lemma, there is a maximal $R \in I$. Clearly, R has no proper real algebraic extensions; thus, R is real closed.

Corollary A.7 *If F is any real field, then F is orderable. Indeed, if $a \in F$ and $-a \notin \sum F^2$, then there is an ordering of F , where $a > 0$.*

Proof By Lemma A.3, $F(\sqrt{a})$ is real. Let R be a real closure of F . We order F by restricting the ordering of R because a is a square in R , $a > 0$.

The following theorem is a version of the Fundamental Theorem of Algebra.

Theorem A.8 *Let R be a real field such that*
i) for all $a \in R$, either \sqrt{a} or $\sqrt{-a} \in R$ and
ii) if $f(X) \in R[X]$ has odd degree, then f has a root in R .
If $i = \sqrt{-1}$, then $K = R(i)$ is algebraically closed.

Proof

Claim 1 Every element of K has a square root in K .

Let $a + bi \in K$. Note that $\frac{a + \sqrt{a^2 + b^2}}{2}$ is nonnegative for any ordering of R . Thus, by i), there is $c \in R$ with

$$c^2 = \frac{a + \sqrt{a^2 + b^2}}{2}.$$

If $d = \frac{b}{2c}$, then $(c + di)^2 = a + bi$.

Let $L \supseteq K$ be a finite Galois extension of R . We must show that $L = K$. Let $G = \text{Gal}(L/R)$ be the Galois group of L/R . Let H be the 2-Sylow subgroup of G .

Claim 2 $G = H$.

Let F be the fixed field of H . Then F/R must have odd degree. If $F = R(x)$, then the minimal polynomial of x over R has odd degree, but the only irreducible polynomials of odd degree are linear. Thus, $F = R$ and $G = H$.

Let $G_1 = \text{Gal}(L/K)$. If G_1 is nontrivial, then there is G_2 a subgroup of G_1 of index 2. Let F be the fixed field of G_2 . Then, F/K has degree 2. But by Claim 1, K has no extensions of degree 2. Thus, G_1 is trivial and $L = K$.

Corollary A.9 *Suppose that R is real. Then R is real closed if and only if $R(i)$ is algebraically closed.*

Proof

(\Rightarrow) By Theorem A.8.

(\Leftarrow) $R(i)$ is the only algebraic extension of R , and it is not real.

Let $(R, <)$ be an ordered field. We say that R has the *intermediate value property* if for any polynomial $p(X) \in R[X]$ if $a < b$ and $p(a) < 0 < p(b)$, then there is $c \in (a, b)$ with $p(c) = 0$.

Lemma A.10 *If $(R, <)$ is an ordered field with the intermediate value property, then R is real closed.*

Proof Let $a > 0$ and let $p(X) = X^2 - a$. Then $p(0) < 0$, and $p(1+a) > 0$; thus, there is $c \in R$ with $c^2 = a$.

Let

$$f(X) = X^n + \sum_{i=0}^{n-1} a_i X^i$$

where n is odd. For M large enough, $f(M) > 0$ and $f(-M) < 0$; thus, there is a c such that $f(c) = 0$.

By Theorem A.8, $R(i)$ is algebraically closed. Because R is real, it must be real closed.

Lemma A.11 *Suppose that R is real closed and $<$ is the unique ordering, then $(R, <)$ has the intermediate value property.*

Proof Suppose $f(X) \in R[X]$, $a < b$, and $f(a) < 0 < f(b)$. We may assume that $f(X)$ is irreducible (for some factor of f must change signs). Because $R(i)$ is algebraically closed, either $f(X)$ is linear, and hence has a root in (a, b) , or

$$f(X) = X^2 + cX + d,$$

where $c^2 - 4d < 0$. But then

$$f(X) = \left(X + \frac{c}{2}\right)^2 + \left(d - \frac{c^2}{4}\right)$$

and $f(x) > 0$ for all x .

We summarize as follows.

Theorem A.12 *The following are equivalent.*

- i) R is real closed.
- ii) For all $a \in R$, either a or $-a$ has a square root in R and every polynomial of odd degree has a root in R .
- iii) We can order R by $a \geq 0$ if and only if a is a square and, with respect to this ordering, R has the intermediate value property.

Finally, we consider the question of uniqueness of real closures. We first note that there are some subtleties. For example, there are nonisomorphic real closures of $F = \mathbf{Q}(\sqrt{2})$. The field of real algebraic numbers is one real closure of F . Because $a + b\sqrt{2} \mapsto a - b\sqrt{2}$ is an automorphism of F , $\sqrt{2}$ is not in $\sum F^2$. Thus, by Corollary B.5, $F(\sqrt{-2})$ is real. Let R be a real closure of F containing $F(\sqrt{-2})$. Then, R is not isomorphic to the real algebraic numbers over F .

This is an example of a more general phenomenon. It is proved by successive applications of Lemmas A.2 and A.3.

Lemma A.13 *If $(F, <)$ is an ordered field, then there is a real closure of F in which every positive element of F is a square.*

Because $\mathbf{Q}(\sqrt{2})$ has two distinct orderings, it has two nonisomorphic real closures. The field $\mathbf{Q}(t)$ of rational functions over \mathbf{Q} has 2^{\aleph_0} orderings and hence 2^{\aleph_0} nonisomorphic real closures.

The next theorem shows that once we fix an ordering of F , there is a unique real closure that induces the ordering.

Theorem A.14 *Let $(F, <)$ be an ordered field. Let R_0 and R_1 be real closures of F such that $(R_i, <)$ is an ordered field extension of $(F, <)$. Then, R_0 is isomorphic to R_1 over F and the isomorphism is unique.*

The proof of Theorem A.14 uses Sturm's algorithm.

Definition A.15 Let R be a real closed field. A *Sturm sequence* is a finite sequence of polynomials f_0, \dots, f_n such that:

- i) $f_1 = f'_0$;
- ii) for all x and $0 \leq i \leq n - 1$, it is not the case that $f_i(x) = f_{i+1}(x) = 0$;
- iii) for all x and $1 \leq i \leq n - 1$, if $f_i(x) = 0$, then $f_{i-1}(x)$ and $f_{i+1}(x)$ have opposite signs;
- iv) f_n is a nonzero constant.

If f_0, \dots, f_n is a Sturm sequence and $x \in \mathbb{R}$, define $v(x)$ to be the number of sign changes in the sequence $f_0(x), \dots, f_n(x)$.

Suppose that $f \in R[X]$ is nonconstant and does not have multiple roots. We define a Sturm sequence as follows:

$$\begin{aligned} f_0 &= f; \\ f_1 &= f'. \end{aligned}$$

Given f_i nonconstant, use the Euclidean algorithm to write

$$f_i = g_i f_{i-1} - f_{i+1}$$

where the degree of f_{i+1} is less than the degree of f_{i-1} . We eventually reach a constant function f_n .

Lemma A.16 *If f has no multiple roots, then f_0, \dots, f_n is a Sturm sequence.*

Proof

iv) If $f_n = 0$, then $f_{n-1} | f_i$ for all i . But f has no multiple roots; thus f and f' have no common factors, a contradiction.

ii) If $f_i(x) = f_{i+1}(x) = 0$, then by induction $f_n(x) = 0$, contradicting iv).

iii) If $1 \leq i \leq n-1$ and $f_i(x) = 0$, then $f_{i-1}(x) = -f_{i+1}(x)$. Thus, $f_{i-1}(x)$ and $f_{i+1}(x)$ have opposite signs.

Theorem A.17 (Sturm's Algorithm) *Suppose that R is a real closed field, $a, b \in R$, and $a < b$. Let f be a polynomial without multiple roots. Let $f = f_0, \dots, f_n$ be a Sturm sequence such that $f_i(a) \neq 0$ and $f_i(b) \neq 0$ for all i . Then, the number of roots of f in (a, b) is equal to $v(a) - v(b)$.*

Proof Let $z_1 < \dots < z_m$ be all the roots of the polynomials f_0, \dots, f_n that are in the interval (a, b) . Choose c_1, \dots, c_{m-1} with $z_i < c_i < z_{i+1}$. Let $a = c_0$ and $b = c_m$. For $0 \leq i \leq m-1$, let r_i be the number of roots of f in the interval (c_i, c_{i+1}) . Clearly, $\sum r_i$ is the number of roots of f in the interval (a, b) . On the other hand,

$$v(a) - v(b) = \sum_{i=0}^{m-1} (v(c_i) - v(c_{i+1})).$$

Thus, it suffices to show that if $c < z < d$ and z is the only root of any f_i in (c, d) , then

$$v(d) = \begin{cases} v(c) - 1 & z \text{ is a root of } f \\ v(c) & \text{otherwise} \end{cases}.$$

If $f_i(b)$ and $f_i(c)$ have different signs, then $f_i(z) = 0$. We need only see what happens at those places.

If z is a root of f_i , $i > 0$, then $f_{i+1}(z)$ and $f_{i-1}(z)$ have opposite signs and f_{i+1} and f_{i-1} do not change signs on $[c, d]$. Thus, the sequences $f_{i-1}(c), f_i(c), f_{i+1}(c)$ and $f_{i-1}(d), f_i(d), f_{i+1}(d)$ each have one sign change. For example, if $f_{i-1}(z) > 0$ and $f_{i-1}(z) < 0$, then these sequences are either $+, +, -$ or $+, -, +$, and in either case both sequences have one sign change.

If z is a root of f_0 , then, because $f'(z) \neq 0$, f is monotonic on (c, d) . If f is increasing on (c, d) , the sequence at c starts $-, +, \dots$ and the sequence at d starts $+, +, \dots$. Similarly, if f is decreasing, the sequence at c starts $+, -, \dots$, and the sequence at b starts $-, -, \dots$. In either case, the sequence at c has one more sign change than the sequence at d . Thus, $v(c) - v(d) = 1$, as desired.

Corollary A.18 *Suppose that $(F, <)$ is an ordered field. Let f be a nonconstant irreducible polynomial over F . If R_0 and R_1 are real closures of F compatible with the ordering, then f has the same number of roots in both R_0 and R_1 .*

Proof Let f_0, \dots, f_n be the Sturm sequence from Lemma A.16. Note that each $f_i \in F[X]$. We can find $M \in F$ such that any root of f_i is in $(-M, M)$ (if $g(X) = X^n + \sum a_i X^i$, then any root of g has absolute value at most $1 + \sum |a_i|$, for example). Then, the number of roots of f in R_i is equal to $v(-M) - v(M)$, but $v(M)$ depends only on F .

Lemma A.19 *Suppose $(F, <)$ is an ordered field and R_0 and R_1 are real closures of F such that $(R_i, <)$ is an ordered field extension of $(F, <)$. If $\alpha \in R_0 \setminus F$, there is an ordered field embedding of $F(\alpha)$ into R_1 fixing F .*

Proof Let $f \in F[X]$ be the minimal polynomial of α over F . Let $\alpha_1 < \dots < \alpha_n$ be all zeros of f in R_0 . By Corollary B.18, f has exactly n zeros $\beta_1 < \dots < \beta_n \in R_1$. Let

$$\sigma : F(\alpha_1, \dots, \alpha_n) \rightarrow F(\beta_1, \dots, \beta_n)$$

be the map obtained by sending α_i to β_i . We claim that σ is an ordered field isomorphism.

For $i = 1, \dots, n-1$, let $\gamma_i = \sqrt{\alpha_{i+1} - \alpha_i} \in R_0$. By the Primitive Element Theorem, there is $a \in F$ such that

$$F(a) = F(\alpha_1, \dots, \alpha_n, \gamma_1, \dots, \gamma_{n-1}).$$

Let $g \in F[X]$ be the minimal polynomial of a over F . By Corollary B.18, g has a zero $b \in R_1$ and there is a field isomorphism $\phi : F(a) \rightarrow F(b)$. Because $F(a)$ contains n zeros of F , so does $F(b)$. Thus $\beta_1, \dots, \beta_n \in F(b)$ and for each i there is a j such that $\phi(\alpha_i) = \beta_j$. But

$$\phi(\gamma_i)^2 = \phi(\alpha_{i+1}) - \phi(\alpha_i).$$

Thus $\phi(\alpha_i) = \beta_i$ for $i = 1, \dots, n$. We still must show that σ is order preserving. Suppose $c \in F(\alpha_1, \dots, \alpha_n)$ and $c > 0$. There is $d \in R_0$ such that $d^2 = c$. Arguing as above, we can find a field embedding

$$\psi : F(\alpha_1, \dots, \alpha_n, d) \subseteq R_1$$

fixing F . As above, $\psi(\alpha_i) = \beta_i$ and $\psi \supseteq \sigma$. Because

$$\psi(d)^2 = \psi(c) = \sigma(c),$$

we have $\sigma(c) > 0$. Thus σ is order preserving.

Proof of Theorem A.14 Let \mathcal{P} be the set of all order preserving $\sigma : K \rightarrow R_1$ where $F \subseteq K \rightarrow R_0$ and $\sigma|_F$ is the identity. By Zorn's Lemma, there is a maximal $\sigma : K \rightarrow R_1$ in \mathcal{P} . By identifying K and $\sigma(K)$ and applying the previous lemma, we see that $K = R_0$. A similar argument shows that $\sigma(K) = R_1$.

Uniqueness follows because the i th root of $f(X)$ in R_0 must be sent to the i th root of $f(X)$ in R_1 .