Math 563: Model Theory

Stephen G. Simpson

May 2, 1998

Department of Mathematics
The Pennsylvania State University
University Park, State College PA 16802

simpson@math.psu.edu

www.math.psu.edu/simpson/courses/math563/

Note: Chapters 12 and 13 are not finished.

Contents

Chapter 1

Sentences and models

1.1 Symbols

1.

We assume the availability of a large supply of *nonlogical symbols* of the following kinds:

1. n-ary *relation symbols* $\underline{R}(\ ,\ldots,\ )$, $n \geq 1$;
2. n-ary *operation symbols* $\underline{o}(\ ,\ldots,\ )$, $n \geq 1$;
3. *constant symbols* $\underline{c}$.

These collections of symbols are assumed to be disjoint.

2.

We make use of the following *logical symbols*:

1. *propositional connectives* $\neg$ (negation), $\wedge, \vee$ (conjunction, disjunction), $\rightarrow, \leftrightarrow$ (implication, biimplication);
2. *quantifiers* $\forall, \exists$ (universal, existential);
3. *equality* $=$;
4. *variables* $v_0, v_1, \ldots, v_n, \ldots$.

Note that $=$ is a logical symbol although syntactically it behaves as a binary relation symbol.

1.2 Formulas

1.

The notion of a *term* is defined inductively as follows. A constant symbol is a term. A variable is a term. If $t_1, \ldots, t_n$ are terms and $\underline{o}$ is an n-ary operation symbol, then $\underline{o}(t_1, \ldots, t_n)$ is a term.

2.

The notion of *atomic formula* is defined as follows. If t_1 and t_2 are terms, then $t_1 = t_2$ is an atomic formula. If $t_1, \ldots, t_n$ are terms and $\underline{R}$ is an n-ary relation symbol, then $\underline{R}(t_1, \ldots, t_n)$ is an atomic formula.

3.

The notion of a *formula* is defined inductively as follows. An atomic formula is a formula. If φ and ψ are formulas then so are $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, $\varphi \rightarrow \psi$, $\varphi \leftrightarrow \psi$. If φ is a formula and v is a variable, then $\forall v\varphi$ and $\exists v\varphi$ are formulas.

We assume familiarity with the concept of a *free variable*, i.e. one not bound by a quantifier. We assume unique readability of formulas.

4.

If S is a set of formulas and/or terms, the *signature* of S is the set of all nonlogical symbols occurring in it. This is sometimes called in the literature the *similarity type* of S. Note that $=$ never bolongs to the signature since it is a logical symbol. We write $\mathrm{sig}(S) =$ signature of S.

5.

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

Examples: The formula $\forall x \exists y (x + y = 0)$ is a sentence. Here $+$ is a binary operation symbol, 0 is a constant symbol, and $=$ is a logical symbol. The formula $x + y = y + x$ is not a sentence.

If we write $x, y, \ldots$ in the same formula, we tacitly assume that $x, y, \ldots$ are distinct variables.

Examples, continued: The formula $\forall x(\exists y(y \cdot y = x) \vee \exists z(z \cdot z = -x))$ is a sentence. It is "logically equivalent" to the sentence $\forall x \exists y(y \cdot y = x \vee y \cdot y = -x)$ but these two sentences are not identical. We asume that the student has some previous acquaintance with the syntactical and semantical notions of logical equivalence. These notions will be defined later.

1.3 Structures

1.

A *structure* is an ordered pair $\mathcal{A} = (|\mathcal{A}|, \Phi)$ where $|\mathcal{A}|$ is a nonempty set, called the *universe* of $\mathcal{A}$, and Φ is a function whose domain is a set of non-logical symbols. The domain of Φ is called the *signature* of $\mathcal{A}$. To each n-ary relation symbol $\underline{R} \in \mathrm{sig}(\mathcal{A})$ we assume that Φ assigns an n-ary relation

$$R \subseteq |\mathcal{A}|^n = \underbrace{|\mathcal{A}| \times \cdots \times |\mathcal{A}|}_{n \text{ times}} .$$

To each n-ary operation symbol $\underline{o} \in \mathrm{sig}(\mathcal{A})$ we assume that Φ assigns an n-ary operation

$$o : |\mathcal{A}|^n \to |\mathcal{A}| .$$

To each constant symbol $\underline{c} \in \mathrm{sig}(\mathcal{A})$ we assume that Φ assigns an individual constant $c \in |\mathcal{A}|$.

Example: the structure of the reals

$$\mathcal{R} = (|\mathcal{R}|, +, -, \cdot, 0, 1, <)$$

where $|\mathcal{R}| = \mathbb{R}$. We cannot include $\div$ because it is not an operation on $|\mathcal{R}|$ (because not everywhere defined). Here the universe is $|\mathcal{R}| = \mathbb{R} = (-\infty, \infty)$; $+, \cdot$ are binary operations; $-$ is a unary operation; $0, 1$ are constants; $<$ is a binary relation.

1.4 Truth

1.

Given a structure $\mathcal{A}$ and a sentence σ such that $\mathrm{sig}(\sigma) \subseteq \mathrm{sig}(\mathcal{A})$, we assume known the meaning of

$$\mathcal{A} \models \sigma \qquad (\mathcal{A} \textit{ satisfies } \sigma,\ \sigma \text{ is } \textit{true} \text{ in } \mathcal{A})\,.$$

For example, $\mathcal{R} \models \forall x \exists y (y \cdot y = x \vee y \cdot y = -x)$ expresses the fact that every real number or its negative is a square. Note that the structure

$$\mathcal{Z} = (\mathbb{Z}, +, \cdot, -, 0, 1, <)\,,$$

where $|\mathcal{Z}| = \mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$, has the same signature as $\mathcal{R}$ but satisfies the negation of the above sentence.

In general, $\mathcal{A} \models \sigma$ means that σ is true in $\mathcal{A}$ when the variables are interpreted as ranging over $|\mathcal{A}|$, the other symbols in σ being given their obvious interpretation.

Another example:

$$\mathcal{Z} \models \forall x (x > 0 \to \exists y_1 \exists y_2 \exists y_3 \exists y_4 (x = y_1^2 + y_2^2 + y_3^2 + y_4^2))$$

and this expresses the fact that every positive integer is the sum of four squares.

1.5 Models and theories

1.

Let S be a set of sentences. A *model* of S is a structure $\mathcal{M}$ such that $\mathcal{M} \models \sigma$ for all $\sigma \in S$, and $\mathrm{sig}(\mathcal{M}) = \mathrm{sig}(S)$.

For example, a *group* can be described as a model $\mathcal{G} = (|\mathcal{G}|, \cdot, ^{-1}, 1)$ of the *axioms of group theory*:

$$\forall x \forall y \forall z ((x \cdot y) \cdot z = x \cdot (y \cdot z))$$
$$\forall x (x \cdot 1 = 1 \cdot x = x)$$
$$\forall x (x \cdot x^{-1} = x^{-1} \cdot x = 1)$$

2.

The class of all models of S is denoted $\mathrm{Mod}(S)$. A sentence τ is said to be a *logical consequence*[1] of S (written $S \models \tau$) if $\mathrm{sig}(\tau) \subseteq \mathrm{sig}(S)$, and $\mathcal{M} \models \tau$ for all $\mathcal{M} \in \mathrm{Mod}(S)$.

A *theory* is a set T of sentences which is consistent and closed under logical consequence; in other words, T has at least one model, and $\tau \in T$ whenever τ is a sentence such that $\mathrm{sig}(\tau) \subseteq \mathrm{sig}(T)$ and $\mathcal{M} \models \tau$ for all $\mathcal{M} \in \mathrm{Mod}(T)$.

For example, the *theory of groups* is the set of all logical consequences of the axioms of group theory. These axioms have many nonobvious logical consequences, e.g. the Jacobi identity

$$((x, y), z^x) \cdot ((y, z), x^y) \cdot ((z, x), y^z) = 1$$

where we use abbreviations $(x, y) = x^{-1} \cdot y^{-1} \cdot x \cdot y$ and $x^y = y^{-1} \cdot x \cdot y$.

3.

A *model class* is a nonempty class of structures all having the same signature. An *elementary model class* is a model class of the form $\mathrm{Mod}(S)$ where S is a consistent set of sentences. There are lots of nonelementary model classes, e.g. the class of finite groups.

If K is a model class, we write $\mathrm{Th}(K)$ for the *theory of* K, i.e. the set of sentences σ such that $\mathrm{sig}(\sigma) \subseteq \mathrm{sig}(K)$ and $\mathcal{M} \models \sigma$ for all $\mathcal{M} \in K$. Note that $\mathrm{Th}(K)$ is a theory and for any theory T we have $T = \mathrm{Th}(\mathrm{Mod}(T))$. There is a natural 1-1 correspondence between theories and elementary model classes.

4.

If T is a theory and $S \subseteq T$, we say that S is a *set of axioms* for T if $T = \mathrm{Th}(\mathrm{Mod}(S))$. If there exists a finite set of axioms for T, we say that T is *finitely axiomatizable.*

For example, the theory of groups is finitely axiomatizable (a finite set of axioms for it is displayed above). We shall see later that the theory of fields of characteristic 0 is not finitely axiomatizable.

---

[1] Note: This definition is somewhat unusual because, for instance, $\forall x(x < 0 \lor x = 0 \lor x > 0)$ is *not* a logical consequence of $\forall x \forall y(x < y \lor x = y \lor x > y)$, owing to the restriction on signature.

Chapter 2

Complete theories

2.1 Definitions and examples

1.

A theory T is *complete* if for all sentences σ, either $\sigma \in T$ or $\neg\sigma \in T$, provided $\mathrm{sig}(\sigma) \subseteq \mathrm{sig}(T)$.

Examples: The theory of groups is not complete. The theory of fields of characteristic 0 is not complete (e.g. $\exists x(x \cdot x = 1 + 1)$ is true in $\mathbb{R}$, false in $\mathbb{Q}$). We shall see later that the theory of algebraically closed fields of characteristic 0 is complete.

2.

Two structures $\mathcal{A}$ and $\mathcal{B}$ are *elementarily equivalent* (written $\mathcal{A} \equiv \mathcal{B}$) if they have the same signature and satisfy the same sentences. In other words, $\mathrm{Th}(\mathcal{A}) = \mathrm{Th}(\mathcal{B})$.

3.

Proposition. For a theory T the following are equivalent.

1. T is complete;

2. all models of T are elementarily equivalent;

3. $T = \mathrm{Th}(\mathcal{A})$ for some structure $\mathcal{A}$.

Proof. Trivial.

4.

Exercise. Prove that the theory of dense linear ordreings without end points is complete. Use the method of elimination of quantifiers, described below.

$$\begin{array}{l}
\text{axioms for linear orderings:} \\
\qquad \forall x \forall y \forall z (x < y \wedge y < z \rightarrow x < z) \\
\qquad \forall x \forall y (x < y \rightarrow \neg y < x) \\
\qquad \forall x \forall y (x < y \vee x = y \vee y < x) \\
\text{dense :} \quad \forall x \forall y (x < y \rightarrow \exists z (x < z < y)) \\
\text{without end points :} \quad \forall x \exists y \exists z (y < x < z)\,.
\end{array}$$

5.

We use the following notational convention: $\varphi(x_1, \ldots, x_n)$ denotes a formula φ whose free variables are among $x_1, \ldots, x_n$.

A theory T is said to *admit elimination of quantifiers* if for all formulas $\varphi(x_1, \ldots, x_n)$, $n \geq 1$, there exists a quantifier free formula $\varphi^*(x_1, \ldots, x_n)$ such that

$$T \models \forall x_1 \ldots \forall x_n (\varphi(x_1, \ldots, x_n) \leftrightarrow \varphi^*(x_1, \ldots, x_n))\,.$$

For the above exercise, show that the theory of dense linear orderings without end points admits elimination of quantifiers. (Do this by induction on the number of quantifiers, working from the inside out.) Once this has been done, completeness follows easily.

Historically, the method of elimination of quantifiers came very early. It is also the most versatile method for showing that algebraic theories are complete and/or decidable. The drawback is that the method is syntactical. We shall develop various model-theoretic methods for proving the same results.

2.2 Vaught's test

Vaught's test is another method for establishing completeness of theories. It is less versatile than quantifier elimination, but much easier to use.

1.

Two structures $\mathcal{A}$ and $\mathcal{B}$ are said to be *isomorphic* (written $\mathcal{A} \cong \mathcal{B}$) if $\mathrm{sig}(\mathcal{A}) = \mathrm{sig}(\mathcal{B})$ and there exists an isomorphic map of $\mathcal{A}$ onto $\mathcal{B}$, i.e. $i : |\mathcal{A}| \to |\mathcal{B}|$ such that

1. i is one-one and onto;
2. $R^{\mathcal{A}}(a_1, \ldots, a_n)$ if and only if $R^{\mathcal{B}}(a_1, \ldots, a_n)$;
3. $i(o^{\mathcal{A}}(a_1, ..., a_n)) = o^{\mathcal{B}}(i(a_1), \ldots, i(a_n))$;
4. $i(c^{\mathcal{A}}) = c^{\mathcal{B}}$.

Note that isomorphic structures are "essentially identical". In particular they satisfy the same sentences, i.e. $\mathcal{A} \cong \mathcal{B}$ implies $\mathcal{A} \equiv \mathcal{B}$.

2.

The *cardinality* or *power* of a structure $\mathcal{A}$ is the cardinality of its universe. For instance, we say that $\mathcal{A}$ is *countable* if and only if $|\mathcal{A}|$ is countable, etc. The power of $\mathcal{A}$ is denoted $\|\mathcal{A}\|$ or $\mathrm{card}(\mathcal{A})$.

3.

A theory T is said to be κ-*categorical* if

(i) T has at least one model of power κ;

(ii) any two such models are isomorphic.

4.

Theorem (Vaught's test for completeness). Let T be a countable theory such that

(i) T has no finite models;

(ii) for some infinite cardinal κ, T is κ-categorical.

Then T is complete.

We shall prove this later as an application of the Löwenheim-Skolem-Tarski theorem. For now we content ourselves with giving various examples of how the theorem is used to prove completeness of specific theories.

2.3 Applications of Vaught's test

1.

We can use Vaught's test to show that the theory of dense linear orderings without end points is complete.

Clearly the theory has no finite models. We shall show that the theory is $\aleph_0$-categorical. Let $\mathcal{A} = (|\mathcal{A}|, <^{\mathcal{A}})$ and and $\mathcal{B} = (|\mathcal{B}|, <^{\mathcal{B}})$ be countable dense linear orderings without end points. To show $\mathcal{A} \cong \mathcal{B}$ we use a back-and-forth argument. Let $|\mathcal{A}| = \{a_n : n \in \omega\}$, $|\mathcal{B}| = \{b_n : n \in \omega\}$. The construction of an isomorphism proceeds in stages. At each stage we have constructed a finite partial isomorphism of $\mathcal{A}$ onto $\mathcal{B}$. *Stage* $2n$. Pick the least k so that $a_k = a'$ is not in the domain of the partial isomorphism so far. Find an element $b' \in |\mathcal{B}|$ satisfying finitely many inequalities so that we can extend the partial isomorphism by adding to it the ordered pair (a', b'). *Stage* $2n+1$. Proceed as in stage $2n$ replacing a's by b's and vice versa. In this way we build up an isomorphism of $\mathcal{A}$ onto $\mathcal{B}$.

2.

We use Vaught's test to show that the theory of nontrivial torsion free divisible Abelian groups is complete.

Abelian groups:

$\forall x \forall y \forall z ((x+y)+z = x+(y+z))$
$\forall x \forall y (x+y = y+x)$
$\forall x (x+0 = x)$
$\forall x (x+(-x) = 0)$

An Abelian group is said to be *divisible* if it satisfies

$$\forall x \exists y (ny = x)$$

for $n = 1, 2, 3, \ldots$. This is an infinite set of axioms. Note that ny is an abbreviation for

$$\underbrace{y + \cdots + y}_{n \text{ times}} .$$

An Abelian group is said to be *torsion free* if it satisfies

$$\forall x (nx = 0 \rightarrow x = 0)$$

for $n = 1, 2, \ldots$. This is again an infinite set of axioms.

An Abelian group is nontrivial if it contains a nonzero element. Clearly any nontrivial torsion free Abelian group is infinite. Thus our theory has no finite models. To prove completeness, we shall apply Vaught's test by showing that our theory is κ-categorical for any uncountable cardinal κ.

An example of a torsion free divisible Abelian group is the additive group of the rationals $(\mathbb{Q}, +, -, 0)$. More generally, if V is any vector space over the field of rational numbers, then the additive group of V is torsion free and divisible.

Fact: Any torsion free divisible Abelian group is of this form.

Proof. Let $\mathcal{A} = (|\mathcal{A}|, +, -, 0)$ be a torsion free divisible Abelian group. For $a \in |\mathcal{A}|$ and $r \in \mathbb{Q}$ we want to define ra. Let $r = m/n$, $m, n \in \mathbb{Z}$, $n > 0$. Define

$$ra = \text{ some } b \text{ with } ma = nb .$$

We need to show that ra is well defined. So suppose $ma = nb_1 = nb_2$. Then $n(b_1 - b_2) = 0$. hence $b_1 - b_2 = 0$ by torsion freeness of $\mathcal{A}$. The vector space axioms

$$\begin{aligned}(r_1+r_2)a &= r_1a+r_2a\\ r(a_1+a_2) &= ra_1+ra_2\\ 1a &= a\\ (-r)a &= -ra\end{aligned}$$

are easily verified. Thus $\mathcal{A}$ is the additive group of a vector space over $\mathbb{Q}$.

We shall need the following facts from the theory of vector spaces. (1) Any vector space V over a field F has a *basis*, i.e. a set $U \subseteq V$ such that every element of V is uniquely expressible in the form

$$r_1u_1+\cdots+r_nu_n\,, \qquad r_i \in F\,,$$

where the u_i are distinct elements of U, $1 \leq i \leq n$. (Proof: use Zorn's lemma.) (2) Any two bases of V over F have the same cardinality, the *dimension* of V. (3) If V_1 and V_2 are vector spaces over F of the same dimension, then they are isomorphic over F.

Let $\mathcal{A}$ and $\mathcal{B}$ be torsion free divisible Abelian groups of power κ where κ is uncountable. As vector spaces over $\mathbb{Q}$ they have bases $U_{\mathcal{A}}$ and $U_{\mathcal{B}}$ respectively. Put $\lambda = \mathrm{card}(U_{\mathcal{A}})$. Since every element of $|\mathcal{A}|$ is of the form $r_1u_1+\cdots+r_nu_n$, $r_i \in \mathbb{Q}$, $u_i \in U_{\mathcal{A}}$, we have $|\mathcal{A}| \leq \lambda \cdot \aleph_0$. From this we easily deduce $\mathrm{card}(U_{\mathcal{A}}) = \lambda = \kappa$. Similarly $\mathrm{card}(U_{\mathcal{B}}) = \kappa$. So by fact (3), $\mathcal{A} \cong \mathcal{B}$. This shows that our theory is κ-categorical. Hence by Vaught's test it is complete.

A corollary of the result we have just proved is that the Abelian groups $(\mathbb{Q},+,-,0)$ and $(\mathbb{R},+,-,0)$ are elementarily equivalent. Here $\mathbb{Q}$ = rationals, $\mathbb{R}$ = reals.

3.

Remark. The theory of nontrivial torsion free divisible Abelian groups is not $\aleph_0$ categorical, since there are countable vector spaces over $\mathbb{Q}$ of dimensions $1,2,3,\ldots$. The theory of dense linear orderings without end points is not κ-categorical, $\kappa > \aleph_0$.

4.

Remark. Let T_p be the theory of algebraically closed fields of characterstic p, where p is a prime or $p = 0$. We shall see later that Vaught's test can be used to show that T_p is complete. Namely, T_p is κ-categorical for all $\kappa > \aleph_0$.

The field axioms are as follows:

$$\begin{array}{ll} \text{axioms for a commutative ring:} & \\ & \forall x \forall y \forall z((x+y)+z = x+(y+z)) \\ & \forall x \forall y(x+y = y+x) \\ & \forall x(x+0 = x) \\ & \forall x(x+(-x) = 0) \\ & \forall x \forall y \forall z((x \cdot y) \cdot z = x \cdot (y \cdot z)) \\ & \forall x \forall y(x \cdot y = y \cdot x) \\ & \forall x \forall y \forall z(x \cdot (y+z) = x \cdot y + x \cdot z) \\ & 0 \neq 1 \\ \text{field axiom:} & \forall x(x \neq 0 \rightarrow \exists y(x \cdot y = 1) \end{array}$$

If we weaken the last axiom to

$$\forall x \forall y((x \neq 0 \wedge y \neq 0) \rightarrow x \cdot y \neq 0)$$

we get the axioms for a *domain* (usually called integral domain). A field (or domain) is said to be of *characteristic* p if p is the least n such that

$$n = \underbrace{1 + \cdots + 1}_{n \text{ times}} = 0\,.$$

Here p will have to be a prime number (since $m \cdot n = 0$ implies $m = 0$ or $n = 0$). If there is no such p, then we say that the field is of *characteristic* 0. For example $\mathbb{Q}$, $\mathbb{R}$, $\mathbb{C}$, $\mathbb{Z}$ are of characteristic 0.

A field $\mathcal{F} = (|\mathcal{F}|, +, -, 0, \cdot, 1)$ is said to be *algebraically closed* if it satisfies

$$\forall x_0 \forall x_1 \ldots \forall x_n(x_n \neq 0 \rightarrow \exists y(x_n y^n + \cdots + x_1 y + x_0 = 0))$$

for $n = 1, 2, 3, \ldots$. For example, the fundamental theorem of algebra asserts that the complex field $\mathcal{C} = (\mathbb{C}, +, -, 0, \cdot, 1)$ is algebraically closed.

5.

Remark. A difficult theorem of Morley says that if a theory T is countable and κ-categorical for some uncountable cardinal κ, then it is κ-categorical for all uncountable cardinals κ. This shows that the class of theories to which Vaught's test applies is rather limited. Many important complete theories are not κ-categorical for any κ. This limitation of Vaught's test will be partly overcome by means of saturated models. (See theorem 7.3.6 below.)

Chapter 3

The compactness theorem

In this chapter and in chapter 4 it will be convenient to assume that formulae have been defined in terms of $\neg$, $\vee$, $\exists$ only; then $\wedge$, $\rightarrow$, $\leftrightarrow$, $\forall$ can be introduced by definition as usual.

3.1 Proof of the compactness theorem

1.

A set S of sentences is said to be *consistent* if it has a model. The following theorem is due to Gödel 1929 in the countable case and Malcev 1936 in the uncountable case.

2.

Theorem (compactness theorem). If every finite subset of S is consistent, then S is consistent.

Proof. We give a proof by transfinite induction in the style of Henkin 1949. Assume that S is *finitely consistent*, i.e. every finite subset of S has a model. Let $\kappa = \max(\aleph_0, \mathrm{card}(S))$. Thus κ is a cardinal; we identify cardinals with initial ordinals. Let $\{\underline{c}_\gamma : \gamma < \kappa\}$ be new constant symbols, and put $W = \mathrm{sig}(S) \cup \{\underline{c}_\gamma : \gamma < \kappa\}$. Let $\{\sigma_\gamma : \gamma < \kappa\}$ be an enumeration of all sentences σ with $\mathrm{sig}(\sigma) \subseteq W$. Fix a variable x and let $\{\varphi_\gamma(x) : \gamma < \kappa\}$ be an enumeration of all formulas φ with only free variable x and $\mathrm{sig}(\varphi) \subseteq W$. By induction on γ define finitely consistent sets of sentences S_γ, $\gamma \leq \kappa$, and a function $h : \kappa \to \kappa$ as follows.

Stage 0: $S_0 = S$.

Stage $2\gamma + 1$: Let $h(\gamma)$ be the least β such that $\underline{c}_\beta$ does not occur in $S_{2\gamma}$ or in $\varphi_\gamma(x)$. Define

$$S_{2\gamma+1} = S_{2\gamma} \cup \{(\exists x \varphi_\gamma(x)) \to \varphi_\gamma(\underline{c}_{h(\gamma)})\}\,.$$

The sentence $(\exists x \varphi_\gamma(x)) \to \varphi_\gamma(\underline{c}_{h(\gamma)})$ is called a *Henkin sentence.* We claim that $S_{2\gamma+1}$ is finitely consistent (assuming $S_{2\gamma}$ was). Let $\mathcal{M}$ be a model of some finite subset of $S_{2\gamma+1}$. We want to expand $\mathcal{M}$ to a model of those sentences plus the Henkin sentence. First, interpret in $|\mathcal{M}|$ arbitrarily the nonlogical symbols in $\varphi_\gamma(x)$ which are not interpreted in $\mathcal{M}$. Call the resulting structure $\mathcal{M}'$. Let $\mathcal{M}''$ be the result of interpreting the Henkin constant $\underline{c}_{h(\gamma)}$ so as to make the Henkin sentence true.

Stage $2\gamma + 2$: Let $S_{2\gamma+2} = S_{2\gamma+1} \cup \{\sigma_\gamma\}$ if this is finitely consistent; otherwise $S_{2\gamma+2} = S_{2\gamma+1} \cup \{\neg\sigma_\gamma\}$. If the former is not fintely consistent then the latter is.

Stage $\delta \leq \kappa$, δ limit ordinal: Let $S_\delta = \bigcup_{\gamma<\delta} S_\gamma$. This is still finitely consistent.

We now build a model $\mathcal{M}$ directly from S_κ. Let T be the set of all variable free terms in S_κ. Define an equivalence relation $\approx$ on T by $t_1 \approx t_2$ if and only if $t_1 = t_2 \in S_\kappa$. Put $|\mathcal{M}| = T/{\approx}$. If $[t_1], \ldots, [t_n] \in T/{\approx}$ and $\underline{R}$ is an n-ary relation symbol, put

$$R([t_1], \ldots, [t_n]) \quad \text{if and only if} \quad \underline{R}(t_1, \ldots, t_n) \in S_\kappa\,.$$

If $\underline{o}$ is an n-ary operation symbol define

$$o([t_1], \ldots, [t_n]) = [\underline{o}(t_1, \ldots, t_n)]\,.$$

For $\underline{c}$ a constant symbol put $c = [\underline{c}]$. Our model $\mathcal{M} = (|\mathcal{M}|, \Phi)$ is defined by $|\mathcal{M}| = T/{\approx}$, $\Phi(\underline{R}) = R$, $\Phi(\underline{o}) = o$, $\Phi(\underline{c}) = c$.

We claim that, for all sentences σ with $\mathrm{sig}(\sigma) \subseteq W$, $\mathcal{M} \models \sigma$ if and only if $\sigma \in S_\kappa$. We prove this by induction on the number of propositional connectives in σ. For atomic σ this holds by definition of $\mathcal{M}$. For $\sigma = \sigma_1 \vee \sigma_2$ or $\sigma = \neg\sigma_1$ it is easy to check. Suppose $\sigma = \exists y \varphi(y)$. If $\mathcal{M} \models \sigma$ then $\mathcal{M} \models \varphi(t)$ for some $t \in T$. Hence by induction $\varphi(t) \in S_\kappa$. Therefore $\exists y \varphi(y) \in S_\kappa$ by finite consistency. Conversely, suppose $\exists y \varphi(y) \in S_\kappa$. By relabeling variables we may assume that $y = x$. Also $\varphi(x) = \varphi_\gamma(x)$ for some $\gamma < \kappa$. Hence $(\exists x \varphi(x)) \to \varphi(\underline{c}_\beta) \in S_\kappa$ where $\beta = h(\gamma)$. Hence $\varphi(\underline{c}_\beta) \in S_\kappa$ by finite consistency. Hence by induction $\mathcal{M} \models \varphi(\underline{c}_\beta)$, hence $\mathcal{M} \models \exists y \varphi(y)$. This completes the proof.

3.2 Some applications to field theory

1.

Theorem (A. Robinson). Suppose σ is a sentence which is true in all fields of characteristic 0. Then σ is true in all fields of characteristic p, for all but finitely many primes p (the "exceptional primes").

Proof. If the conclusion fails, then $\neg\sigma$ is consistent with the field axioms plus $\underbrace{1 + \cdots + 1}_{p} = 0$, for infinitely many primes p. But, for primes p, $\underbrace{1 + \cdots + 1}_{p} = 0$ implies $\underbrace{1 + \cdots + 1}_{n} \neq 0$ for all $n < p$. Hence $\neg\sigma$ is finitely consistent with the theory of fields of characteristic 0. Hence, by the compactness theorem, $\neg\sigma$ holds in some field of characteristic 0, contradiction.

2.

Corollary. The theory of fields of characteristic 0 is not finitely axiomatizable.

3.

An *affine variety* is a set $X \subseteq \mathbb{C}^n$ such that

$$X = \{(z_1, \ldots, z_n) : f_i(z_1, \ldots, z_n) = 0,\ 1 \leq i \leq k\}$$

where f_i, $1 \leq i \leq k$ are polynomials in n variables with coefficients from $\mathbb{C}$. A mapping $G : X \to \mathbb{C}^n$ is said to be *polynomial* if there exist polynomials g_j, $1 \leq j \leq n$ in n variables with coefficients from $\mathbb{C}$, such that for all $(z_1, \ldots, z_n) \in X$,

$$G(z_1, \ldots, z_n) = (g_1(z_1, \ldots, z_n), \ldots, g_n(z_1, \ldots, z_n)) \,.$$

Theorem (Ax). Let $G : X \to X$ be a polynomial map of an affine variety in $\mathbb{C}^n$ into itself. If G is one-one, then G is onto.

Proof. More generally, we shall prove that the theorem holds for an arbitrary algebraically closed field in place of $\mathbb{C}$. Besides the compactness theorem, we shall use completeness of the theory T_p of algebraically closed fields of characteristic p where p is 0 or a prime.

For prime p let $\overline{\mathbb{F}}_p$ be the algebraic closure of the prime field $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$. Note that any finite set of elements in $\overline{\mathbb{F}}_p$ is contained in a finite subfield of $\overline{\mathbb{F}}_p$. Hence the theorem holds for $\overline{\mathbb{F}}_p$.

Observe that the theorem holds for a particular field $\mathcal{F}$ if and only if $\mathcal{F}$ satisfies a certain collection of sentences. If one of these sentences σ is false in $\mathbb{C}$, then by completeness of T_0 we have $T_0 \models \neg\sigma$. Hence $T_p \models \neg\sigma$ for all sufficiently large primes p. Hence $\overline{\mathbb{F}}_p \models \neg\sigma$. This is a contradiction.

For more applications of the compactness theorem in field theory, see Robinson's book and more recent literature. For applications to "local properties" in group theory, see Malcev's book.

3.3 The Löwenheim-Skolem-Tarski theorem

1.

Note that the proof of the compactness theorem also established the following result: If S is (finitely) consistent then S has a model of power $\leq \max(\aleph_0, \mathrm{card}(S))$.

2.

Theorem (Löwenheim-Skolem-Tarski). Let S be a set of sentences which has an infinite model. Let κ be a cardinal $\geq \max(\aleph_0, \mathrm{card}(S))$. Then S has a model of power κ.

Proof. Let $\{\underline{c}_\gamma : \gamma < \kappa\}$ be new constant symbols. Put

$$S^* = S \cup \{\underline{c}_\beta \neq \underline{c}_\gamma : \beta < \gamma < \kappa\}.$$

Since S has an infinite model, S^* is finitely consistent. Hence by the compactness theorem S^* is consistent. Applying the remark above we find that

S^* has a model of power $\leq \max(\aleph_0, \mathrm{card}(S^*)) = \kappa$. This model must have power exactly κ.

3.

Corollary (Vaught's test). Let T be a theory which

1. has no finite models
2. is κ-categorical for some $\kappa \geq \mathrm{card}(T)$.

Then T is complete.

Proof. Suppose not. Let $\mathcal{A}, \mathcal{B} \in \mathrm{Mod}(T)$ such that $\mathcal{A} \not\equiv \mathcal{B}$. By the Löwenheim-Skolem-Tarski theorem find $\mathcal{A}'$, $\mathcal{B}'$ of power κ such that $\mathcal{A} \equiv \mathcal{A}'$, $\mathcal{B} \equiv \mathcal{B}'$. Then $\mathcal{A}' \not\equiv \mathcal{B}'$. Hence $\mathcal{A}' \not\cong \mathcal{B}'$ so T is not κ-categorical.

Chapter 4

Decidability

4.1 Recursively axiomatizable theories

1.

Let V be a fixed recursive[1] signature. The formulas φ with $\mathrm{sig}(\varphi) \subseteq V$ can be Gödel numbered in the usual way. This permits notions from recursion theory to be introduced. For instance, a set F of formulas is said to be *recursive* if its set of Gödel numbers $\{\sharp(\varphi) : \varphi \in F\}$ is recursive, etc.

2.

Definition. A theory T is said to be *recursively axiomatizable* if (i) $\mathrm{sig}(T)$ is recursive, and (ii) T has a recursive set of axioms.

Remark. The notion of recursive axiomatizability is an important generalization of that of finite axiomatizability. Most theories which arise in practice have finite signature and are recursively axiomatizable, e.g. the theory of groups, the theory of fields of characteristic 0, first order Peano arithmetic, ZF set theory. Examples of theories which are not recursively axiomatizable are: the theory of finite groups; $\mathrm{Th}(\mathcal{Z})$ where $\mathcal{Z} = (\mathbb{Z}, +, -, \cdot, 0, 1, <)$.

[1] i.e. V is a recursive set, and we can recursively identify the elements of V as n-ary relation or operation symbols or constant symbols. This is the case if V is finite.

3.

A very important general result about recursively axiomatizable theories is the following, which is a reformulation of Gödel's completeness theorem.

Theorem (essentially Gödel 1929). Let T be a recursively axiomatizable theory. Then T is recursively enumerable.

This theorem would usually be proved, following Gödel, by showing that T can be generated from its axioms and a certain explicit recursive set of "logical axioms" by means of a certain explicit recursive set of "logical rules". One would then apply Church's thesis to conclude that T is recursively enumerable. Such a proof would give more or less information depending on the choice of logical axioms and rules. We shall give a "bare bones" proof which avoids these concepts entirely.

4.

Definition. A sentence τ is *logically valid* if $\mathcal{A} \models \tau$ for all structures $\mathcal{A}$ with $\mathrm{sig}(\mathcal{A}) \subseteq \mathrm{sig}(\tau)$.

5.

Lemma. Let V be a fixed recursive signature. The set of all logically valid sentences τ with $\mathrm{sig}(\tau) \subseteq V$ is recursively enumerable.

Proof. We may safely assume that V contains no operation symbols. We may also assume that the only logical symbols are $\neg, \vee, \exists$. (Thus we are dispensing with $=$.)

Let $\{\underline{c}_n : n \in \omega\}$ be a recursive set of new constant symbols and put $W = V \cup \{\underline{c}_n : n \in \omega\}$. Let $\{\varphi_n(v_{i_n}) : n \in \omega\}$ be a recursive enumeration of all formulas $\varphi(v_i)$ with exactly one free variable and with $\mathrm{sig}(\varphi) \subseteq W$. We can choose[2] our enumeration so that

$$\mathrm{sig}(\varphi_n(v_{i_n})) \subseteq V \cup \{\underline{c}_0, \ldots, \underline{c}_{2n-1}\}\,.$$

We shall work with Henkin sentences $(\exists v_{i_n} \varphi_n(v_{i_n})) \to \varphi(\underline{c}_{2n})$. Let $S = \{\sigma_n : n \in \omega\}$ be a recursive enumeration of all sentences σ with $\mathrm{sig}(\sigma) \subseteq W$.

[2]This amounts to arranging things so that, in the proof of the compactness theorem 3.1, the function h is given by $h(n) = 2n$.

Given a sentence τ with $\mathrm{sig}(\tau) \subseteq V$, we have

τ is logically valid

$\Leftrightarrow \neg\tau$ has no model

$\Leftrightarrow \neg\tau$ has no countable model (Löwenheim-Skolem theorem)

$\Leftrightarrow$ there is no structure $\mathcal{A}$ such that $|\mathcal{A}| = \{\underline{c}_n : n \in \omega\}$ and $\mathrm{sig}(\mathcal{A}) = W$ and $\mathcal{A} \models \neg\tau$ and, for all n, $\mathcal{A} \models (\exists v_{i_n} \varphi_n(v_{i_n})) \to \varphi_n(\underline{c}_{2n})$

$\Leftrightarrow$ there is no function $f : S \to \{0,1\}$ such that for all n, $f(\neg\sigma_n) = 1 - f(\sigma_n)$, and for all m and n, $f(\sigma_m \vee \sigma_m) = \max(f(\sigma_m), f(\sigma_n))$ and $f(\exists v_{i_n} \varphi_n(v_{i_n})) = f(\varphi_n(\underline{c}_{2n}) \geq f(\varphi_n(\underline{c}_m))$.

$\Leftrightarrow \forall f : S \to \{0,1\}\, \exists n R(f \restriction S_n, \tau)$ where $S_n = \{\sigma_0, \ldots, \sigma_{n-1}\}$ and R is a fixed primitive recursive relation

$\Leftrightarrow \exists N\, \forall f : S \to \{0,1\}\, (\exists n \leq N)\, R(f \restriction S_n, \tau)$ (by König's lemma)

$\Leftrightarrow \exists N\, P(N, \tau)$ where P is a fixed primitive recursive relation.

This proves the lemma.

Remark. The set of logically valid sentences is *not* recursively enumerable ("Church's theorem").

6.

Proof of the theorem. Given a recursively axiomatizable theory T, let $A = \{\tau_n : n \in \omega\}$ be a recursive set of axioms for T. Put $A_n = \{\tau_0, \ldots, \tau_{n-1}\}$. If σ is a sentence with $\mathrm{sig}(\sigma) \subseteq \mathrm{sig}(T)$, we have

$\sigma \in T$

$\Leftrightarrow A \models \sigma$

$\Leftrightarrow \exists n\, A_n \models \sigma$ (compactness theorem)

$\Leftrightarrow \exists n \left[\bigwedge_{i=0}^{n-1} \tau_i \to \sigma \text{ is logically valid} \right]$.

By the lemma, the expression in square brackets is recursively enumerable. Hence so is T.

7.

Exercise (Craig). Prove the converse of the above theorem: If a theory T with $\mathrm{sig}(T)$ recursive is recursively enumerable, then it is recursively axiomatizable.

4.2 Decidable theories

1.

Definition. A theory T is *decidable* if (i) $\mathrm{sig}(T)$ is recursive, (ii) T is recursive.

Remark. Every decidable theory is recursively axiomatizable, but the converse fails, e.g. first order Peano arithmetic. there even exist finitely axiomatizable undecidable theories, e.g. Robinson's Q, the theory of groups, the theory of fields.

2.

Remark. The first order theories which are important in algebra tend to be decidable. We make a short table.

Decidable	Undecidable
Abelian groups	(finite) groups
Boolean algebras	(finite) distributive lattices
linear orderings	
*algebraically closed fields	fields of characteristic p
$^*\mathrm{Th}(\mathbb{R})$ = real closed fields	fields of characteristic 0
finite fields	ordered fields
*differentially closed fields of char. 0	
ordered Abelian groups	
$\mathrm{Th}(\mathbb{Q}_p)$, $\mathbb{Q}_p = p$-adic rationals	

Note: A * indicates that these decidability results will be proved later.

A useful sufficient condition for decidability is the following:

3.

Theorem. Let T be a theory which is recursively aximatizable and complete. Then T is decidable.

Proof. By the previous theorem, T is recursively enumerable. Also its complement

$$\begin{aligned}\overline{T} &= \{\sigma : \sigma \text{ is a sentence, } \mathrm{sig}(\sigma) \subseteq \mathrm{sig}(T),\, \sigma \notin T\} \\ &= \{\sigma : \neg\sigma \in T\} \qquad \text{(by completeness)}\end{aligned}$$

is recursively enumerable. Hence T is recursive.

4.

Examples. We have seen in exercise 2.1.4 that the theory of dense linear orderings without end points is complete. It is also recursively (in fact finitely) axiomatizable. Hence by the previous theorem, it is complete.

Remark. Each of the following theories will be proved complete later: algebraically closed fields of characteristic 0 or a prime p, real closed fields, differentially closed fields of characteristic 0. Decidability then follows by the previous theorem.

5.

Remark. The most versatile method for proving that an algebraic theory T is decidable is quantifier elimination. In order to decide whether $\sigma \in T$ one constructs an equivalent quantifier-free sentence σ^*. It should be easy to decide whether $\sigma^* \in T$.

4.3 Decidable models

1.

Definition. A countable structure $\mathcal{A}$ is *decidable* if (i) $\mathrm{sig}(\mathcal{A})$ is finite; (ii) there exists an enumeration of the universe of $\mathcal{A}$, $|\mathcal{A}| = \{a_k : k \in \omega\}$, such that

$$\{(\sharp(\varphi(v_1, \ldots, v_n)), \langle k_1, \ldots, k_n\rangle) : \mathcal{A} \models \varphi(a_{k_1}, \ldots, a_{k_n})\}$$

is recursive.

2.

Definition. A countable structure $\mathcal{A}$ is *computable* if (i) $\mathrm{sig}(\mathcal{A})$ is finite, and (ii) as above but restricted to atomic formulas φ.

For example, the structure $\mathcal{Z} = (\mathbb{Z}, +, -, \cdot, 0, 1, <)$ is computable but not decidable. The structure $(\mathbb{Q}, <)$ is computable and hence decidable in view of the following proposition.

3.

Proposition. Let T be a recursively axiomatizable theory which admits elimination of quantifiers. If $\mathcal{M} \in \mathrm{Mod}(T)$ is computable, then it is decidable.

Proof. To decide whether $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ find a quantifier free φ^* such that $T \models \forall x_1 \ldots \forall x_n(\varphi \leftrightarrow \varphi^*)$. Since T is recursively enumerable, we can find φ^* recursively. So $\mathcal{M} \models \varphi(a_1, \ldots, a_n)$ if and only if $\mathcal{M} \models \varphi^*(a_1, \ldots, a_n)$. Since φ^* is a Boolean combination of atomic formulas, we can recursively decide whether $\mathcal{M} \models \varphi^*(a_1, \ldots, a_n)$.

Example. It is fairly easy to see that the field $\overline{\mathbb{Q}}$ of algebraic numbers is recursive. However, it is also a model of the theory of algebraically closed fields of characteristic 0. We shall see later that this theory admits elimination of quantifiers. Hence $\overline{\mathbb{Q}}$ is decidable. Similarly, all countable algebraically closed fields are computable and hence decidable.

4.

Theorem. Let T be a decidable theory such that $\mathrm{sig}(T)$ is finite. Then T has a decidable model.

Proof. We imitate the proof of the compactness theorem 3.1.2. Put $V = \mathrm{sig}(T)$ and let $\{\underline{c}_m : m \in \omega\}$ be a recursive list of new constant symbols. Put $W = V \cup \{\underline{c}_m : m \in \omega\}$. Fix a variable x and let $\{\varphi_n(x) : n \in \omega\}$ be a recursive enumeration of all formulas $\varphi(x)$ with only one free variable x

and with $\mathrm{sig}(\varphi) \subseteq W$. Let $\{\sigma_n : n \in \omega\}$ be a recursive enumeration of all sentences σ with $\mathrm{sig}(\sigma) \subseteq W$. Perform the following recursive construction.

Stage 0: Let $S_0 = T$.

Stage $2n+1$: Let $h(n)$ be the least m such that $\underline{c}_m$ does not occur in $S_{2n} \cup \{\varphi_n(x)\}$. Put $S_{2n+1} = S_{2n} \cup \{(\exists x \varphi_n(x)) \to \varphi_n(\underline{c}_{h(n)})\}$.

State $2n+2$: Put $S_{2n+2} = S_{2n+1} \cup \{\sigma_n\}$ if this is consistent; $S_{2n+2} = S_{2n+1} \cup \{\neg\sigma_n\}$ otherwise. To see that we can make this decision recursively, note that $S_{2n+1} = T \cup \{\tau_0, \ldots, \tau_{2n}\}$ for some finitely many sentences $\tau_0, \ldots, \tau_{2n}$. Let $\widetilde{\tau}_0, \ldots, \widetilde{\tau}_{2n}, \widetilde{\sigma}_n$ be the result of replacing the new constant symbols $\underline{c}_m$ by new variables z_m. Then

$$S_{2n+1} \cup \{\sigma_n\} \text{ is consistent}$$

$$\Leftrightarrow \ T \cup \exists z_0 \ldots z_j \left(\bigwedge_{i=0}^{2n} \widetilde{\tau}_i \wedge \widetilde{\sigma}_n\right) \text{ is consistent}$$

$$\Leftrightarrow \ \neg\exists z_0 \ldots z_j \left(\bigwedge_{i=0}^{2n} \widetilde{\tau}_i \wedge \widetilde{\sigma}_n\right) \notin T,$$

and we can decide this recursively.

At the end of the construction, $S_\omega = \bigcup_{n\in\omega} S_n$ is a complete recursive theory. As in the proof of 3.1.2 we can build a model $\mathcal{M} = (|\mathcal{M}|, \Phi)$ where $|\mathcal{M}| = T/{\approx}$, $T = \{$Gödel numbers of variable free terms$\}$, $t_1 \approx t_2$ if and only if $t_1 \approx t_2 \in S_\omega$, etc. $\mathcal{M}$ is decidable because we can identify $|\mathcal{M}|$ with the set of least elements of equivalence classes under $\approx$, and then $\mathcal{M} \models \varphi(t_1, \ldots, t_n) \Leftrightarrow \varphi(t_1, \ldots, t_n) \in S_\omega$.

5.

Remark. There exist examples of recursively (even finitely) axiomatizable theories with no decidable (even computable) model. For example, we may take as axioms a finite fragment of Peano arithmetic together with a false Σ^0_1 sentence.

Chapter 5

Elementary extensions

5.1 Definition and examples

1.

Definition. Let $\mathcal{A}, \mathcal{B}$ be structures with $\mathrm{sig}(\mathcal{A}) = \mathrm{sig}(\mathcal{B})$ and $|\mathcal{A}| \subseteq |\mathcal{B}|$. We say that $\mathcal{A} \subseteq_e \mathcal{B}$ if for all formulas $\varphi(x_1, \ldots, x_n)$ with only the free variables shown, and all $a_1, \ldots, a_n \in |\mathcal{A}|$, $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$ if and only if $\mathcal{B} \models \varphi(a_1, \ldots, a_n)$. We then say that $\mathcal{A}$ is an *elementary substructure* of $\mathcal{B}$, or $\mathcal{B}$ is an *elementary extension* of $\mathcal{A}$.

2.

Let $\mathcal{A}, \mathcal{B}$ be as above. We say that $\mathcal{A} \subseteq \mathcal{B}$ ($\mathcal{A}$ is a *substructure* of $\mathcal{B}$, $\mathcal{B}$ is an *extension* of $\mathcal{A}$) if the above condition holds for atomic formulas φ.

3.

Example. Algebra is full of examples of substructures. E.g. if $\mathcal{A}$ and $\mathcal{B}$ are groups, $\mathcal{A} \subseteq \mathcal{B}$ if and only if $\mathcal{A}$ is a subgroup of $\mathcal{B}$. Similarly for rings, fields, linear orderings, etc.

4.

Example. Let $\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1, <)$ and $\mathcal{R} = (\mathbb{R}, +, -, \cdot, 0, 1, <)$. Then clearly $\mathcal{Q} \subseteq \mathcal{R}$, but $\mathcal{Q} \not\subseteq_e \mathcal{R}$ since, for example, $\mathcal{R} \models \exists x(x \cdot x = 1 + 1)$ while

$\mathcal{Q}$ does not. This example actually shows that $\mathcal{Q} \not\equiv \mathcal{R}$, i.e. $\mathcal{Q}$ and $\mathcal{R}$ are not elementarily equivalent (definition 2.1.2).

5.

Proposition. If $\mathcal{A} \subseteq_e \mathcal{B}$ then $\mathcal{A} \equiv \mathcal{B}$.

Proof. Obvious.

The converse does not hold, e.g.

6.

Example. Let $\mathbb{N} = \{0, 1, 2, \dots\}$, $\mathbb{P} = \{1, 2, \dots\}$. Then $(\mathbb{P}, <) \subseteq (\mathbb{N}, <)$, and $(\mathbb{P}, <) \equiv (\mathbb{N}, <)$ since they are isomorphic. However $(\mathbb{P}, <) \models \forall x(1 = x \vee 1 < x)$ and $(\mathbb{N}, <) \models \exists x(x < 1)$ so $(\mathbb{P}, <) \not\subseteq_e (\mathbb{N}, <)$.

However, this phenomenon does not occur in theories which admit elimination of quantifiers:

7.

Proposition. If $\mathcal{A}, \mathcal{B} \in \mathrm{Mod}(T)$ where T admits elimination of quantifiers, then $\mathcal{A} \subseteq \mathcal{B}$ implies $\mathcal{A} \subseteq_e \mathcal{B}$.

Proof. Obvious.

8.

Examples. Let X be any densely ordered subset of the real numbers such that X has no first or last element. E.g. $X = (0, 1) \cup (1, 2)$ or $X = (0, 1) \cup \mathbb{Q}$. Then $(X, <) \subseteq_e (\mathbb{R}, <)$ because the theory of dense linear ordering without endpoints admits elimination of quantifiers.

9.

Example. The fact that the theory of algebraically closed fields admits elimination of quantifiers yields the following result:

Let $\overline{\mathbb{Q}}$ be the field of algebraic numbers. Let

$$f_1(x_1,\ldots,x_n) = \cdots = f_k(x_1,\ldots,x_n) = 0 \neq g(x_1,\ldots,x_n)$$

be a finite system of equations and inequations with coefficients in $\overline{\mathbb{Q}}$. If the system has a solution in some extension field of $\mathbb{Q}$, then it has one in $\overline{\mathbb{Q}}$.

Proof. We may assume that the extension field is algebraically closed. Hence it is an extension of $\overline{\mathbb{Q}}$, and by quantifier elimination the extension is elementary. This gives the result immediately if we look at the formula

$$\exists x_1 \ldots \exists x_n \left[\bigwedge_{i=1}^{k} f_i(x_1,\ldots,x_n) = 0 \wedge g(x_1,\ldots,x_n) \neq 0 \right].$$

In other words, if a variety is nonempty, then it has a point with coordinates in $\overline{\mathbb{Q}}$. This is closely related to Hilbert's Nullstellensatz. (See chapter 6.)

9.

Definition (A. Robinson). A theory T is said to be *model complete* if for all $\mathcal{A}, \mathcal{B} \in \mathrm{Mod}(T)$, $\mathcal{A} \subseteq \mathcal{B}$ implies $\mathcal{A} \subseteq_e \mathcal{B}$.

The previous proposition says that if T admints elimination of quantifiers, then T is model complete. The converse is not true, e.g. $T = \mathrm{Th}(\mathbb{Z}, +, -, 0, 1, <)$ = Presburger arithmetic. This theory is model complete but does not admint elimination of quantifiers. However, it admits elimination of quantifiers if we add defined relations $\equiv_m$, $m \geq 2$.

10.

Example. We shall show later that the theory of algebraically closed fields is model complete and admits elimination of quantifiers. These concepts have algebraic meaning: model completeness expresses the *Hilbert Nullstellensatz*, while quantifier elimination is an abstract formulation of the theory of *resultants*.

5.2 Existence of elementary extensions

1.

Theorem. Any infinite structure $\mathcal{A}$ has a proper elementary extension $\mathcal{A}^*$.

Proof. Let $\mathcal{A} = (|\mathcal{A}|, \Phi)$. Introduce new constant symbols $\underline{a}$ for each $a \in |\mathcal{A}|$ and form the *elementary diagram* of $\mathcal{A}$, i.e. the set of all sentences $\varphi(\underline{a}_1, \ldots, \underline{a}_n)$, $\text{sig}(\varphi(x_1, \ldots, x_n)) \subseteq \text{sig}(\mathcal{A})$, $a_1, \ldots, a_n \in |\mathcal{A}|$, such[1] that $\mathcal{A} \models \varphi(\underline{a}_1, \ldots, \underline{a}_n)$. Note that an elementary extension of $\mathcal{A}$ is virtually the same thing as a model of the elementary diagram of $\mathcal{A}$.

Let $\underline{c}$ be a new constant symbol and let $S =$ (elementary diagram of $\mathcal{A}) \cup \{\underline{c} \neq \underline{a} : a \in |\mathcal{A}|\}$. Since $|\mathcal{A}|$ is infinite, S is finitely consistent. [Details: Let S_0 be a finite subset of S. Let $c_0 \in |\mathcal{A}|$, $c_0 \neq a$ for all $\underline{a}$ mentioned in S_0. Then $\mathcal{A} \models S_0$ if we interpret $\underline{a}$ as a, $\underline{c}$ as c_0.] Therefore by the compactness theorem, S is consistent, i.e. has a model $\mathcal{B} = (|\mathcal{B}|, \Psi)$. Put $\mathcal{A}^* = (|\mathcal{A}^*|, \Phi^*)$ where $|\mathcal{A}^*| = |\mathcal{B}|$, $\Phi^* = \Psi \upharpoonright \text{sig}(\mathcal{A})$. Clearly $\mathcal{A} \subseteq_e \mathcal{A}^*$ since $\mathcal{B} \models$ elementary diagram of $\mathcal{A}$. Also $|\mathcal{A}| \neq |\mathcal{A}^*|$ since $c = \Psi(\underline{c}) \in |\mathcal{A}^*| \setminus |\mathcal{A}|$.

2.

Example. Let $\mathcal{R} = (\mathbb{R}, +, -, \cdot, 0, 1, <, \cdots)$ where $\cdots$ represents other structure on $\mathbb{R}$. Let $\mathcal{R}^*$ be a proper elementary extension of $\mathcal{R}$. Then $\mathcal{R}^*$ will contain *infinitesimals*, i.e. quantities δ such that $0 < \delta < r$ for all $r \in \mathbb{R}$, $r > 0$. Elements of $\mathcal{R}^*$ are sometimes called *hyperreal numbers.* They can be used to give a rigorous development of calculus based on infinitesimals. This is the beginning of a subject known as *nonstandard analysis* (A. Robinson).

3.

Exercise. Prove the upward Löwenheim-Skolem-Tarski theorem: Let $\mathcal{A}$ be an infinite structure. Let κ be a cardinal $\geq \max(\|\mathcal{A}\|, |\text{sig}(\mathcal{A})|)$. Then $\mathcal{A}$ has a proper elementary extension of power κ.

(Hint: Combine the elementary diagram method (used in the proof of the previous theorem) with the proof of the Löwenheim-Skolem-Tarski theorem in §3.3.)

[1] To be precise, instead of $\mathcal{A} \models \varphi(\underline{a}_1, \ldots, \underline{a}_n)$ we should write $\mathcal{A}_{|\mathcal{A}|} \models \varphi(\underline{a}_1, \ldots, \underline{a}_n)$ where $\mathcal{A}_{|\mathcal{A}|} = (|\mathcal{A}|, \Phi \cup \{(\underline{a}, a) : a \in |\mathcal{A}|\})$.

In addition to the compactness theorem, another useful method for constructing elementary extensions is the method of elementary chains (Tarski/Vaught):

4.

A *chain* is a collection of structures $(\mathcal{A}_\alpha)_{\alpha<\delta}$ where δ is a limit ordinal, such that $\mathcal{A}_\alpha \subseteq \mathcal{A}_\beta$ for all $\alpha < \beta < \delta$.

5.

Proposition. Given a chain as above, there is one and only one structure $\mathcal{A}_\delta = \bigcup_{\alpha<\delta} \mathcal{A}_\alpha$ with the following properties:

1. $|\mathcal{A}_\delta| = \bigcup_{\alpha<\delta} |\mathcal{A}_\alpha|$;
2. $\mathcal{A}_\alpha \subseteq \mathcal{A}_\delta$ for all $\alpha < \delta$.

Proof. Obvious. Note that $R^{\mathcal{A}_\delta} = \bigcup_{\alpha<\delta} R^{\mathcal{A}_\alpha}$, $o^{\mathcal{A}_\delta} = \bigcup_{\alpha<\delta} o^{\mathcal{A}_\alpha}$, and $c^{\mathcal{A}_\delta} = c^{\mathcal{A}_0}$.

6.

Theorem (Tarski/Vaught elementary chain principle). Given an *elementary chain*, i.e. a chain $(\mathcal{A}_\alpha)_{\alpha<\delta}$ such that $\mathcal{A}_\alpha \subseteq_e \mathcal{A}_\beta$ for all $\alpha < \beta < \delta$, put $\mathcal{A}_\delta = \bigcup_{\alpha<\delta} \mathcal{A}_\alpha$. Then $\mathcal{A}_\alpha \subseteq_e \mathcal{A}_\delta$ for all $\alpha < \delta$.

Proof. Show by induction on the length of a formula $\varphi(x_1, \ldots, x_n)$ that, for all $\alpha < \delta$ and $a_1, \ldots, a_n \in |\mathcal{A}_\alpha|$, $\mathcal{A}_\alpha \models \varphi(a_1, \ldots, a_n)$ if and only if $\mathcal{A}_\delta \models \varphi(a_1, \ldots, a_n)$. This is trivial if φ is atomic or a Boolean combination of shorter formulas. Assume that $\varphi(x_1, \ldots, x_n) = \exists y \psi(x_1, \ldots, x_n, y)$. If $\mathcal{A}_\alpha \models \exists y \psi(a_1, \ldots, a_n, y)$, let $a_{n+1} \in |\mathcal{A}_\alpha|$ be such that $\mathcal{A}_\alpha \models \psi(a_1, \ldots, a_n, a_{n+1})$. Then $\mathcal{A}_\delta \models \psi(a_1, \ldots, a_n, a_{n+1})$ by induction, so $\mathcal{A}_\delta \models \exists y \psi(a_1, \ldots, a_n, y)$. Conversely, suppose $\mathcal{A}_\delta \models \exists y \psi(a_1, \ldots, a_n, y)$. Then $\mathcal{A}_\delta \models \psi(a_1, \ldots, a_n, b)$ for some $b \in |\mathcal{A}_\delta| = \bigcup_{\alpha<\delta} |\mathcal{A}_\alpha|$. Let $\beta < \delta$ be such that $b \in |\mathcal{A}_\beta|$. Then $\mathcal{A}_\beta \models \psi(a_1, \ldots, a_n, b)$ by induction. Hence $\mathcal{A}_\beta \models \exists y \psi(a_1, \ldots, a_n, y)$. Hence $\mathcal{A}_\alpha \models \exists y \psi(a_1, \ldots, a_n, y)$ since $\mathcal{A}_\alpha \subseteq_e \mathcal{A}_\beta$.

Later we shall use this theorem plus compactness to construct saturated models.

5.3 Elementary monomorphisms

1.

Let $\mathcal{A}, \mathcal{B}$ be structures with $\mathrm{sig}(\mathcal{A}) = \mathrm{sig}(\mathcal{B})$. We say that $f : |\mathcal{A}| \to |\mathcal{B}|$ is an *elementary embedding* or *elementary monomorphism* if f is one-one and, for all formulae $\varphi(x_1, \ldots, x_n)$ and all $a_1, \ldots, a_n \in |\mathcal{A}|$, $\mathcal{A} \models \varphi(a_1, \ldots, a_n) \Leftrightarrow \mathcal{B} \models \varphi(f(a_1), \ldots, f(a_n))$.

2.

If the above equivalence holds for atomic φ, we say that f is an *embedding* or *monomorphism* of $\mathcal{A}$ into $\mathcal{B}$, or an *isomorphism* of $\mathcal{A}$ *into* (not necessarily onto) $\mathcal{B}$. For example, what is usually called a monomorphism of groups, rings, etc. is a monomorphism in this sense.

The concepts of monomorphism and elementary monomorphism are essentially just trivial variants of the concepts of extension and elementary extension.

3.

Theorem. If $\mathcal{A} \equiv \mathcal{B}$ then there exists a structure $\mathcal{C}$ such that both $\mathcal{A}$ and $\mathcal{B}$ are elementarily embeddable into $\mathcal{C}$. (The converse is obvious.)

Proof. Let $S =$ (elementary diagram of $\mathcal{A}$) $\cup$ (elementary diagram of $\mathcal{B}$). [We assume that $\{\underline{a} : a \in |\mathcal{A}|\} \cap \{\underline{b} : b \in |\mathcal{B}|\} = \emptyset$.] Clearly any model of S yields a $\mathcal{C}$ as desired. So by the compactness theorem it suffices to show that S is finitely consistent. Let S_0 be a finite subset of S. Say $S_0 = \{\varphi(\underline{a}_1, \ldots, \underline{a}_m), \psi(\underline{b}_1, \ldots, \underline{b}_n)\}$ where $\mathrm{sig}(\varphi \wedge \psi) \subseteq \mathrm{sig}(\mathcal{A}) = \mathrm{sig}(\mathcal{B})$, $a_1, \ldots, a_m \in |\mathcal{A}|$, $b_1, \ldots, b_n \in |\mathcal{B}|$. Since $\mathcal{A} \models \exists x_1 \ldots x_m \varphi(x_1, \ldots, x_m)$ and $\mathcal{B} \models \exists y_1 \ldots y_n \psi(y_1, \ldots, y_n)$ and $\mathcal{A} \equiv \mathcal{B}$, we see that $\exists x_1 \ldots x_m \varphi(x_1, \ldots, x_m) \wedge \exists y_1 \ldots y_n \psi(y_1, \ldots, y_n)$ is consistent. Hence S_0 is consistent, Q.E.D.

4.

From the above we can derive a criterion for model completeness:

Theorem (Robinson's test). A theory T is model complete $\Leftrightarrow$ for any $\mathcal{A}, \mathcal{B} \in \mathrm{Mod}(T)$ and monomorphism $\mathcal{A} \to \mathcal{B}$, we can find an elementary extension

$\mathcal{A}^*$ of $\mathcal{A}$ and a monomorphism $\mathcal{B} \to \mathcal{A}^*$ so that the diagram

$$\begin{tikzcd} \mathcal{B} \arrow[dr] & \\ \mathcal{A} \arrow[u] \arrow[r, phantom, "\subseteq_e"] & \mathcal{A}^* \end{tikzcd}$$

commutes.

Proof. $\Rightarrow$: Since T is model complete, the monomorphism $\mathcal{A} \to \mathcal{B}$ is elementary, so we may take $\mathcal{A}^* = \mathcal{B}$.

$\Leftarrow$: Given a monomorphism $\mathcal{A} \to \mathcal{B}$ where $\mathcal{A}, \mathcal{B} \in \mathrm{Mod}(T)$. Use the hypothesis repeatedly to construct a pair of elementary chains

$$\begin{tikzcd} \mathcal{B} & = & \mathcal{B}_0 \arrow[r, "e"] \arrow[dr] & \mathcal{B}_1 \arrow[r, "e"] \arrow[dr] & \cdots \arrow[r, "e"] \arrow[dr] & \mathcal{B}_n \arrow[r, "e"] \arrow[dr] & \cdots \\ \mathcal{A} \arrow[u] & = & \mathcal{A}_0 \arrow[r, "e"] & \mathcal{A}_1 \arrow[r, "e"] \arrow[u] & \cdots \arrow[r, "e"] & \mathcal{A}_n \arrow[r, "e"] \arrow[u] & \cdots \end{tikzcd}$$

where all triangles commute. Hence $\mathcal{A}_\omega = \bigcup_{n\in\omega} \mathcal{A}_n \cong \bigcup_{n\in\omega} \mathcal{B}_n = \mathcal{B}_\omega$ so we have by Tarski/Vaught

$$\begin{tikzcd} \mathcal{A}_\omega \arrow[r, phantom, "\cong"] & \mathcal{B}_\omega \\ \mathcal{A} \arrow[u, "e"] \arrow[r] & \mathcal{B} \arrow[u, "e"] \end{tikzcd}$$

whence $\mathcal{A} \to \mathcal{B}$ is elementary. So T is model complete.

Chapter 6

Algebraically closed fields

6.1 Simple field extensions

1.

Let $\mathcal{A}$ be any structure. For each $a \in |\mathcal{A}|$ introduce a new constant symbol $\underline{a}$. The *diagram* of $\mathcal{A}$ is the set of atomic sentences $\varphi(\underline{a}_1, \ldots, \underline{a}_n)$ which hold in $\mathcal{A}$. Recall that an *extension* of $\mathcal{A}$ is any structure $\mathcal{B} \supseteq \mathcal{A}$. Note that an extension of $\mathcal{A}$ is virtually the same thing as a model of the diagram of $\mathcal{A}$.

2.

Let $\mathcal{A}$ be a field. A *field extension* of $\mathcal{A}$ is an extension $\mathcal{B} \supseteq \mathcal{A}$ such that $\mathcal{B}$ is a field. A field extension of $\mathcal{A}$ is thus virtually the same thing as a model of (field axioms) $\cup$ (diagram of $\mathcal{A}$).

3.

If $\mathcal{A}, \mathcal{B}$ are fields, $\mathcal{A} \subseteq \mathcal{B}$, and $b \in |\mathcal{B}|$, let $\mathcal{A}[b]$ be the smallest substructure of $\mathcal{B}$ containing $|\mathcal{A}| \cup \{b\}$, i.e. the subring generated by $|\mathcal{A}| \cup \{b\}$. Let $\mathcal{A}(b)$ be the smallest subfield of $\mathcal{B}$ containing $|\mathcal{A}| \cup \{b\}$. A field extension of $\mathcal{A}$ is said to be *simple* if it is of the form $\mathcal{A}(b)$.

We want to survey all possible simple field extensions of $\mathcal{A}$. This survey is the main piece of algebraic information that we need for our model-theoretic analysis.

4.

Let $\mathcal{A}$ be a field and consider terms $t(x)$ such that $\mathrm{sig}(t(x)) \subseteq \mathrm{sig}(\text{diagram of } \mathcal{A}))$ and x is the only free variable in $t(x)$. Two such terms $t_1(x)$ and $t_2(x)$ are said to be *equivalent* if

$$(\text{field axioms}) \cup (\text{diagram of } \mathcal{A}) \models \forall x(t_1(x) = t_2(x)) \,.$$

The set of equivalence classes is naturally a commutative ring and is denoted $\mathcal{A}[x]$. We shall see that $\mathcal{A}[x]$ is a domain.

5.

A *polynomial* is a term $f(x)$ as above with either $f(x) \equiv 0$ or $f(x) = \underline{a}_n x^n + \cdots + \underline{a}_1 x + \underline{a}_0$, $a_i \in |\mathcal{A}|$. Clearly each equivalence class in $\mathcal{A}[x]$ contains a polynomial. We claim that each equivalence class contains only one polynomial. Thus the polynomials are a system of representatives for the equivalence classes in $\mathcal{A}[x]$.

Obviously the polynomials over $\mathcal{A}$ form a domain, since if $f(x) = \underline{a}_m x^m + \cdots + \underline{a}_1 x + \underline{a}_0$, $g(x) = \underline{b}_n x^n + \cdots + \underline{b}_1 x + \underline{b}_0$, $\underline{a}_m \neq 0 \neq \underline{b}_n$, then $f(x) \cdot g(x) = \underline{a}_m \underline{b}_n x^{m+n} + \cdots \neq 0$. We now use the following fact: any domain is embeddable in a field, its *field of quotients*. In particular, the set of polynomials over $\mathcal{A}$ is embedded in a field extension of $\mathcal{A}$. Hence distinct polynomials represent distinct equivalence classes in $\mathcal{A}[x]$, as claimed above.

6.

It follows that $\mathcal{A}[x]$ is a domain. The quotient field of $\mathcal{A}[x]$ is denoted $\mathcal{A}(x)$. This is in agreement with our earlier notation for simple extensions. $\mathcal{A}(x)$ is our first example of a simple extension of $\mathcal{A}$.

7.

The *degree* of a nonzero polynomial $f(x) \in \mathcal{A}[x]$ is $\deg(f) = n$ where $f(x) = \underline{a}_n x^n + \cdots + \underline{a}_1 x + \underline{a}_0$, $a_n \neq 0$. The degree of the zero polynomial, $\deg(0)$, is undefined.

Lemma (division algorithm). If $f(x), g(x) \in \mathcal{A}[x]$, $g \neq 0$, there exist $q(x), r(x) \in \mathcal{A}[x]$ such that $f(x) = g(x) \cdot q(x) + r(x)$ and either $r \equiv 0$ or $\deg(r) < \deg(g)$.

Proof. As usual.

In the above lemma, if $r \equiv 0$ we say that g *divides* f.

8.

Definition. For $f(x) \in \mathcal{A}[x]$ we say that $f(x)$ is *nonconstant* if $\deg(f) > 0$. We say that $f(x)$ is *irreducible* if f is nonconstant and there is no nonconstant $g(x) \in \mathcal{A}[x]$ of lower degree than $f(x)$ such that $g(x)$ divides $f(x)$.

9.

Let $f(x) \in \mathcal{A}[x]$ be nonconstant. For $g_1, g_2 \in \mathcal{A}[x]$ say $g_1 = g_2 \mod f$ if f divides $g_1 - g_2$. Let $\mathcal{A}[x]/(f(x))$ be the set of equivalence classes mod f. Clearly $\mathcal{A}[x]/(f(x))$ has the structure of a commutative ring with $\mathcal{A}$ as a subfield. Clearly each equivalence class mod f contains one and only one polynomial $g(x) \in \mathcal{A}[x]$ such that $\deg(g) < \deg(f)$ or $g \equiv 0$. Thus $\mathcal{A}[x]/(f(x))$ may be viewed as a vector space of dimension $\deg(f)$ over $\mathcal{A}$.

10.

Lemma. If $f(x)$ is irreducible then $\mathcal{A}[x]/(f(x))$ is a field.

Proof. Given $g \in \mathcal{A}[x]$, $\deg(g) < \deg(f)$, $g \neq 0 \mod f$, to find an inverse mod f. Put

$$I = \{sf + tg : s, t \in \mathcal{A}[x]\}.$$

Let $h \in I$, $h \not\equiv 0$ be of least possible degree. Apply the division algorithm to get $f(x) = g(x) \cdot q(x) + r(x)$. Note that $r \in I$, hence $\deg(r) < \deg(h)$ is impossible. Hence $r \equiv 0$, i.e. h divides f. Also $\deg(h) \leq \deg(g) < \deg(f)$. Since f is irreducible it follows that h is a constant. $h \neq 0$ since $f \neq 0$. We now have $sf + tg = h \in |\mathcal{A}|$, so $h^{-1}tg = 1 \mod f$, so $h^{-1}t$ is the desired inverse, Q.E.D.

11.

In the above lemma, let b be the equivalence class of x in $\mathcal{A}[x]/(f(x))$. Then $\mathcal{A}(b) = \mathcal{A}[x]/(f(x))$ is a simple field extension of $\mathcal{A}$. Note also that $\mathcal{A}[b] =$

$\mathcal{A}(b)$ in this case, so the dimension of $\mathcal{A}(b)$ over $\mathcal{A}$ as a vector space is finite, being just $\deg(f)$.

We now have two kinds of examples of simple field extensions of $\mathcal{A}$. The next theorem says that there are no others.

12.

Theorem. Let $\mathcal{A}(b)$ be any simple field extension of $\mathcal{A}$. Then either
1. $\mathcal{A}(b) \cong \mathcal{A}[x]/(f(x))$ for some irreducible $f \in \mathcal{A}[x]$, or
2. $\mathcal{A}(b) \cong \mathcal{A}(x)$.

In both cases the isomorphism is the identity on $\mathcal{A}$ and sends b to x.

Proof. Case 1: b is *algebraic* over $\mathcal{A}$, i.e. $f(b) = 0$ for some nonzero $f \in \mathcal{A}[x]$. Let f be of least degree such that this holds. Clearly f is irreducible.

We claim that $g(b) = 0$ if and only if f divides g. One direction is obvious. For the other direction, suppose $g(b) = 0$. The division algorithm gives $g(x) = f(x) \cdot q(x) + r(x)$. Hence $r(b) = f(b) \cdot q(b) + r(b) = g(b) = 0$. Hence we cannot have $\deg(r) < \deg(f)$. Hence $r \equiv 0$, i.e. f divides g. This proves the claim.

It follows that $g_1(b) = g_2(b)$ if and only if f divides $g_1 - g_2$. This matches the definition of $\mathcal{A}[x]/f$, so we have $\mathcal{A}[b] \cong \mathcal{A}[x]/f$. But by lemma 10 the latter is a field, hence so is the former, i.e. $\mathcal{A}[b] = \mathcal{A}(b)$.

Case 2: not case 1, i.e. b is *transcendental* over $\mathcal{A}$. Thus $g_1(b) = g_2(b)$ if and only if $g_1(x) \equiv g_2(x)$. Thus $\mathcal{A}[b] \cong \mathcal{A}[x]$ as rings. Hence the quotient fields are also isomorphic, i.e. $\mathcal{A}(b) \cong \mathcal{A}(x)$.

13.

Corollary. Let $\mathcal{A}(b_1)$ and $\mathcal{A}(b_2)$ be two simple extensions of $\mathcal{A}$. If $f(b_1) = f(b_2) = 0$ for some ireducible $f \in \mathcal{A}[x]$, then $\mathcal{A}(b_1) \cong \mathcal{A}(b_2)$.

14.

Corollary. b is algebraic over $\mathcal{A}$ if and only if $\mathcal{A}(b)$ is finite dimensional over $\mathcal{A}$.

6.2 Algebraic closure

1.

A field extension $\mathcal{B} \supseteq \mathcal{A}$ is *algebraic over* $\mathcal{A}$ if each $b \in |\mathcal{B}|$ is algebraic over $\mathcal{A}$.

Lemma. If $\mathcal{A} \subseteq_{\mathrm{alg}} \mathcal{B} \subseteq_{\mathrm{alg}} \mathcal{C}$ then $\mathcal{A} \subseteq_{\mathrm{alg}} \mathcal{C}$.

Proof. Given $c \in |\mathcal{C}|$, to show that $|\mathcal{A}(c) : \mathcal{A}|$ (the dimension of $\mathcal{A}(c)$ over $\mathcal{A}$ as a vector space) is finite. Let $g \in \mathcal{B}[x]$ be nonzero such that $g(c) = 0$. Put $\mathcal{B}_0 = \mathcal{A}(b_0)(b_1)\cdots(b_n)$ where $g(x) = b_n x^n + \cdots + b_1 x + b_0$. Hence $|\mathcal{B}_0(c) : \mathcal{A}|$ is finite. Hence $|\mathcal{A}(c) : \mathcal{A}|$ is finite.

$$\begin{tikzcd}
\mathcal{B} \arrow[r, "\text{finite}"] & \mathcal{B}(c) \\
\mathcal{A} \arrow[u] \arrow[r, "\text{finite}"'] & \mathcal{A}(c) \arrow[u]
\end{tikzcd}$$

2.

Definition. A field $\mathcal{B}$ is *algebraically closed* if every nonconstant $g(x) \in \mathcal{B}[x]$ has a root in $\mathcal{B}$, i.e. $g(b) = 0$ for some $b \in |\mathcal{B}|$. For any field $\mathcal{A}$, an *algebraic closure* of $\mathcal{A}$ is a field extension $\mathcal{B} \supseteq \mathcal{A}$ such that (i) $\mathcal{B}$ is algebraic over $\mathcal{A}$, and (ii) $\mathcal{B}$ is algebraically closed.

3.

Proposition. Every field $\mathcal{A}$ has an algebraic closure.

Proof. Note first that for any nonconstant $g(x) \in \mathcal{A}[x]$ we can adjoin a root, i.e. we can find a simple algebraic field extension $\mathcal{A}(b)$ such that $g(b) = 0$. (Just factor $g(x)$ into irreducibles and apply the construction of Lemma 6.1.10 to one of these irreducible factors.)

Now let $g_\alpha(x)$, $\alpha < \beta$ be an enumeration of all the nonconstant $g(x) \in \mathcal{A}[x]$. Put $\mathcal{A}_0 = \mathcal{A}$, $\mathcal{A}_{\alpha+1} = \mathcal{A}_\alpha(b_\alpha)$ where $g_\alpha(b_\alpha) = 0$, and $\mathcal{A}_\delta = \bigcup_{\alpha<\delta} \mathcal{A}_\alpha$ for limit ordinals $\delta \le \beta$. Finally put $\mathcal{B}_1 = \mathcal{A}_\beta$. By Lemma 6.2.1 $\mathcal{B}_1$ is algebraic over $\mathcal{A}$, and by construction every nonconstant $g(x) \in \mathcal{A}[x]$ has a root in $\mathcal{B}_1$.

Repeating this construction ω times, we get

$$\mathcal{A} = \mathcal{B}_0 \subseteq_{\text{alg}} \mathcal{B}_1 \subseteq_{\text{alg}} \cdots \subseteq_{\text{alg}} \mathcal{B}_n \subseteq_{\text{alg}} \cdots ,$$

$n \in \omega$, such that each nonconstant $g(x) \in \mathcal{B}_n[x]$ has a root in $\mathcal{B}_{n+1}$. Thus $\mathcal{B} = \bigcup_{n\in\omega} \mathcal{B}_n$ is an algebraic closure of $\mathcal{A}$.

4.

Lemma. If $\mathcal{A}$ is algebraically closed and $\mathcal{A} \subseteq_{\text{alg}} \mathcal{B}$, then $\mathcal{A} = \mathcal{B}$.

Proof. Since $\mathcal{A}$ is algebraically closed, the only irreducible polynomials in $\mathcal{A}[x]$ are *linear*, i.e. of degree 1. Hence by Theorem 6.1.12 every element of an algebraic extension of $\mathcal{A}$ actually belongs to $\mathcal{A}$.

5.

Proposition. Any two algebraic closures of a field $\mathcal{A}$ are isomorphic over $\mathcal{A}$.

Proof. Let $\mathcal{B}_1$ and $\mathcal{B}_2$ be two algebraic closures of $\mathcal{A}$.

$$\begin{tikzcd} \mathcal{B}_1 & & \mathcal{B}_2 \\ & \mathcal{A} \arrow[ul] \arrow[ur] & \end{tikzcd}$$

If $\mathcal{A}$ is algebraically closed, then $\mathcal{B}_1 = \mathcal{B}_2 = \mathcal{A}$ by the previous lemma. By Zorn's lemma or transfinite induction, it suffices to find $b_i \in |\mathcal{B}_i| \setminus |\mathcal{A}|$, $i = 1, 2$, such that $\mathcal{A}(b_1) \cong \mathcal{A}(b_2)$ over $\mathcal{A}$. Since $\mathcal{A}$ is not algebraically closed, let $f \in \mathcal{A}[x]$ be nonconstant with no root in $\mathcal{A}$. We may assume that f is irreducible. Take $b_i \in |\mathcal{B}_i|$ so that $f(b_i) = 0$, $i = 1, 2$. Then $\mathcal{A}(b_1) \cong \mathcal{A}(b_2)$ over $\mathcal{A}$ by Corollary 6.1.13. This completes the proof.

By the above proposition we are justified in writing $\overline{\mathcal{A}} =$ *the algebraic closure of* $\mathcal{A}$ (unique up to isomorphism over $\mathcal{A}$). By the proof of Proposition 6.2.3 we have $\|\overline{\mathcal{A}}\| = \max(\|\mathcal{A}\|, \aleph_0)$.

6.

Corollary. If $\mathcal{B} \supseteq \mathcal{A}$ is algebraically closed, then $\overline{\mathcal{A}}$ is isomorphic over $\mathcal{A}$ to a subfield of $\mathcal{B}$.

$$\begin{tikzcd} & & \mathcal{B} \\ \overline{\mathcal{A}} \arrow[urr, dashed] & \supseteq & \mathcal{A} \arrow[u] \end{tikzcd}$$

Proof. Let $\mathcal{C}$ be the subfield of $\mathcal{B}$ with $|\mathcal{C}| = \{b \in |\mathcal{B}| : b \text{ is algebraic over } \mathcal{A}\}$. By Lemma 6.2.4 $\mathcal{C}$ is algebraically closed. Hence $\mathcal{C}$ is an algebraic closure of $\mathcal{A}$. Hence by Proposition 6.2.5 we have that $\overline{\mathcal{A}}$ is isomorphic to $\mathcal{C}$ over $\mathcal{A}$.

6.3 Completeness and model completeness

In this section we show that the theory T_p of algebraically closed fields of characteristic p ($p = 0$ or prime) is complete and model complete.

1.

Lemma. Let $\mathcal{A}$ be a field and let $\mathcal{B}_1, \mathcal{B}_2$ be algebraically closed extensions of $\mathcal{A}$ such that $\|\mathcal{B}_1\| = \|\mathcal{B}_2\| = \kappa > \max(\aleph_0, \|\mathcal{A}\|)$. Then $\mathcal{B}_1$ and $\mathcal{B}_2$ are isomorphic over $\mathcal{A}$.

$$\begin{tikzcd} \mathcal{B}_1 & & \mathcal{B}_2 \\ & \mathcal{A} \arrow[ul] \arrow[ur] & \end{tikzcd}$$

Proof. We use a back-and-forth argument. Let $|\mathcal{B}_i| = \{b_i^\gamma : \gamma < \kappa\}$, $i = 1, 2$. At stage γ of our construction we will have a partial isomorphism $i_\gamma : \mathcal{A}_1^\gamma \cong$

$\mathcal{A}_2^\gamma$ where $\mathcal{A} \subseteq \mathcal{A}_i^\gamma \subseteq \mathcal{B}_i$, $\|\mathcal{A}_i^\gamma\| < \kappa$, $i = 1, 2$.

$$\begin{tikzcd} & \mathcal{B}_1 & & \mathcal{B}_2 \\ i_\gamma : & \mathcal{A}_1^\gamma \arrow[u] & \cong & \mathcal{A}_2^\gamma \arrow[u] \\ & & \mathcal{A} \arrow[ul] \arrow[ur] & \end{tikzcd}$$

Stage 0: $\mathcal{A}_1^0 = \mathcal{A}_2^0 = \mathcal{A}$, $i_0 =$ identity.

Stage $2\gamma + 1$: If b_1^γ is algebraic over $\mathcal{A}_1^{2\gamma}$, let $f \in \mathcal{A}_1^{2\gamma}[x]$ be irreducible such that $f(b_1^\gamma) = 0$, and take $b \in |\mathcal{B}_2|$ such that $i_{2\gamma}(f)(b) = 0$. If b_1^γ is transcendental over $\mathcal{A}_1^{2\gamma}$, take $b \in |\mathcal{B}_2| \setminus \left|\overline{\mathcal{A}}_2^{2\gamma}\right|$; this is possible since $\left\|\overline{\mathcal{A}}_2^{2\gamma}\right\| < \kappa$, and b is then transcendental over $\mathcal{A}_2^{2\gamma}$ by Lemma 6.2.4. In either case, by Theorem 6.1.12 and Corollary 6.1.13 we can define an isomorphism

$$i_{2\gamma+1} : \mathcal{A}_1^{2\gamma+1} = \mathcal{A}_1^{2\gamma}(b_1^\gamma) \cong \mathcal{A}_2^{2\gamma}(b) = \mathcal{A}_2^{2\gamma+1}$$

by $i_{2\gamma+1} \supseteq i_{2\gamma}$, $i_{2\gamma+1}(b_1^\gamma) = b$.

Stage $2\gamma + 2$: Same as stage $2\gamma + 1$ except reverse the roles of $\mathcal{B}_1$ and $\mathcal{B}_2$.

Stage limit $\delta \le \kappa$: $\mathcal{A}_i^\delta = \bigcup_{\gamma<\delta} \mathcal{A}_i^\gamma$, $i = 1, 2$, $i_\delta = \bigcup_{\gamma<\delta} i_\gamma$. Finally we have $i_\kappa : \mathcal{B}_1 = \mathcal{A}_1^\kappa \cong \mathcal{A}_2^\kappa = \mathcal{B}_2$, Q.E.D.

2.

Theorem. T_p is κ-categorical for all $\kappa > \aleph_0$.

Proof. Let $\mathcal{A}$ be the smallest field of characteristic p, i.e. $\mathcal{A} = \mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1)$ if $p = 0$, $\mathcal{A} = \mathbb{Z}/p$ if p is prime. Note that $\mathcal{A}$ is a subfield of any field of characteristic p. Let $\mathcal{B}_1, \mathcal{B}_2 \in \mathrm{Mod}(T_p)$, $\|\mathcal{B}_1\| = \|\mathcal{B}_2\| = \kappa$. Then we have that $\mathcal{B}_1 \cong \mathcal{B}_2$ over $\mathcal{A}$ by Lemma 6.3.1.

3.

Corollary. T_p is complete.

Proof. Immediate by Vaught's test 2.2.4.

4.

Corollary. T_p is decidable.

Proof. Immediate from theorem 4.2.3.

5.

Exercise. Prove that the theory of algebraically closed fields (of all characteristics) is decidable.

6.

Theorem (A. Robinson). T_p is model complete.

Proof. Given a monomorphism $\mathcal{A} \to \mathcal{B}$ where $\mathcal{A}, \mathcal{B}$ are algebraically closed, we want to show that this monomorphism is elementary. Let κ be a cardinal $> \|\mathcal{B}\|$. By Exercise 5.2 we can get elementary extensions $\mathcal{A}^*$ of $\mathcal{A}$ and $\mathcal{B}^*$ of $\mathcal{B}$ such that $\|\mathcal{A}^*\| = \|\mathcal{B}^*\| = \kappa$. Then $\mathcal{A}^* \cong \mathcal{B}^*$ over $\mathcal{A}$ by Lemma 6.3.1.

$$\begin{tikzcd} \mathcal{A}^* \arrow[r, phantom, "\cong"] & \mathcal{B}^* \\ \mathcal{A} \arrow[u, "e"] \arrow[r] & \mathcal{B} \arrow[u, "e"'] \end{tikzcd}$$

Since the diagram commutes, the monomorphism $\mathcal{A} \to \mathcal{B}$ is elementary, Q.E.D.

7.

Corollary (Hilbert). Let $\mathcal{A}$ be a field and let $\overline{\mathcal{A}}$ be its algebraic closure. Let

$$\left.\begin{array}{c} g_1(x_1, \ldots, x_n) = \cdots = g_m(x_1, \ldots, x_n) = 0 \\ f(x_1, \ldots, x_n) \neq 0 \end{array}\right\} \qquad (*)$$

be a finite system of polynomial equations and inequations with coefficients in $\mathcal{A}$. If $(*)$ has a solution in some field extension of $\mathcal{A}$, then it has a solution in $\overline{\mathcal{A}}$.

(*Algebraic reformulation:* The algebraic points on a variety are dense with respect to the Zariski topology on that variety.)

Proof. This is immediate from model completeness of the theory of algebraically closed fields. Consider the sentence

$$\sigma \quad = \quad \exists x_1 \cdots x_n \left(\bigwedge_{i=0}^{m} g_i = 0 \wedge f \neq 0 \right) .$$

If $\mathcal{A} \subseteq \mathcal{B} \models \sigma$ then $\overline{\mathcal{B}} \models \sigma$. By Corollary 6.2.6 we have $\mathcal{A} \subseteq \overline{\mathcal{A}} \subseteq \overline{\mathcal{B}}$. Hence $\overline{\mathcal{A}} \models \sigma$ by model completeness.

6.4 Hilbert's Nullstellensatz

1.

Let $\mathcal{A}$ be a field and let $\mathcal{R} = \mathcal{A}[x_1, \ldots, x_n]$ be the ring of polynomials in n variables $x_1, \ldots, x_n$ over $\mathcal{A}$. A *zero* of $f \in \mathcal{R}$ is an n-tuple $b_1, \ldots, b_n \in |\mathcal{B}|$, where $\mathcal{B}$ is a field extension of $\mathcal{A}$, such that $f(b_1, \ldots, b_n) = 0$. An *algebraic zero* of f is a zero whose coordinates $b_1, \ldots, b_n$ lie in $\overline{\mathcal{A}}$, the algebraic closure of $\mathcal{A}$.

2.

Theorem (Nullstellensatz). Suppose $f, g_1, \ldots, g_m \in \mathcal{R}$ such that every common algebraic zero of $g_1, \ldots, g_m$ is a zero of f. Then there exists a nonnegative integer k such that

$$f^k = \sum_{i=1}^{m} p_i g_i$$

where $p_1, \ldots, p_m \in \mathcal{R}$. (The converse holds trivially.)

Proof. Let I be the set of $f \in \mathcal{R}$ for which the conclusion holds. Suppose the conclusion fails, i.e. $f \notin I$. then we claim that I is an *ideal* in $\mathcal{R}$, i.e.

(i) $g, h \in I \Rightarrow g + h \in I$;

(ii) $g \in I, h \in \mathcal{R} \Rightarrow g \cdot h \in I$;

(iii) $0 \in I$ and $1 \notin I$.

To prove (i), suppose $g, h \in I$, say $g^k = \sum p_i g_i$ and $h^l = \sum q_i g_i$. Then

$$(g+h)^{k+l} = \sum_{i+j=k+l} \binom{k+l}{i} g^i h^j \,.$$

Since $i + j = k + l$, we must have either $i \geq k$ or $j \geq l$. Therefore each term on the right hand side contains g^k or h^l. Therefore each term is of the form $\sum r_i g_i$. So $(g+h)^{k+l}$ is also of this form. So $g + h \in I$. To prove (ii) just note that if $g^k = \sum p_i g_i$ then $(g \cdot h)^k = \sum h^k p_i g_i$ so $g \cdot h \in I$. Part (iii) is obvious since $f \notin I$.

Thus I is an ideal containing no power of f. By Zorn's lemma let $J \supseteq I$ be an ideal containing no power of f and maximal with this property.

We claim that J is a prime ideal[1], i.e. $g \notin J, h \notin J \Rightarrow g \cdot h \notin J$. If $g \notin J$ then the ideal generated by $J \cup \{g\}$ is properly larger than J, so some power of f belongs to it, say $f^k = sg + u$ where $s \in \mathcal{R}$, $u \in J$. Similarly, if $h \notin J$ then $f^l = th + v$ where $t \in \mathcal{R}$, $v \in J$. Hence

$$f^{k+l} = (sg+u)\cdot(th+v) = stgh + \underbrace{sgv + thu + uv}_{\in J} \,.$$

If $gh \in J$ then $f^{k+l} \in J$, contradiction. This proves the claim.

Put $\mathcal{R}_1 = \mathcal{R}/J =$ the quotient ring of $\mathcal{R}$ by J. Since J is prime, $\mathcal{R}_1$ is a domain. Let $\mathcal{B}$ be the quotient field of $\mathcal{R}_1$. We claim that $\mathcal{A}$ is canonically isomorphic to a subfield of $\mathcal{B}$. This is clear since $|\mathcal{A}| \cap J = \{0\}$ (since $|\mathcal{A}| \cap J \neq \{0\}$ implies $1 \in J$).

Let $b_1, \ldots, b_n \in |\mathcal{B}|$ correspond to $x_1, \ldots, x_n \in \mathcal{R}$, i.e. $b_i = [x_i]_J$. Then for any $g \in \mathcal{R}$, $g(b_1, \ldots, b_n) = 0$ if and only if $g \in J$.

In particular, $g_1(b_1, \ldots, b_n) = \cdots = g_m(b_1, \ldots, b_n) = 0$ and $f(b_1, \ldots, b_n) \neq 0$. Hence by model completeness (Corollary 6.3.7) we can find $a_1, \ldots, a_n \in |\overline{\mathcal{A}}|$ such that $g_1(a_1, \ldots, a_n) = \cdots = g_m(a_1, \ldots, a_n) = 0$ and $f(a_1, \ldots, a_n) \neq 0$. Q.E.D.

3.

Given the statement of Hilbert's Nullstellensatz, it is natural to ask whether a bound on the exponent k can be computed. We can use Gödel's completeness theorem 4.1.3 to show that this is the case:

[1]The Nullstellensatz actually implies that J is a maximal ideal, i.e. $\mathcal{R}/J$ is a field.

Theorem (effective bounds for Hilbert's Nullstellensatz). We can find a recursive function $K : \omega^3 \to \omega$ with the following property:

> Let $\mathcal{A}$ be a field and let $f, g_1, \ldots, g_m \in \mathcal{R} = \mathcal{A}[x_1, \ldots, x_n]$ be such that every common algebraic zero of $g_1, \ldots, g_m$ is a zero of f. If $f, g_1, \ldots, g_m$ are of degree $\leq d$, then $f^k = \sum p_i g_i$ for some $k \leq K(m, n, d)$ and some $p_i \in \mathcal{R}$ of degree $\leq K(m, n, d)$.

Proof. Let T be the theory of algebraically closed fields with a distinguished subfield. Let σ_{mndk} be a sentence in the language of T asserting that, for all polynomials $f, g_1, \ldots, g_m$ of degree $\leq d$ in n variables $x_1, \ldots, x_n$ with coefficients in the distinguished subfield, if all common zeros of $g_1, \ldots, g_m$ are zeros of f, then there exists polynomials $p_1, \ldots, p_m$ of degree $\leq k$ in variables $x_1, \ldots, x_n$ with coefficients in the distinguished subfield, such that $f^l = \sum p_i g_i$ for some $l \leq k$.

We claim that $\forall m, n, d\, \exists k\, \sigma_{mndk} \in T$. To see this, assume the contrary for some fixed m, n, d and consider a theory $T' \supseteq T$ with new constant symbols intended to denote the coefficients of $f, g_1, \ldots, g_m$. Let T' say that for all k, σ_{mndk} fails for this choice of $f, g_1, \ldots, g_m$. By the compactness theorem, T' is consistent. But any model of T' would give a counterexample to Hilbert's Nullstellensatz.

Let $R(m, n, d, k) \Leftrightarrow \sigma_{mndk} \in T$. By Gödel's completeness theorem, T is recursively enumerable. Hence so is R. Hence we can find a recursive function $K : \omega^3 \to \omega$ such that $R(m, n, d, K(m, n, d))$ for all m, n, d. This completes the proof.

4.

One could actually show that the bounding function K can be taken to be primitive recursive. This would be done by giving an explicit quantifier elimination procedure for the theory of algebraically closed fields.

Chapter 7

Saturated models

7.1 Element types

1.

Given a theory T and a nonnegative integer n, let $F_n(T)$ be the set of all formulas $\varphi(v_1, \ldots, v_n)$ with no free variables other than $v_1, \ldots, v_n$, such that $\mathrm{sig}(\varphi) \subseteq \mathrm{sig}(T)$. We say that $Y \subseteq F_n(T)$ is *consistent over* T if there exist $\mathcal{A} \in \mathrm{Mod}(T)$ and $a_1, \ldots, a_n \in |\mathcal{A}|$ such that $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$ for all $\varphi \in Y$. Note that by compactness, if each finite subset of Y is consistent over T then so is Y.

2.

A *realization* of $Y \subseteq F_n(T)$ is an n-tuple $a_1, \ldots, a_n \in |\mathcal{A}|$, $\mathcal{A} \in \mathrm{Mod}(T)$, such that $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$ for all $\varphi \in Y$. We then say that Y is *realized* in the model $\mathcal{A}$.

We say that $\psi \in F_n(T)$ is a *logical consequence of* $Y \subseteq F_n(T)$ *over* T if every realization of Y is a realization of ψ.

3.

Definition. An *n-type over* T is a set $p \subseteq F_n(T)$ which is consistent over T and closed under logical consequence over T.

4.

Definition. An n-type p over T is said to be *complete* if, for all $\varphi \in F_n(T)$, either $\varphi \in p$ or $\neg\varphi \in p$. The set of all complete n-types over T is denoted $S_n(T)$.

5.

Remark. Let us say that $\varphi, \psi \in F_n(T)$ are equivalent over T if

$$T \models \forall v_1 \cdots \forall v_n(\varphi \leftrightarrow \psi)\,.$$

The equivalence classes form a Boolean algebra $B_n(T)$ where the Boolean operations are given by

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

$$[\varphi] + [\psi] = [\varphi \vee \psi]$$

$$-[\varphi] = [\neg\varphi]$$

$$1 = [v_1 = v_1]\,,\ 0 = [v_1 \neq v_1]$$

(For $n = 0$ this Boolean algebra is sometimes known as the *Lindenbaum algebra* of T.) A complete n-type over T is essentially the same thing as an ultrafilter on $B_n(T)$. Thus $S_n(T)$ is just the Stone space of $B_n(T)$.

6.

Example. Suppose T admits elimination of quantifiers. Then an n-type over T is determined by the quantifier-free formulas in it (at least for $n \geq 1$). For example, let T be the theory of dense linear order without end points. A complete n-type over T is determined by specifying the order relations among $v_1, \ldots, v_n$. Thus $|S_n(T)| = \sum_{k=1}^{n} k^n$. For example $|S_1(T)| = 1$, $|S_2(T)| = 3$, $|S_3(T)| = 13$, etc.

We shall see later that a theory T is $\aleph_0$-categorical if and only if $\forall n(S_n(T)$ is finite). Here T is assumed to be countable and complete and to have an infinite model.

7.

Remark. Let T be a complete theory and let p be a complete n-type over T, $n \geq 1$. It is easy to see that every finite subset of p is realized in every model of T. (Consider sentences of the form $\exists v_1 \cdots \exists v_n(\varphi_1 \wedge \cdots \wedge \varphi_k)$ where $\varphi_1, \ldots, \varphi_k \in p$.) However, p itself need not be realized in a given model of T. For example, let $T = T_0 =$ the theory of algebraically closed fields of characteristic zero. Let $p \in S_1(T)$ be a 1-type saying that v_1 is transcendental, i.e. $a_n v_1^n + \cdots + a_1 v_1 + a_0 \neq 0$ for $a_i \in \mathbb{Z}$, $a_n \neq 0$. Clearly every finite subset of p is realized in every algebraically closed field of characteristic 0. But p itself is not realized in $\overline{\mathcal{Q}} =$ the algebraic closure of the rational field $\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1)$.

8.

Proposition. Let $\mathcal{A}$ be an infinite model of a complete theory T and assume $|T| \leq \|\mathcal{A}\|$.

(i) Given $p \in S_n(T)$ we can find an elementary extension of $\mathcal{A}$ of the same power as $\mathcal{A}$ in which p is realized.

(ii) $\mathcal{A}$ has an elementary extension of power $\max(\|\mathcal{A}\|, |S_n(T)|)$ in which every $p \in S_n(T)$ is realized.

Proof. (i) Let $\underline{c}_1, \ldots, \underline{c}_n, \underline{a}\,(a \in |\mathcal{A}|)$ be new constant symbols. Since each finite subset of p is realized n $\mathcal{A}$, we have that

$$\text{(elementary diagram of } \mathcal{A}) \cup \{\varphi(\underline{c}_1, \ldots, \underline{c}_n) : \varphi \in p\}$$

is finitely consistent. Hence by the compactness and Löwenheim-Skolem theorems, this has a model of cardinality $\max(\|\mathcal{A}\|, |p|) = \|\mathcal{A}\|$. This is an elementary extension of $\mathcal{A}$ and the interpretation of $\underline{c}_1, \ldots, \underline{c}_n$ realizes p.

(ii) We use an elementary chain argument. Let $S_n(T) = \{p_\alpha : \alpha < \kappa\}$ where $\kappa = |S_n(T)|$. Put $\mathcal{A}_0 = \mathcal{A}$, $\mathcal{A}_{\alpha+1} =$ an elementary extension of $\mathcal{A}_\alpha$ of the same power as $\mathcal{A}_\alpha$ in which p_α is realized; $\mathcal{A}_\delta = \bigcup_{\alpha<\delta} \mathcal{A}_\alpha$ for limit $\delta \leq \kappa$; finally $\mathcal{B} = \mathcal{A}_\kappa$. We can show by induction on $\alpha \leq \kappa$ that $\|\mathcal{A}_\alpha\| \leq \max(\|\mathcal{A}\|, |\alpha|)$. Hence $\|\mathcal{B}\| \leq \max(\|\mathcal{A}\|, \kappa)$. If need be, use the upward Löwenheim-Skolem theorem to raise the cardinality of $\mathcal{B}$ to $\max(\|\mathcal{A}\|, \kappa)$.

7.2 Saturated models

1.

Given a structure $\mathcal{B} = (|\mathcal{B}|, \Phi)$ and a set $X \subseteq |\mathcal{B}|$, introduce new constant symbols $\underline{a}\,(a \in X)$ and let $\mathcal{B}_X = (|\mathcal{B}|, \Phi_X)$ where $\Phi_X = \Phi \cup \{(\underline{a}, a) : a \in X\}$. Thus $\mathcal{B}_X$ is just like $\mathcal{B}$ but with its signature expanded to include constant symbols denoting the elements of X. $\mathrm{Th}(\mathcal{B}_X)$ is just the set of all sentences $\varphi(\underline{a}_1, \ldots, \underline{a}_n)$ with $a_1, \ldots, a_n \in X$, $\mathrm{sig}(\varphi) \subseteq \mathrm{sig}(\mathcal{B})$, such that $\mathcal{B} \models \varphi(a_1, \ldots, a_n)$.

2.

Definition. Let κ be an infinite cardinal. We say that $\mathcal{B}$ is *κ-saturated* if for all $X \subseteq |\mathcal{B}|$ of cardinality less than κ, every complete 1-type over X is realized in $\mathcal{B}$.

(To be precise we should say: every complete 1-type over $\mathrm{Th}(\mathcal{B}_X)$ is realized in $\mathcal{B}_X$.)

3.

Exercise. Show that if $\mathcal{B}$ is κ-saturated, then for all $X \subseteq |\mathcal{B}|$ of cardinality $< \kappa$ and all n-types p over X $(n \geq 1)$, p is realized in $\mathcal{B}$.

4.

Example. Let $\mathcal{B} = (|\mathcal{B}|, <)$ be a linear ordering, and let κ be an infinite cardinal. We say that $\mathcal{B}$ is *κ-dense* if for every pair of sets $X_0, X_1 \subseteq |\mathcal{B}|$ of cardinality $< \kappa$, if $X_0 < X_1$ we can find $b \in |\mathcal{B}|$ so that $X_0 < b < X_1$. For example, $(\mathbb{Q}, <)$ is $\aleph_0$-dense (in our previous terminology, dense without endpoints). Also, $(\mathbb{R}, <)$ is not $\aleph_1$-dense, as may be seen by taking $X_0 = \{0\}$, $X_1 = \{1/2, 1/4, 1/8, \ldots\}$.

Proposition. Let $\mathcal{B}$ be a dense linear ordering without end points. $\mathcal{B}$ is κ-dense if and only if $\mathcal{B}$ is κ-saturated.

Proof. This is easy using the fact that the theory of dense linear orderings without end points admits elimination of quantifiers.

5.

We note that a finite structure is κ-saturated for all infinite κ. If an infinite structure $\mathcal{B}$ is κ-saturated, then $\|\mathcal{B}\| \geq \kappa$.

6.

Definition. A structure $\mathcal{B}$ is *saturated* if it is $\|\mathcal{B}\|$-saturated.

For example, $(\mathbb{Q}, <)$ is $\aleph_0$-saturated, hence saturated.

7.

Proposition. (i) every uncountable algebraically closed field is saturated. (ii) $\overline{\mathbb{Q}}$ and $\overline{\mathbb{F}}_p$ are not saturated. (iii) There exist countable saturated algebraically closed fields of every characteristic.

Proof. Let $X \subseteq |\mathcal{B}|$ where $\mathcal{B}$ is an algebraically closed field. A complete 1-type over X is virtually the same thing as a simple extension of the subfield $\mathcal{A}$ generated by X. We may therefore apply our classification of simple field extensions (see §6.1). The details are left to the reader.

8.

Theorem (uniqueness of saturated models). Let $\mathcal{A}$ and $\mathcal{B}$ be saturated models of the same power. If $\mathcal{A}$ and $\mathcal{B}$ are elementarily equivalent, then they are isomorphic.

Proof. We use a back-and-forth argument. Let $\|\mathcal{A}\| = \|\mathcal{B}\| = \kappa$ and fix well orderings of $|\mathcal{A}|$ and $|\mathcal{B}|$ of order type κ. We shall define enumerations $|\mathcal{A}| = \{a_\gamma : \gamma < \kappa\}$ and $|\mathcal{B}| = \{b_\gamma : \gamma < \kappa\}$ in such a way that $(\mathcal{A}, a_\gamma)_{\gamma<\kappa} \equiv (\mathcal{B}, b_\gamma)_{\gamma<\kappa}$.

Stage $\gamma < \kappa$, γ even: We have inductively $(\mathcal{A}, a_\alpha)_{\alpha<\gamma} \equiv (\mathcal{B}, b_\alpha)_{\alpha<\gamma}$. Let a_γ be the least element of $|\mathcal{A}|$ (with respect to the fixed well ordering) different from a_α, $\alpha < \gamma$. Let

$$p_\gamma \in S_1(\mathrm{Th}((\mathcal{A}, a_\alpha)_{\alpha<\gamma})) = S_1(\mathrm{Th}((\mathcal{B}, b_\alpha)_{\alpha<\gamma}))$$

be the complete 1-type over $\{a_\alpha : \alpha < \gamma\}$ realized by a_γ. Since $\mathcal{B}$ is κ-saturated we can find $b_\gamma \in |\mathcal{B}|$ realizing p_γ over $\{b_\alpha : \alpha < \gamma\}$. Thus

$$(\mathcal{A}, a_\alpha)_{\alpha \leq \gamma} \equiv (\mathcal{B}, b_\alpha)_{\alpha \leq \gamma}\,.$$

Stage $\gamma < \kappa$, γ odd: Reverse the roles of $\mathcal{A}, \mathcal{B}$.

Finally $(\mathcal{A}, a_\gamma)_{\gamma<\kappa} \equiv (\mathcal{B}, b_\gamma)_{\gamma<\kappa}$ and these enumerations exhaust the universes of $|\mathcal{A}|$ and $|\mathcal{B}|$ respectively since the elements were chosen by means of fixed well orderings of order type κ.

Thus $a_\gamma \mapsto b_\gamma$ is an isomorphism of $\mathcal{A}$ onto $\mathcal{B}$.

9.

Theorem (universality of saturated models). Let $\mathcal{B}$ be a κ-saturated model and $\mathcal{A} \equiv \mathcal{B}$, $\|\mathcal{A}\| \leq \kappa$. Then there exists an elementary monomorphism of $\mathcal{A}$ into $\mathcal{B}$.

Proof. Similar to the previous proof.

10.

Example. For algebraically closed fields, theorems 8 and 9 have the following significance. Let $p = 0$ or a prime. There is exactly one countable saturated algebraically closed field. Every countable algebraically closed field of characteristic p is embeddable into this one.

7.3 Existence of saturated models

1.

We begin with an existence theorem for countable saturated models.

Theorem (Vaught). Let T be a complete countable theory. The following are equivalent.

(1) T has a countable saturated model;

(2) $|S_n(T)| \leq \aleph_0$ for all n.

Proof. (1) $\Rightarrow$ (2): Easy. Let $\mathcal{A}$ be the countable saturated model of T. Then each $p \in S_n(T)$ is realized in $\mathcal{A}$. Hence $S_n(T)$ is countable.

(2) $\Rightarrow$ (1): Note first that (2) implies

(3) for all finite $X \subseteq |\mathcal{B}|$, $\mathcal{B} \in \mathrm{Mod}(T)$, one has $|S_1(X)| \leq X_0$.

We shall prove (3) $\Rightarrow$ (1) by an elementary chain argument.

Stage 0. Let $\mathcal{A}_0$ be any countable model of T.

Stage $n+1$. We have a countable $\mathcal{A}_n \in \mathrm{Mod}(T)$. Let $\{X_n^i : i \in \omega\}$ be an enumeration of the finite subsets of $|\mathcal{A}_n|$. Put $\mathcal{A}_n^0 = \mathcal{A}_n$, $\mathcal{A}_n^{i+1} =$ a countable elementary extension of $\mathcal{A}_n^i$ in which all complete 1-types over X_n^i are realized. ($\mathcal{A}_n^{i+1}$ exists because $S_1(X_n^i)$ is countable.) Finally put $\mathcal{A}_{n+1} = \bigcup_{i\in\omega} \mathcal{A}_n^i$. This completes stage n.

Finally put $\mathcal{B} = \bigcup_{n\in\omega} \mathcal{A}_n$. We claim that $\mathcal{B}$ is $\aleph_0$-saturated. Let $X \subseteq \mathcal{B}$ be finite. Then $X \subseteq |\mathcal{A}_n|$ for some n. Hence $X = X_n^i$ for some i. Hence every 1-type over X is realized in $\mathcal{A}_n^{i+1}$. This is an elementary submodel of $\mathcal{B}$, so these types are realized in $\mathcal{B}$, Q.E.D.

2.

Example. We give an example of a countable complete theory T with no countable saturated model. Let $T = \mathrm{Th}(\mathcal{Q}_{\mathbb{Q}})$ where $\mathcal{Q} = (\mathbb{Q}, <)$. Then $|S_1(T)| = 2^{\aleph_0}$ because each Dedekind cut in $(\mathbb{Q}, <)$ gives rise to a different complete 1-type over T. (There are other complete 1-types over T which do not correspond to Dedekind cuts.)

3.

The basic theorem on the existence of uncountable saturated models is the following:

Theorem (Morley and Vaught). Let $\mathcal{A}$ be an infinite model and let κ be an infinite cardinal such that $\kappa \geq \mathrm{sig}(\mathcal{A})$. Then $\mathcal{A}$ has a κ^+-saturated elementary extension of power $\|\mathcal{A}\|^\kappa$.

Proof. We use an elementary chain argument as in the proof of the previous theorem. Put $\lambda = \|\mathcal{A}\|^\kappa$ and note that $\lambda^\kappa = \lambda$. We construct an elementary chain in which each model has power λ.

Stage 0: Let $\mathcal{A} \subseteq_e \mathcal{A}_0$ where $\|\mathcal{A}_0\| = \lambda$.

Stage $\gamma + 1$, $\gamma < \kappa^+$: We have $\mathcal{A}_\gamma$ of power λ. Let $\{X_\gamma^\alpha : \alpha < \lambda\}$ be an enumeration of all $X \subseteq |\mathcal{A}_\gamma|$ of cardinality $\leq \kappa$. Stage $\gamma+1$ has λ substages: Put $\mathcal{A}_\gamma^0 = \mathcal{A}_\gamma$; $\mathcal{A}_\gamma^{\alpha+1}$ an elementary extension of $\mathcal{A}_\gamma^\alpha$ realizing all 1-types over X_γ^α (we can do this because $|S_1(X_\gamma^\alpha)| \leq 2^\kappa \leq \lambda$); and $\mathcal{A}_\gamma^\beta = \bigcup_{\alpha<\beta} \mathcal{A}_\gamma^\alpha$ for limit $\beta \leq \lambda$. Finally put $\mathcal{A}_{\gamma+1} = \mathcal{A}_\gamma^\lambda$.

Stage $\delta \leq \kappa^+$, δ limit: Put $\mathcal{A}_\delta = \bigcup_{\gamma<\delta} \mathcal{A}_\gamma$.

Finally put $\mathcal{B} = \mathcal{A}_{\kappa^+}$. Clearly $\mathcal{A} \subseteq_e \mathcal{B}$ and $\|\mathcal{B}\| = \lambda$. We claim $\mathcal{B}$ is κ^+-saturated. Suppose $X \subseteq |\mathcal{B}|$, $|X| \leq \kappa$. Since κ^+ is a regular cardinal, $X \subseteq \mathcal{A}_\gamma$ for some $\gamma < \kappa^+$. Hence $X = X_\gamma^\alpha$ for some $\alpha < \lambda$. So every 1-type over X is realized in $\mathcal{A}_\gamma^{\alpha+1} \subseteq_e \mathcal{B}$. This completes the proof.

4.

Corollary. Let T be a complete theory of cardinality $\leq \kappa$, and suppose that T has an infinite model. Then T has a κ^+-saturated model of power 2^κ.

Proof. By Löwenheim-Skolem let $\mathcal{A}$ be a model of T of power κ. Apply the previous theorem to get a κ^+-saturated elementary extension of power $\kappa^\kappa = 2^\kappa$.

5.

Corollary. Assume the Generalized Continuum Hypothesis. Let T be a complete theory of cardinality $\leq \kappa$. Then T has a saturated model of power κ^+. This model is unique up to isomorphism.

Proof. The G.C.H. tells us that $2^\kappa = \kappa^+$. Apply the previous corollary and the uniqueness result for saturated models (§7.2).

6.

The previous corollary may be applied to give the following useful variant of Vaught's test. Unlike Vaught's test itself, the following provides a necessary and sufficient condition for completeness.

Theorem. Assume $2^{\aleph_0} = \aleph_1$. Let T be a countable complete theory with no finite models. T is complete $\Leftrightarrow$ any two saturated models of T of power $\aleph_1$ are isomorphic.

Proof. Immediate from the previous corollary.

7.

Remark. The G.C.H. is needed for the above results. For example, let T be the complete theory of dense linear ordering without end points. Then any $\aleph_1$-saturated model of T has cardinality $\geq 2^{\aleph_0}$. Thus we do not get satisfactory results unless we assume $2^{\aleph_0} = \aleph_1$.

However, the G.C.H. can be eliminated from most applications of saturated models, by observing that the conclusions are usually absolute. Thus, if we are trying to prove that a certain theory T is complete, we may assume G.C.H., apply the variant of Vaught's test referred to above, and then eliminate G.C.H. by noting that completeness is an absolute property of T.

Actually, there are several ways to avoid the need to assume the G.C.H.

(1) Assume G.C.H. and then eliminate it by absoluteness arguments.

(2) Assume the existence of inaccessible cardinals (some people might consider this assumption more reasonable than G.C.H.).

(3) Use "special" models (Morley-Vaught).

(4) Use "recursively saturated" models (Barwise-Schlipf).

We shall simply assume the G.C.H. when needed.

7.4 Preservation theorems

In this section we present a typical application of saturated models, namely to the proofs of "preservation theorems". Preservation theorems relate model theoretic properties of T to syntactic properties of axioms for T.

1.

The following syntactical classification of formulas is occasionally useful.

Definition. A formula $\psi(\overline{y})$ is Σ_k if it is of the form

$$\psi(\overline{y}) = \exists\overline{x}_1\forall\overline{x}_2\cdots\overline{x}_k\,\varphi(\overline{x}_1,\overline{x}_2,\ldots,\overline{x}_k,\overline{y})$$

where φ is quantifier free and $\overline{x}_i = x_{i1}\ldots x_{in_i}$. Thus a Σ_k formula consists of k alternating blocks of quantifiers followed by a quantifier free matrix, and the first block is existential. The class of Π_k formulas is defined similarly except that the first block is universal.

2.

Let Γ be a class of formulas (e.g. $\Gamma = \Sigma_k$ or Π_k). A Γ-*theory* is a theory T such that $T \cap \Gamma$ is a set of axioms for T. We write $\mathcal{A} \xRightarrow{\Gamma} \mathcal{B}$ to mean that $\mathrm{sig}(\mathcal{A}) = \mathrm{sig}(\mathcal{B})$ and $\mathrm{Th}(\mathcal{A}) \cap \Gamma \subseteq \mathrm{Th}(\mathcal{B}) \cap \Gamma$.

3.

Lemma (localization lemma). Let Γ be a class of formulas which is closed under disjunction. Let T be a theory such that $\mathrm{Mod}(T)$ is closed under $\xRightarrow{\Gamma}$, i.e. $\mathcal{A} \in \mathrm{Mod}(T)$, $\mathcal{A} \xRightarrow{\Gamma} \mathcal{B}$ implies $\mathcal{B} \in \mathrm{Mod}(T)$. Then T is a Γ-theory.

Proof. Let $\mathcal{B} \models T \cap \Gamma$, $\mathrm{sig}(\mathcal{B}) = \mathrm{sig}(T)$. We must show that $\mathcal{B} \models T$. Suppose not. Let $\Delta = \{\sigma \in \Gamma : \mathcal{B} \models \neg\sigma\}$. Then for each $\mathcal{A} \in \mathrm{Mod}(T)$ there exists $\sigma \in \Delta$ such that $\mathcal{A} \models \sigma$. Hence by compactness $T \models \bigvee_{i=1}^k \sigma_i$ for some $\sigma_1, \ldots, \sigma_k \in \Delta$. Hence $\bigvee_{i=1}^k \sigma_i \in T \cap \Gamma$. Hence $\mathcal{B} \models \bigvee_{i=1}^k \sigma_i$, a contradiction.

4.

If $\mathrm{sig}(\mathcal{A}) = \mathrm{sig}(\mathcal{B})$, we write $f : \mathcal{A} \xrightarrow{\Gamma} \mathcal{B}$ (f is a Γ-*morphism* of $\mathcal{A}$ into $\mathcal{B}$) to mean that $f : |\mathcal{A}| \to |\mathcal{B}|$ and f preserves satisfaction of Γ formulas, i.e. for all Γ formulas $\varphi(x_1, \ldots, x_n)$ and $a_1, \ldots, a_n \in |\mathcal{A}|$, if $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$ then $\mathcal{B} \models \varphi(f(a_1), \ldots, f(a_n))$.

Examples:

(1) If $\Gamma = \Pi_0 = \Sigma_0 = \{$quantifier free formulas$\}$, then $f : \mathcal{A} \xrightarrow{\Gamma} \mathcal{B}$ means that f is a monomorphism of $\mathcal{A}$ into $\mathcal{B}$.

(2) If $\Gamma = \{$atomic formulas$\}$, then $f : \mathcal{A} \xrightarrow{\Gamma} \mathcal{B}$ means that f is a *homomorphism* of $\mathcal{A}$ into $\mathcal{B}$. (We have not previously defined the concept of homomorphism, so this may be taken as the definition. If $\mathcal{A}$ and $\mathcal{B}$ are groups or rings, this coincides with the usual definition of homomorphism.)

(3) If $\Gamma = \{$all formulas$\}$ then $f : \mathcal{A} \xrightarrow{\Gamma} \mathcal{B}$ means of course that f is an elementary embedding of $\mathcal{A}$ into $\mathcal{B}$.

5.

We now present our first preservation theorem.

Theorem (Tarski). Suppose that T is preserved under substructures, i.e. every substructure of a model of T is a model of T. Then T is a universal (i.e. Π_1-) theory.

Proof. By the localization lemma, it suffices to prove that if $\mathcal{B} \in \mathrm{Mod}(T)$, $\mathcal{B} \xRightarrow{\Pi_1} \mathcal{A}$, then $\mathcal{A} \in \mathrm{Mod}(T)$. Let $\mathcal{B} \equiv \mathcal{B}^*$ where $\mathcal{B}^*$ is κ-saturated, $\kappa = \|\mathcal{A}\|$. We have $\mathcal{A} \xRightarrow{\Sigma_1} \mathcal{B}$, hence $\mathcal{A} \xRightarrow{\Sigma_1} \mathcal{B}^*$.

Claim: There exists a monomorphism of $\mathcal{A}$ into $\mathcal{B}^*$. To see this, let $|\mathcal{A}| = \{a_\gamma : \gamma < \kappa\}$. As in the proof of Theorem 7.2.9 we choose $b_\gamma \in |\mathcal{B}^*|$, $\gamma < \kappa$, so that the embedding $a_\gamma \mapsto b_\gamma$ is a Σ_1-morphism. *Stage* $\gamma < \kappa$*:* Let Y_γ be the set of Σ_1 formulas $\psi(\underline{a}_{\alpha_1}, \dots, \underline{a}_{\alpha_n}, v_1)$, $\alpha_1 < \cdots < \alpha_n < \gamma$, such that $\mathcal{A} \models \psi(a_1, \dots, a_n, a_\gamma)$. By inductive hypothesis we have

$$(\mathcal{A}, a_\alpha)_{\alpha<\gamma} \xRightarrow{\Sigma_1} (\mathcal{B}^*, b_\alpha)_{\alpha<\gamma} \ .$$

Hence $\mathcal{B}^* \models \exists v_1 \psi(b_{\alpha_1}, \dots, b_{\alpha_n}, v_1)$ for each $\psi(\underline{a}_{\alpha_1}, \dots, \underline{a}_{\alpha_n}, v_1) \in Y_\gamma$. Hence by κ-saturation there exists $b_\gamma \in |\mathcal{B}^*|$ such that $\mathcal{B}^* \models \psi(b_{\alpha_1}, \dots, b_{\alpha_n}, b_\gamma)$ for all $\psi(\underline{a}_{\alpha_1}, \dots, \underline{a}_{\alpha_n}, v_1) \in Y_\gamma$.

[Details: To show that Y_γ is realized in $(\mathcal{B}^*, b_\alpha)_{\alpha<\gamma}$ it suffices to show that each finite subset is, i.e. $(\mathcal{B}^*, b_\alpha)_{\alpha<\gamma} \models \exists v \bigwedge_{i=1}^{k} \psi_i(b_{\alpha_1}, \dots, b_{\alpha_n}, v)$ for all $\psi_1, \dots, \psi_k \in Y_\gamma$. But this is a Σ_1 sentence, true in $(\mathcal{A}, a_\alpha)_{\alpha<\gamma}$, hence true in $(\mathcal{B}^*, b_\alpha)_{\alpha<\gamma}$.]

Now clearly

$$(\mathcal{A}, a_\alpha)_{\alpha\le\gamma} \xRightarrow{\Sigma_1} (\mathcal{B}^*, b_\alpha)_{\alpha\le\gamma}$$

so the induction hypothesis is preserved.

Finally

$$(\mathcal{A}, a_\gamma)_{\gamma<\kappa} \xRightarrow{\Sigma_1} (\mathcal{B}^*, b_\gamma)_{\gamma<\kappa}$$

so in particular $a_\gamma \mapsto b_\gamma$ is a monomorphism of $\mathcal{A}$ into $\mathcal{B}^*$. This proves the claim.

Now $\mathcal{B}^* \in \mathrm{Mod}(T)$, hence $\mathcal{A} \in \mathrm{Mod}(T)$ since $\mathrm{Mod}(T)$ is closed under substructure. This completes the proof.

6.

The above argument actually established the following result.

Lemma (embedding lemma). Let Γ be a class of formulas which is closed under $\wedge$ and $\exists$ (e.g. $\Gamma = \Sigma_k$, $k \geq 1$). If $\mathcal{A} \xRightarrow{\Gamma} \mathcal{B}$ and $\mathcal{B}$ is κ-saturated, $\kappa \geq \|\mathcal{A}\|$, then $\exists f : \mathcal{A} \xrightarrow{\Gamma} \mathcal{B}$.

We can use this lemma for $\Gamma = \Sigma_2$ to prove the following preservation theorem.

7.

Theorem (Chang-Łos-Susko). Let T be a theory such that $\mathrm{Mod}(T)$ is closed under unions of ω-chains, i.e. $\mathcal{A}_0 \subseteq \cdots \subseteq \mathcal{A}_n \subseteq \cdots$ $(n \in \omega)$, $\mathcal{A}_n \in \mathrm{Mod}(T)$ implies $\bigcup_{n\in\omega} \mathcal{A}_n \in \mathrm{Mod}(T)$. Then T is a Π_2-theory. (The converse is obvious.)

Proof. By the localization lemma, in order to show that T is a Π_2-theory, it suffices to show that if $\mathcal{A} \in \mathrm{Mod}(T)$ and $\mathcal{A} \xRightarrow{\Pi_2} \mathcal{B}$ then $\mathcal{B} \in \mathrm{Mod}(T)$.

We have $\mathcal{B} \xRightarrow{\Sigma_2} \mathcal{A} \in \mathrm{Mod}(T)$. Let $\mathcal{A} \equiv \mathcal{A}_1$ where $\mathcal{A}_1$ is $\|\mathcal{B}\|$-saturated. Then by the embedding lemma we can find a Σ_2 embedding $f : \mathcal{B} \xrightarrow{\Sigma_2} \mathcal{A}_1$. Let $\mathcal{B} \subseteq_e \mathcal{B}_1$, $\mathcal{B}_1$ $\|\mathcal{A}_1\|$-saturated. Then

$$(\mathcal{A}_1, f(b))_{b\in|\mathcal{B}|} \xRightarrow{\Pi_2} (\mathcal{B}_1, b)_{b\in|\mathcal{B}|}$$

so by the embedding lemma again we can find an embedding $g : \mathcal{A}_1 \xrightarrow{\Sigma_1} \mathcal{B}_1$ such that the following diagram commutes:

$$\begin{tikzcd} & \mathcal{A}_1 \arrow[d, "\Sigma_1"] \\ \mathcal{B} \arrow[ur, "\Sigma_2"] \arrow[r, phantom, "\subseteq_e"] & \mathcal{B}_1 \end{tikzcd}$$

Repeating ω times we get

$$\begin{tikzcd} & & \mathcal{A}_1 \arrow[d, "\Sigma_1"] & \mathcal{A}_2 \arrow[d, "\Sigma_1"] & & & \mathcal{A}_{n+1} \arrow[d, "\Sigma_1"] & \\ \mathcal{B} \arrow[r, phantom, "="] & \mathcal{B}_0 \arrow[ur, "\Sigma_2"] \arrow[r, phantom, "\subseteq_e"] & \mathcal{B}_1 \arrow[ur, "\Sigma_2"] \arrow[r, phantom, "\subseteq_e"] & \mathcal{B}_2 \arrow[r, phantom, "\subseteq_e"] & \cdots \arrow[r, phantom, "\subseteq_e"] & \mathcal{B}_n \arrow[ur, "\Sigma_2"] \arrow[r, phantom, "\subseteq_e"] & \mathcal{B}_{n+1} \arrow[r, phantom, "\subseteq_e"] & \cdots . \end{tikzcd}$$

So finally $\mathcal{B}_\omega = \bigcup_{n\in\omega}\mathcal{B}_n = \bigcup_{n\in\omega}\mathcal{A}_n$. Since $\mathcal{A}_n \in \mathrm{Mod}(T)$ we have by assumption $\bigcup_{n\in\omega}\mathcal{A}_n \in \mathrm{Mod}(T)$. Since $\mathcal{B} \subseteq_e \mathcal{B}_\omega$, it follows that $\mathcal{B} \in \mathrm{Mod}(T)$. This completes the proof.

8.

Corollary (Robinson). If T is model complete then T is a Π_2-theory.

Proof. If T is model complete, then any chain of models of T is an elementary chain. Hence $\mathrm{Mod}(T)$ is closed under unions of chains, so we apply the theorem.

For example, the theory of algebraically closed fields is model complete, and the given axioms for this theory are Π_2.

9.

We now prove Lyndon's preservation theorem for positive sentences. A formula is said to be *positive* if it is built up from atomic formulas using only $\forall$, $\exists$, $\wedge$, $\vee$ (and not using $\neg$, $\rightarrow$, $\leftrightarrow$). A *positive theory* is a theory with a set of positive sentences as axioms.

Theorem (Lyndon). Let T be a theory which is preserved under homomorphic images, i.e. any homomorphic image of a model of T is a model of T. Then T is a positive theory. (The converse is easy.)

For example, the theory of groups is positive, and each homomorphic image of a group is a group. The theory of commutative rings is not positive because of the presence of the axiom $0 \neq 1$. If we drop this axiom, the theory becomes positive, and every homomorphic image of a commutative ring is a commutative ring.

Proof. By the localization lemma, it suffices to show that if $\mathcal{A} \in \mathrm{Mod}(T)$, $\mathcal{A} \xRightarrow{\text{pos}} \mathcal{B}$, then $\mathcal{B} \in \mathrm{Mod}(T)$. Let $\mathcal{A} \equiv \mathcal{A}^*$, $\mathcal{B} \equiv \mathcal{B}^*$ where $\|\mathcal{A}^*\| = \|\mathcal{B}^*\| = \kappa$ and $\mathcal{A}^*, \mathcal{B}^*$ are κ-saturated. (We use the Generalized Continuum Hypothesis to obtain saturated models. By absoluteness, the use of G.C.H. can be eliminated.)

We have $\mathcal{A}^* \xRightarrow{\text{pos}} \mathcal{B}^*$. We claim that $\mathcal{B}^*$ is a homomorphic image of $\mathcal{A}^*$. The proof is similar to that of the embedding lemma (7.4.6), except we need a back-and-forth argument.

[*Details:* Choose well orderings of $|\mathcal{A}^*|$ and $|\mathcal{B}^*|$ of order type κ. We construct enumerations $|\mathcal{A}^*| = \{a_\gamma : \gamma < \kappa\}$, $|\mathcal{B}^*| = \{b_\gamma : \gamma < \kappa\}$.

Stage γ, γ even: We have

$$(\mathcal{A}^*, a_\alpha)_{\alpha<\gamma} \xRightarrow{\text{pos}} (\mathcal{B}^*, b_\alpha)_{\alpha<\gamma} \ .$$

Let a_γ be the least element of $|\mathcal{A}^*|$ not among $\{a_\alpha : \alpha < \gamma\}$. Let Y_γ be the set of positive formulas $\varphi(\underline{a}_{\alpha_1}, \ldots, \underline{a}_{\alpha_n}, v)$ such that $\mathcal{A}^* \models \varphi(a_{\alpha_1}, \ldots, a_{\alpha_n}, a_\gamma)$. As in the proof of the embedding lemma, we have $\mathcal{A}^* \models \exists v \bigwedge_{i=1}^k \varphi_i(a_{\alpha_1}, \ldots, a_{\alpha_n}, v)$, hence $\mathcal{B}^* \models \exists v \bigwedge_{i=1}^k \varphi_i(b_{\alpha_1}, \ldots, b_{\alpha_n}, v)$ for each finite set of formulas $\varphi_1, \ldots, \varphi_k \in Y_\gamma$. Hence by saturation of $\mathcal{B}^*$ it follows that Y_γ is realized in $(\mathcal{B}^*, b_\alpha)_{\alpha<\gamma}$, i.e. there exists $b_\gamma \in |\mathcal{B}^*|$ such that

$$(\mathcal{A}^*, a_\alpha)_{\alpha\leq\gamma} \xRightarrow{\text{pos}} (\mathcal{B}^*, b_\alpha)_{\alpha\leq\gamma} \ .$$

Stage γ, γ odd: We have

$$(\mathcal{A}^*, a_\alpha)_{\alpha<\gamma} \xRightarrow{\text{pos}} (\mathcal{B}^*, b_\alpha)_{\alpha<\gamma} \ ,$$

hence

$$(\mathcal{B}^*, b_\alpha)_{\alpha<\gamma} \xRightarrow{\text{neg}} (\mathcal{A}^*, a_\alpha)_{\alpha<\gamma} \ ,$$

where a *negative formula* is defined to be the negation of a positive one. Let b_γ be the least element of $|\mathcal{B}^*|$ not among $\{b_\alpha : \alpha < \gamma\}$. According to the DeMorgan laws, the class of negative formulas is closed under $\exists$ and $\wedge$. Hence we can apply the same argument as before to find $a_\gamma \in |\mathcal{A}^*|$ such that

$$(\mathcal{B}^*, b_\alpha)_{\alpha\leq\gamma} \xRightarrow{\text{neg}} (\mathcal{A}^*, a_\alpha)_{\alpha\leq\gamma} \ ,$$

i.e.

$$(\mathcal{A}^*, a_\alpha)_{\alpha\leq\gamma} \xRightarrow{\text{pos}} (\mathcal{B}^*, b_\alpha)_{\alpha\leq\gamma} \ .$$

Note that $a_\gamma \neq a_\alpha$, $\alpha < \gamma$, since $v \neq \underline{a}_\alpha$ is a negative formula.

Finally we get enumerations $|\mathcal{A}^*| = \{a_\gamma : \gamma < \kappa\}$ and $|\mathcal{B}^*| = \{b_\gamma : \gamma < \kappa\}$ such that $\alpha < \gamma < \kappa \Rightarrow a_\alpha \neq a_\gamma$ and

$$(\mathcal{A}^*, a_\gamma)_{\gamma<\kappa} \overset{\text{pos}}{\Longrightarrow} (\mathcal{B}^*, b_\gamma)_{\gamma<\kappa} \ .$$

Hence in particular the mapping $a_\gamma \mapsto b_\gamma$ is a homomorphism of $\mathcal{A}^*$ onto $\mathcal{B}^*$.]

Since $\mathcal{A} \equiv \mathcal{A}^* \in \mathrm{Mod}(T)$, it follows by hypothesis that $\mathcal{B}^* \in \mathrm{Mod}(T)$. Hence $\mathcal{B} \in \mathrm{Mod}(T)$, Q.E.D.

For more on preservation theorems, see Chang and Keisler, *Model Theory*.

Chapter 8

Elimination of quantifiers

8.1 The model completion of a theory

1.

Definition (Robinson). Let T and T^* be theories with the same signature. We say that T^* is a *model completion* of T if

(i) $T \subseteq T^*$;

(ii) every model of T can be extended to a model of T^*;

(iii) for any $\mathcal{A} \in \mathrm{Mod}(T)$, $T^* \cup (\text{diagram of } \mathcal{A})$ is complete.

(Recall that the diagram of $\mathcal{A}$ is the set of all quantifier free sentences $\varphi(\underline{a}_1, \ldots, \underline{a}_n)$, $\mathrm{sig}(\varphi) \subseteq \mathrm{sig}(\mathcal{A})$, $a_1, \ldots, a_n \in |\mathcal{A}|$, $\mathcal{A} \models \varphi(\underline{a}_1, \ldots, \underline{a}_n)$.)

Condition (iii) is equivalent to

(iii) if $\mathcal{A} \in \mathrm{Mod}(T)$, $\mathcal{A} \subseteq \mathcal{B}_i \in \mathrm{Mod}(T^*)$, $i = 1, 2$, then $\mathcal{B}_1$ and $\mathcal{B}_2$ are "elementarily equivalent over $\mathcal{A}$."

$$\begin{tikzcd} \mathcal{B}_1 \arrow[r, phantom, "\equiv"] & & \mathcal{B}_2 \\ & \mathcal{A} \arrow[ul] \arrow[ur] & \end{tikzcd}$$

2.

Example. The theory of algebraically closed fields is a model completion of the theory of fields.

Proof. (i) is trivial, and (ii) holds since every field can be extended to an algebraically closed field (Proposition 6.2.3). For (iii), let $\mathcal{A}$ be a field and let $\mathcal{B}_1$ and $\mathcal{B}_2$ be two algebraically closed fields extending $\mathcal{A}$. We want to show that $\mathcal{B}_1 \equiv \mathcal{B}_2$ over $\mathcal{A}$. By the upward Löwenheim-Skolem theorem let $\mathcal{B}_i \subseteq_e \mathcal{B}_i^*$, $i = 1, 2$, where $\|\mathcal{B}_1^*\| = \|\mathcal{B}_2^*\| = \kappa$ and $\kappa > \max(\aleph_0, \|\mathcal{A}\|)$. By Lemma 6.3.1 we have that $\mathcal{B}_1^* \cong \mathcal{B}_2^*$ over $\mathcal{A}$. Hence $\mathcal{B}_1 \equiv \mathcal{B}_2$ over $\mathcal{A}$.

$$\begin{tikzcd}
\mathcal{B}_1^* \arrow[rr, phantom, "\cong"] & & \mathcal{B}_2^* \\
\mathcal{B}_1 \arrow[u, "e"] & & \mathcal{B}_2 \arrow[u, "e"] \\
& \mathcal{A} \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

3.

Remarks. The following facts about model completions are obvious.

(a) If T^* is a model completion of T then T^* is model complete.

(b) T is model complete if and only if T is its own model completion.

4.

Theorem (Robinson). Let T_0, T_1 be model completions of T. Then $T_0 = T_1$.

Proof. Given $\mathcal{A} \in \mathrm{Mod}(T_0)$, we shall show that $\mathcal{A} \in \mathrm{Mod}(T_1)$. We use an elementary chain argument. Let $\mathcal{A} = \mathcal{A}_0 \in \mathrm{Mod}(T_0)$. By (i) and (ii) of the definition, we can find $\mathcal{A}_1 \in \mathrm{Mod}(T_1)$, $\mathcal{A}_0 \subseteq \mathcal{A}_1$. More generally, given $\mathcal{A}_{2n} \in \mathrm{Mod}(T_0)$ let $\mathcal{A}_{2n} \subseteq \mathcal{A}_{2n+1} \in \mathrm{Mod}(T_1)$ and $\mathcal{A}_{2n+1} \subseteq \mathcal{A}_{2n+2} \in \mathrm{Mod}(T_0)$. Put $\mathcal{B} = \bigcup_{n\in\omega} \mathcal{A}_n$. Since T_0 is model complete, we have $\mathcal{A}_{2n} \subseteq_e \mathcal{A}_{2n+2}$ so by the elementary chain principle, $\mathcal{A} = \mathcal{A}_0 \subseteq_e \mathcal{B}$. But since T_1 is model complete, we also have $\mathcal{A}_{2n+1} \subseteq_e \mathcal{A}_{2n+3}$ so $\mathcal{B} \in \mathrm{Mod}(T_1)$. Hence $\mathcal{A} \in \mathrm{Mod}(T_1)$, Q.E.D.

Thus the model completion of T is unique if it exists.

4.

Another example. The theory of dense linear order without end points is the model completion of the theory of linear order. (This is an easy consequence of quantifier elimination.)

8.2 Substructure completeness

1.

Definition. A theory T is *substructure complete* if $T \cup$ (diagram of $\mathcal{A}$) is complete whenever $\mathcal{A}$ is a substructure of a model of T.

2.

Definition. T admits *elimination of quantifiers* if for all $n \geq 1$ and all $\varphi \in F_n(T)$, there exists a quantifier free formula $\varphi^* \in F_n(T)$ such that $T \models \forall v_1 \ldots \forall v_n[\varphi(v_1, \ldots, v_n) \leftrightarrow \varphi^*(v_1, \ldots, v_n)]$.

3.

Theorem. Let T be a theory. The following are equivalent.

(i) T is the model completion of a universal (i.e. Π_1) theory.

(ii) T is substructure complete.

(iii) T admits elimination of quantifiers.

Proof. (i) $\Rightarrow$ (ii). Let T be the model completion of a universal theory U. Let $\mathcal{A} \subseteq \mathcal{B} \in \mathrm{Mod}(T)$. Then $\mathcal{A} \in \mathrm{Mod}(U)$ since U is universal. Hence $T \cup$ (diagram of $\mathcal{A}$) is complete, by the definition of model completion.

(ii) $\Rightarrow$ (iii). Assume that T is substructure complete. Given $\varphi \in F_n(T)$, to find a quantifier free $\varphi^* \in F_n(T)$ such that $T^* \models \varphi \leftrightarrow \varphi^*$. We proceed as in the proof of the localization lemma 7.4.3. Let Y be the set of quantifier free $\psi \in F_n(T)$ such that $T \models \varphi \rightarrow \psi$. If φ^* does not exist, then by compactness we can find $\mathcal{B} \in \mathrm{Mod}(T)$ and $a_1, \ldots, a_n \in |\mathcal{B}|$ so that $\mathcal{B} \models$

$\neg\varphi(a_1, \ldots, a_n) \wedge \psi(a_1, \ldots, a_n)$ for all $\psi \in Y$. Let $\mathcal{A}$ be the substructure of $\mathcal{B}$ generated by $a_1, \ldots, a_n$. By substructure completeness we have

$$T \cup (\text{diagram of } \mathcal{A}) \models \neg\varphi(\underline{a}_1, \ldots, \underline{a}_n)\,.$$

Hence there exists a quantifier free formula $\theta \in F_n(T)$ such that $\theta(\underline{a}_1, \ldots, \underline{a}_n) \in$ diagram of $\mathcal{A}$ and

$$T \models \forall v_1 \ldots v_n\, [\theta(v_1, \ldots, v_n) \to \neg\varphi(v_1, \ldots, v_n)]\,,$$

i.e. $T \models \forall v_1 \ldots v_n[\varphi \to \neg\theta]$. Hence $\neg\theta \in Y$. Hence $\mathcal{B} \models \neg\theta(a_1, \ldots, a_n)$. This is a contradiction.

(iii) $\Rightarrow$ (i). Suppose T admits elimination of quantifiers. Let U be the theory whose axioms are the universal sentences of T. We claim that T is a the model completion of U. If $\mathcal{A} \in \mathrm{Mod}(U)$ then an easy compactness argument shows that $T \cup (\text{diagram of } \mathcal{A})$ is consistent. Hence $\mathcal{A}$ is extendible to a model of T. The last part of the definition of model completion is immediate from elimination of quantifiers.

(Note: $\mathrm{Mod}(U)$ is just the class of substructures of models of T.)

4.

Example. The theory of dense linear ordering without end points is the model completion of the theory of linear orderings. the latter theory is universal so we conclude:

Theorem. The theory of dense linear ordering without endpoints admits elimination of quantifiers.

(This can also be proved by a direct syntactical argument. See Exercise 2.1.4.)

5.

Example. We have seen that the theory T of algebraically closed fields is the model completion of the theory of fields. Unfortunately, the theory of fields is not universal. (The key axiom is $\forall x(x \neq 0 \to \exists y(x \cdot y = 1))$.) However, T is also the model completion of the theory of domains, and the latter theory is universal. (The key axiom is $\forall x \forall y((x \neq 0 \wedge y \neq 0) \to x \cdot y \neq 0)$.)

(To see that T is the model completion of the theory of domains, just note that any field containing a domain also contains its fraction field, and apply Example 8.1.2. See also Example 8.3.3 below.)

Hence we have:

Theorem (Tarski). The theory of algebraically closed fields admits elimination of quantifiers. (In the language of $+, \cdot, -, 0, 1$.)

6.

This theorm has algebraic applications; e.g.

Corollary (elimination theory). Let S be a finite system of equations and inequations in n unknowns (variables) $x_1, \ldots, x_n$ with coefficients (constants) $\underline{c}_1, \ldots, \underline{c}_N$. Then we can effectively find a finite set $S_1, \ldots, S_k$ of systems of equations and inequations in $\underline{c}_1, \ldots, \underline{c}_N$ alone, such that for any $c_1, \ldots, c_N$ in an algebraically closed field, S is solvable if and only if one of $S_1, \ldots, S_k$ holds.

Proof. Regard $\underline{c}_1, \ldots, \underline{c}_N$ as variables and effectively eliminate quantifiers from the formula $\exists x_1 \cdots \exists x_n\, S$. Put the resulting quantifier free formula into disjunctive normal form.

Special cases of the above result are known classically. For example, in van der Waerden (vol. 1, §27) one finds the following result. Let S be a system of two equations

$$\begin{aligned} a_0 x^m + a_1 x^{m-1} + \cdots + a_m &= 0 \\ b_0 x^n + b_1 x^{n-1} + \cdots + b_n &= 0\,. \end{aligned}$$

These two polynomials have a common root if and only if the determinant

$$\begin{vmatrix} a_0 & a_1 & \ldots & a_m & & \\ & a_0 & a_1 & \ldots & a_m & \\ & & a_0 & a_1 & \ldots & a_m \\ b_0 & b_1 & \ldots & b_n & & \\ & b_0 & b_1 & \ldots & b_n & \\ & & b_0 & b_1 & \ldots & b_n \end{vmatrix}$$

vanishes. This $(m+n) \times (m+n)$ determinant is called the *resultant* of the system S. In general, the theory of "resultants" or "elimination theory" is a branch of classical algebra concerned with explicit determination of $S_1, \ldots, S_k$ given S as in the corollary.

7.

Recall that a variety $X \subseteq \mathbb{C}^n$ is the solution set of a system of equations in n unknowns with complex coefficients.

Corollary (Tarski). Let $X \subseteq \mathbb{C}^n$ be a variety. Let $Y \subseteq \mathbb{C}^k$ be the image of X under the projection $(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_k)$, $1 \leq k \leq n$. Then Y is a Boolean combination of varieties.

(This is already nontrivial for $k = 2$, $n = 3$.)

8.3 The role of simple extensions

In this section we discuss the role of simple extensions in quantifier elimination.

1.

Definition. Let U be a universal theory. Let $\mathcal{A} \subseteq \mathcal{B} \in \mathrm{Mod}(U)$. Given $b \in |\mathcal{B}|$, we write $\mathcal{A}[b]$ for the substructure of $\mathcal{B}$ generated by $|\mathcal{A}| \cup \{b\}$. This is called a simple extension of $\mathcal{A}$.

2.

Theorem (L. Blum). Let U be a universal theory and let T be a theory with $\mathrm{sig}(T) = \mathrm{sig}(U)$, $U \subseteq T$, such that every model of U is extendible to a model of T. The following are equivalent.

(i) T is the model completion of U.

(ii) Suppose $\mathcal{A} \subseteq \mathcal{B} \in \mathrm{Mod}(T)$, $\mathcal{B}$ κ^+-saturated, where $\kappa = \max(\aleph_0, \|\mathcal{A}\|)$. Then up to isomorphism over $\mathcal{A}$, $\mathcal{B}$ contains every simple extension of $\mathcal{A}$.

(iii) Suppose $\mathcal{A} \subseteq \mathcal{B}_i \in \mathrm{Mod}(T)$, $i = 1, 2$. If $\mathcal{B}_2$ is κ-saturated, $\kappa \geq \|\mathcal{B}_1\|$, then there exists an embedding of $\mathcal{B}_1$ into $\mathcal{B}_2$ such that the following diagram commutes.

$$\begin{tikzcd} \mathcal{B}_1 \arrow[rr, dashed] & & \mathcal{B}_2 \\ & \mathcal{A} \arrow[ul] \arrow[ur] & \end{tikzcd}$$

Proof. (i) $\Rightarrow$ (ii). Immediate from the definitions.

(ii) $\Rightarrow$ (iii). Let $|\mathcal{B}_1| = \{b_\alpha : \alpha < \kappa\}$ and define a tower of simple extensions $\mathcal{A}_0 = \mathcal{A}$, $\mathcal{A}_{\alpha+1} = \mathcal{A}_\alpha[b_\alpha]$, $\mathcal{A}_\beta = \bigcup_{\alpha<\beta} \mathcal{A}_\alpha$ for limit $\beta \leq \kappa$. Apply (ii) κ times to get an embedding of $\mathcal{B}_1$ into $\mathcal{B}_2$.

(iii) $\Rightarrow$ (i). Given $\mathcal{A} \subseteq \mathcal{B}_i \in \mathrm{Mod}(T)$, $i = 1, 2$. To show $\mathcal{B}_1 \equiv \mathcal{B}_2$ over $\mathcal{A}$. We use an elementary chain argument.

$$\begin{tikzcd} & \mathcal{B}_1 \arrow[r, "e"] \arrow[ddr] & \mathcal{B}_1' \arrow[r, "e"] \arrow[ddr] & \mathcal{B}_1'' \arrow[r, "e"] & \cdots \\ \mathcal{A} \arrow[ur] \arrow[dr] & & & & \\ & \mathcal{B}_2 \arrow[r, "e"] \arrow[uur] & \mathcal{B}_2' \arrow[r, "e"] \arrow[uur] & \mathcal{B}_2'' \arrow[r, "e"] & \cdots \end{tikzcd}$$

First choose $\mathcal{B}_2' \supseteq_e \mathcal{B}_2$ sufficiently saturated, and apply (ii) to get $\mathcal{B}_1 \to \mathcal{B}_2'$ over $\mathcal{A}$. Then choose $\mathcal{B}_1' \supseteq_e \mathcal{B}_1$ sufficiently saturated and apply (ii) to get $\mathcal{B}_2 \to \mathcal{B}_1'$ over $\mathcal{A}$. Etc. Finally we get

$$\begin{tikzcd} & \mathcal{B}_1 \arrow[r, "e"] & \mathcal{B}_1^* & = & \bigcup_{n\in\omega} \mathcal{B}_1^{(n)} \\ \mathcal{A} \arrow[ur] \arrow[dr] & & \cong & & \\ & \mathcal{B}_2 \arrow[r, "e"] & \mathcal{B}_2^* & = & \bigcup_{n\in\omega} \mathcal{B}_2^{(n)} \end{tikzcd}$$

and the diagram gives an isomorphism of $\mathcal{B}_1^*$ onto $\mathcal{B}_2^*$ over $\mathcal{A}$. Hence $\mathcal{B}_1 \equiv \mathcal{B}_2$ over $\mathcal{A}$.

3.

Example. We can use the above criterion to prove (using a minimum of algebra) that the theory of algebraically closed fields is the model completion of the theory of domains. Suppose $\mathcal{A} \subseteq \mathcal{B}$, $\mathcal{B}$ algebraically closed and κ^+-saturated, where $\kappa \geq \max(\aleph_0, \|\mathcal{A}\|)$. We must show that $\mathcal{B}$ contains every simple extension $\mathcal{A}[b]$ of the domain $\mathcal{A}$. All such extensions are contained in simple field extensions of the quotient field of $\mathcal{A}$. Using our classification of simple field extensions (§6.1) we see immediately that $\mathcal{B}$ contains all such extensions.

In the next chapter we shall use this criterion to show that the theory of real closed ordered fields is the model completion of the theory of ordered fields.

Chapter 9

Real closed ordered fields

9.1 Ordered fields

1.

The *axioms for ordered fields* are as follows:

(a) the field axioms (as in example 2.3.4)

(b) the axioms for linear order:

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

(c) special axioms:

$$\forall x \forall y \forall z(x < y \rightarrow x + z < y + z)$$
$$\forall x \forall y((0 < x \wedge 0 < y) \rightarrow 0 < x \cdot y)$$

An *ordered field* is a model of these axioms. An *ordered domain* is defined similarly.

2.

Examples.
$\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1, <)$ = the ordered field of rationals.

$\mathcal{R} = (\mathbb{R}, +, -, \cdot, 0, 1, <) =$ the ordered field of real numbers.
$\mathcal{Z} = (\mathbb{Z}, +, -, \cdot, 0, 1, <) =$ the ordered field of rationals.

3.

The following facts are easily proved.

(a) In an ordered domain, we have

$$x < y \leftrightarrow y - x > 0$$
$$x > 0 \leftrightarrow -x < 0$$
$$x \neq 0 \rightarrow x^2 > 0$$
$$(x < y \wedge z > 0) \rightarrow x \cdot z < y \cdot z$$

(b) Any ordered field (or ordered domain) has characteristic 0. (Because $\underbrace{1 + 1 + \cdots + 1}_{n} > 0$.)

(c) The ordering of an ordered field is dense (because $x < y$ implies $x < (x + y)/2 < y$).

(d) Any ordered field contains $\mathcal{Q}$ as an ordered subfield. Any ordered domain contains $\mathcal{Z}$ as an ordered subdomain.

4.

Proposition. Let $\mathcal{A}$ be an ordered domain and let $\mathcal{A}_1$ be its field of quotients. Then there is one and only one way to order $\mathcal{A}_1$ so that it becomes an ordered field with $\mathcal{A}$ as an ordered subdomain.

Proof. We are forced to order $\mathcal{A}_1$ by

$$\frac{x}{y} > 0 \quad \Leftrightarrow \quad x \cdot y = \frac{x}{y} \cdot y^2 > 0\,.$$

It is easy to check that this makes $\mathcal{A}_1$ an ordered field, etc.

For example, there is only one way to order the rational field so as to make it an ordered field.

5.

Definition. Let $\mathcal{F}$ be an ordered field and let $f \in \mathcal{F}[x]$. We say that $\mathcal{F}$ has the *intermediate value property* (IVP) for f if $a, b \in |\mathcal{A}|$, $a < b$, $f(a) < 0 < f(b)$ imply $f(c) = 0$ for some $c \in |\mathcal{F}|$, $a < c < b$.

We say that $\mathcal{F}$ is *real closed* if it has the IVP for all $f \in \mathcal{F}[x]$.

6.

Examples. The rational field $\mathcal{Q}$ does not have the IVP for $f(x) = x^2 - 2$. The real field $\mathcal{R}$ does have IVP for $x^2 - 2$. The real field $\mathcal{R}$ is in fact real closed. The complex field $\mathbb{C}$ is *not* real closed, indeed it cannot be ordered (since $i^2 = -1$).

7.

The property of being real closed is first order, so we may speak of *the theory of real closed ordered fields*, RCOF. The axioms of RCOF are those for ordered fields plus the following, for each $n \geq 1$:

$$\begin{aligned} &\forall w_0 \cdots w_n xy \left[(x < y \wedge w_n x^n + \cdots + w_0 < 0 < w_n y^n + \cdots + w_0)\right. \\ &\qquad \left.\rightarrow \exists z\,(x < z < y \wedge w_n z^n + \cdots + w_0 = 0)\right]. \end{aligned}$$

8.

Our main goal is to show that RCOF is the model completion of the theory of ordered fields, and hence of ordered domains. From this it will follow that RCOF admits elimination of quantifiers, and is complete and decidable (Tarski).

9.

Lemma. Let $\mathcal{F}$ be an ordered field which is not real closed. Then there is an ordered field $\mathcal{F}_1 \supsetneqq \mathcal{F}$ such that $\mathcal{F}_1$ is algebraic over $\mathcal{F}$.

Proof. Pick $f \in \mathcal{F}[x]$ such that the IVP fails for f but holds for all $g \in \mathcal{F}[x]$ of smaller degree. Let $a, b \in \mathcal{F}[x]$ be such that $a < b$, $f(a) < 0 < f(b)$, $f(c) \neq 0$ for all $c \in [a, b] = \{c \in |\mathcal{F}| : a \leq c \leq b\}$. We employ a sort of "Dedekind cut" construction. Put $A = \{a' \in [a, b] : f(c) < 0 \text{ for all } c \in [a, a']\}$ and

$B = [a,b] \setminus A$. Obviously $A, B \neq \emptyset$, $[a,b] = A \cup B$, and $a' < b'$ for all $a' \in A$, $b' \in B$.

Claim: A has no greatest element and B has no least element.

Proof: This is clear since f is "continuous" (by an ϵ-δ-argument).

Claim: For every $g \in \mathcal{F}[x]$ of degree $< \deg(f)$ we can find $a' \in A$, $b' \in B$ so that g has constant sign (positive, negative, or zero) on the interval $[a', b']$.

Proof: If $g \equiv 0$ there is nothing to prove. Otherwise g has only finitely many roots in $|\mathcal{F}|$. Find $a' \in A$, $b' \in B$ such that $[a', b']$ does not contain any of these roots. By the IVP for g, we conclude that g does not change sign on $[a', b']$.

Claim: f is irreducible. For suppose $f = f_1 \cdot f_2$, $\deg(f_i) < \deg(f)$, $i = 1, 2$. We must have $f_i(a)f_i(b) < 0$ for at least one i. Then by the IVP for f_i we must have $f_i(c) = 0$ for some $c \in [a', b']$. Hence $f(c) = 0$, contradiction.

Let $\mathcal{F}_1 = \mathcal{F}(\alpha) = \mathcal{F}[x]/f$ where $f(\alpha) = 0$. (See Lemma 6.1.10.) To define the order relation on $\mathcal{F}(\alpha)$, recall that a typical element of $\mathcal{F}(\alpha)$ is of the form $g(\alpha)$ where $\deg(g) < \deg(f)$. Put $g(\alpha) > 0$ if $g > 0$ on some interval $[a', b']$, $a' \in A$, $b' \in B$; $g(\alpha) < 0$ if $g < 0$ on some interval $[a', b']$, $a' \in A$, $b' \in B$; $g(\alpha) = 0$ if $g \equiv 0$.

Claim: this makes $\mathcal{F}(\alpha)$ an ordered field. The only nontrivial axiom to be checked is the product rule: the product of positive elements is positive. Suppose $g_i(\alpha) > 0$, $\deg(g_i) < \deg(f)$, $i = 1, 2$. Let $g_1(\alpha) \cdot g_2(\alpha) = g_3(\alpha)$, $\deg(g_3) < \deg(f)$. We must show $g_3(\alpha) > 0$. We have

$$g_1(x) \cdot g_2(x) - g_3(x) \;=\; f(x) \cdot h(x)\,.$$

Put $m = \deg(f)$. The left side has degree $\leq 2m - 2$. Hence so does the right side, so $\deg(h) \leq m - 2$. Let $a' \in A$, $b' \in B$ be such that g_1, g_2, g_3, h all have constant sign on $[a', b']$.

Case 1: $h > 0$ on $[a', b']$. Let $c \in [a', b']$ be such that $f(c) < 0$. then $g_3(c) > 0$ since

$$\underbrace{g_1(c) \cdot g_2(c)}_{\text{pos}} \;=\; g_3(c) + \underbrace{f(c) \cdot h(c)}_{\text{neg}}\,,$$

hence $g_3 > 0$ on $[a', b']$. Hence $g_3(\alpha) > 0$.

Case 2: $h < 0$ on $[a', b']$. Take $c \in [a', b']$ so that $f(c) > 0$. Then by an argument similar to the above, $g_3(c) > 0$, so $g_3(\alpha) > 0$.

Case 3: $h = 0$ on $[a', b']$. Hence $h \equiv 0$ so $g_3 = g_1 \cdot g_2 > 0$ on $[a', b']$, hence $g_3(\alpha) > 0$. This completes the proof.

10.

Definition. Let $\mathcal{F}$ be an ordered field. A *real closure* of $\mathcal{F}$ is an ordered field $\mathcal{G} \supseteq \mathcal{F}$ so that (1) $\mathcal{G}$ is algebraic over $\mathcal{F}$; (2) $\mathcal{G}$ is real closed.

11.

Proposition. Any ordered field $\mathcal{F}$ has a real closure.

Proof. This is an easy consequence of the previous lemma plus Zorn's lemma (or transfinite induction). Let $\mathcal{B}$ be the algebraic closure of $\mathcal{F}$ regarded as a field. Form a transfinite sequence of extensions $\mathcal{F} = \mathcal{F}_0$, $\mathcal{F}_{\alpha+1}$ = proper ordered extension of $\mathcal{F}_\alpha$, $\mathcal{F}_\beta = \bigcup_{\alpha<\beta} \mathcal{F}_\beta$ for limit β, all within $\mathcal{B}$. The process must eventually stop, so by lemma 9.1.9 we then have a real closure of $\mathcal{F}$.

9.2 Uniqueness of real closure

The purpose of this section is to prove that the real closure of an ordered field $\mathcal{F}$ is unique up to isomorphism over $\mathcal{F}$. The usual proofs of this result use either Galois theory or Sturm's test. We give a more elementary argument based on Rolle's theorem.

1.

Lemma (Rolle's theorem). Let $\mathcal{G}$ be a real closed ordered field, $g \in \mathcal{G}[x]$. Between any two distinct roots of g lies at least one root of g'.

(Here g' is the formal derivative of g. If $g(x) = a_n x^n + \cdots + a_1 x + a_0$ then $g'(x) = n a_n x^{n-1} + \cdots + a_1$.)

Proof. Given $a < b$, $g(a) = g(b) = 0$. Say $g(x) = (x-a)^k (x-b)^l p(x)$ where $p \in \mathcal{G}[x]$ and $p(a) \neq 0$, $p(b) \neq 0$. We may safely assume that $g(c) \neq 0$ for $a < c < b$. Hence $p(c) \neq 0$ for $c \in [a, b]$. Hence p has constant sign on $[a, b]$. We have

$$g(x) = (x-a)^{k-1}(x-b)^{l-1} q(x)$$

where

$$q(x) = [k(x-b) + l(x-a)]p(x) + (x-a)(x-b)p'(x)\,.$$

We have $q(a) = k(a-b)p(a)$, $q(b) = l(b-a)p(b)$ so $q(a)$ and $q(b)$ have opposite sign, hence $q(c) = 0$ for some $a < c < b$. Hence $g'(c) = 0$.

2.

Lemma (maximum principle). Let $\mathcal{G}$ be a real closed ordered field, $g \in \mathcal{G}[x]$, $a, b \in \mathcal{G}$, $a < b$. Then g takes on a maximum value in the interval $[a, b]$, say at γ. If $a < \gamma < b$ then $g'(\gamma) = 0$.

Proof. If g is a linear function there is nothing to prove. Otherwise g' has only finitely many roots in the interval $[a, b]$. Let

$$a = c_0 < c_1 < \cdots < c_k < c_{k+1} = b$$

be the set consisting of a, b, and all the roots of g' in $[a, b]$. We claim that g takes on a maximum value at one of these c_i's. Suppose not. Then there exists $d \in [a, b]$ such that $g(d) > g(c_i)$ for all i. Then $c_i < d < c_{i+1}$ for some $i \leq k$. Hence by Rolle's theorem there must exist e such that $c_i < e < c_{i+1}$ and $g'(e) = 0$. This is a contradiction.

3.

Lemma. Let $\mathcal{F}$ be an ordered field. Let $g \in \mathcal{F}[x]$ be a nonconstant polynomial such that $\mathcal{F}$ satisfies the IVP for all polynomials of degree $\leq \deg(g)$. Suppose $g(\alpha) = 0$ where α lies in some ordered field extension of $\mathcal{F}$. Then actually $\alpha \in |\mathcal{F}|$.

Proof. We prove the lemma by induction on $n = \deg(g)$. Assume that the conclusion fails, i.e. $\alpha \notin |\mathcal{F}|$. It follows that g is irreducible. (Otherwise we could apply the induction hypothesis to the factors and get a contradiction.) In particular g has no roots in $|\mathcal{F}|$, so by the IVP g has constant sign on $|\mathcal{F}|$. Without loss of generality, assume that $g < 0$ on $|\mathcal{F}|$.

Since $g(\alpha) = 0$, an easy calculation shows that $|\alpha| < 1 + |a_1| + \cdots + |a_n| = M$ where $g(x) = x^n + a_1 x^{n-1} + \cdots + a_n$. Thus we have $g(-M) < 0$, $g(\alpha) = 0$, $g(M) < 0$, and $-M < \alpha < M$. Let $\mathcal{G}$ be a RCOF extending $\mathcal{F}$ (by proposition 9.1.11). By the maximum principle we must have $g(\gamma) \geq 0$, $g'(\gamma) = 0$ for some $\gamma \in |\mathcal{G}|$ with $-M < \gamma < M$. By the induction hypothesis, $g'(\gamma) = 0$ gives $\gamma \in |\mathcal{F}|$. But then $g(\gamma) \geq 0$ contradicts the fact that $g < 0$ on $|\mathcal{F}|$.

4.

Proposition. If $\mathcal{G}$ is a RCOF and $\mathcal{G}_1$ is an ordered field extension of $\mathcal{G}$ which is algebraic over $\mathcal{G}$, then $\mathcal{G}_1 = \mathcal{G}$.

Proof. Immediate from the previous lemma.

5.

Proposition. Let $\mathcal{F}$ be an ordered field. Then any two real closures of $\mathcal{F}$ are isomorphic over $\mathcal{F}$.

Proof. Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be real closures of $\mathcal{F}$. If $\mathcal{F}$ is real closed, then $\mathcal{G}_1 = \mathcal{G}_2 = \mathcal{F}$ by the previous proposition. Therefore, by Zorn's lemma or transfinite induction, it suffices to show that if $\mathcal{F}$ is not real closed, then we can find $\alpha_i \in |\mathcal{G}_i| \setminus |\mathcal{F}|$, $i = 1, 2$, such that $\mathcal{F}(\alpha_1) \cong \mathcal{F}(\alpha_2)$ over $\mathcal{F}$ as ordered fields.

In order to find α_1, α_2 we imitate the proof of lemma 9.1.9. Let $f \in \mathcal{F}[x]$ be such that the IVP fails for f but holds for all g of lower degree. Let a, b, A, B be as in the proof of lemma 9.1.9. Since $\mathcal{G}_1$ is real closed, $|\mathcal{G}_1| \setminus |\mathcal{F}|$ contains α_1 such that $f(\alpha_1) = 0$ and $a' < \alpha_1 < b'$ for all $a' \in A$, $b' \in B$. We claim that these facts determine the structure of $\mathcal{F}(\alpha_1)$ as an ordered field. First of all, the field structure of $\mathcal{F}(\alpha_1)$ is determined by the fact that $\mathcal{F}(\alpha_1) \cong \mathcal{F}[x]/f$ since f is irreducible. Let $g \in \mathcal{F}[x]$ be of lower degree than f such that $g > 0$ on some interval $[a', b']$ in $\mathcal{F}$ where $a' \in A$, $b' \in B$. By the previous lemma we know that any root of g lying in $|\mathcal{G}_1|$ already lies in $|\mathcal{F}|$. Hence $g > 0$ on $[a', b']$ in $\mathcal{G}_1$. Hence in particular $g(\alpha_1) > 0$. Since every element of $\mathcal{F}(\alpha_1)$ is of the form $g(\alpha_1)$, $g \in \mathcal{F}[x]$, $\deg(g) < \deg(f)$, our claim is proved. Similarly we can find $\alpha_2 \in |\mathcal{G}_2| \setminus |\mathcal{F}|$ such that $f(\alpha_2) = 0$ and $a' < \alpha_2 < b'$ for all $a' \in A$, $b' \in B$, and we then have $\mathcal{F}(\alpha_1) \cong \mathcal{F}(\alpha_2)$ over $\mathcal{F}$ as ordered fields. This completes the proof.

6.

Definition. If $\mathcal{F}$ is an ordered field, we denote its real closure by $\overline{\mathcal{F}}$. (The previous result shows that $\overline{\mathcal{F}}$ is unique up to isomorphism over $\mathcal{F}$.)

7.

Corollary. Let $\mathcal{F}$ be an ordered field and let $\mathcal{G}$ be a RCOF extending $\mathcal{F}$. Then $\mathcal{G}$ contains (an isomorphic copy of) $\overline{\mathcal{F}}$.

Proof. Let $\mathcal{F}_1$ be the algebraic closure of $\mathcal{F}$ within $\mathcal{G}$. Clearly $\mathcal{F}_1$ is a real closure of $\mathcal{F}$. By the previous result it follows that $\overline{\mathcal{F}}$ is isomorphic to $\mathcal{F}_1$ over $\mathcal{F}$.

9.3 Quantifier elimination for RCOF

1.

Lemma. Let $\mathcal{G}$ be an RCOF and let $\mathcal{G}(\alpha)$ be a simple ordered field extension of $\mathcal{G}$, $\alpha \notin |\mathcal{G}|$. Then α is transcendental over $\mathcal{G}$, and the ordering of $\mathcal{G}$ is completely determined by the sets $A = \{a \in |\mathcal{G}| : a < \alpha\}$ and $B = \{b \in |\mathcal{G}| : \alpha < b\}$.

Proof. The fact that α is transcendental over $\mathcal{G}$ follows immediately from the uniqueness of the real closure of $\mathcal{G}$. The ordering of $\mathcal{G}(\alpha)$ is completely determined by that of $\mathcal{G}[\alpha]$ which is completely determined by the following claim.

Claim: Let $g \in \mathcal{G}[x]$.

(i) If $A = \emptyset$ (respectively $B = \emptyset$) then the sign of $g(-\alpha)$ (respectively $g(\alpha)$) is the same as that of the leading coefficient of g.

(ii) If $A \neq \emptyset \neq B$, then $g(\alpha) > 0$ if and only if there exist $a \in A$, $b \in B$ such that $g(c) > 0$ for all $c \in |\mathcal{G}|$, $a < c < b$.

Part (i) is obvious, and part (ii) is an easy consequence of lemma 9.2.3. (One may compare the proof of lemma 9.1.9, but note that in the present situation it is possible that A has a greatest element.)

2.

Theorem. The theory of real closed ordered fields is the model completion of the theory of ordered fields.

Proof. It suffices to show that if $\mathcal{F} \subseteq \mathcal{F}_1$ are ordered fields and if $\mathcal{F} \subseteq \mathcal{G}$ where $\mathcal{G}$ is a real closed ordered field and $\mathcal{G}$ is κ-saturated, $\kappa = \|\mathcal{F}_1\|$, then $\mathcal{F}_1$ is embeddable into $\mathcal{G}$ over $\mathcal{F}$. (Cf. theorem 8.3.2.) By transfinite induction we may safely assume that either (1) $\mathcal{F}_1$ is algebraic over $\mathcal{F}$, or (2) $\mathcal{F}_1 = \mathcal{F}(\alpha)$, α transcendental over $\mathcal{F}$, $\alpha \notin |\mathcal{F}|$, and $\mathcal{F}$ is real closed.

In case (1), we have $\mathcal{F}_1 \subseteq \overline{\mathcal{F}}$ and by corollary 9.2.7 $\overline{\mathcal{F}}$ is embeddable into $\mathcal{G}$ over $\mathcal{F}$. In case (2), put $A = \{a \in |\mathcal{F}| : a < \alpha\}$ and $B = \{b \in |\mathcal{F}| : \alpha < b\}$. Consider the set of formulas

$$\{\underline{a} < v : a \in A\} \cup \{v < \underline{b} : b \in B\} \cup \{\underline{f}(v) \neq 0 : f \in \mathcal{F}[x], f \not\equiv 0\}\,.$$

This set of formulas is finitely realizable in $\mathcal{F}$. Hence by saturation it is realized in $\mathcal{G}$, say by $\beta \in |\mathcal{G}|$. By the previous lemma it follows that $\mathcal{F}(\alpha)$ is order isomorphic to $\mathcal{F}(\beta)$ over $\mathcal{F}$. This completes the proof.

3.

Corollary (Tarski). The theory of real closed ordered fields is complete and decidable.

Proof. Any RCOF has an ordered subfield isomorphic to

$$\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1, <)\,.$$

By the previous theorem, any two RCOFs are elementary equivalent over $\mathcal{Q}$. In particular any two RCOFs are elementarily equivalent, so the theory RCOF is complete. Decidability follows by theorem 4.2.3.

(Tarski actually gave a decision procedure based on a specific explicit quantifier elimination procedure.)

4.

Corollary (Tarski's transfer principle). Any first order sentence which is true in the real field $\mathcal{R} = (\mathbb{R}, +, -, \cdot, 0, 1, <)$ is true in any real closed ordered field.

5.

Examples. (1) The fact that the complex field $\mathcal{C}$ is algebraically closed (i.e. the fundamental theorem of algebra) is expressible as a set of first order sentences in the theory of RCOF. (We identify $a \pm b\sqrt{-1} \in \mathbb{C}$ with an ordered pair $(a, b) \in \mathbb{R}^2$.) It follows by Tarski's principle that, if $\mathcal{G}$ is any RCOF, then $\mathcal{G}(\sqrt{-1}) = \mathcal{G}[x]/(x^2 + 1)$ is algebraically closed.

(2) Bott and Milnor have used homotopy theory to prove that the only finite dimensional nonassociative algebras over $\mathbb{R}$ are those of dimensions 1, 2, 4, and 8 respectively (i.e. $\mathbb{R}$, $\mathbb{C}$, quaternions, octonians). By Tarski's transfer principle, the same holds for any RCOF.

6.

Corollary (Tarski). The first order theory of the real field

$$\mathcal{R} = (\mathbb{R}, +, -, \cdot, 0, 1, <)$$

is decidable.

(In other words, there is a decision procedure for "high school mathematics", including plane and solid geometry since the latter can be interpreted into the first order theory of $\mathcal{R}$ using Cartesian coordinates.)

7.

Corollary (Tarksi). Given a finite system S of equations and inequalities in n variables with rational coefficients, we can decide recursively whether S has a solution in reals.

8.

Theorem (Tarksi). The theory of real closed ordered fields admits elimination of quantifiers.

Proof. From proposition 9.1.4 and theorem 9.3.3 it follows that RCOF is the model completion of the theory of ordered domains. The latter theory is universal so by theorem 8.2.3 it follows that RCOF admits elimination of quantifiers.

9.

Corollary (Tarski). The theory of real closed ordered fields is model complete.

10.

Corollary (Artin). Let S be a finite system of equations and inequalities in n unknowns with coefficients in an ordered field $\mathcal{F}$. If S has a solution in some ordered field extension of $\mathcal{F}$, then S has a solution in $\overline{\mathcal{F}}$, the real closure of $\mathcal{F}$.

Proof. Immediate from model completeness plus corollary 9.2.7.

In the next section we shall apply this result to obtain Artin's solution of Hilbert's 17th problem.

9.4 The solution of Hilbert's 17th problem

1.

Let $\mathcal{F}$ be an ordered field, $\overline{\mathcal{F}}$ its real closure. Put $\mathcal{B} = \mathcal{F}(x_1, \ldots, x_n) =$ the field of rational functions over $\mathcal{F}$ in n variables $x_1, \ldots, x_n$. Hilbert's 17th problem is concerned with the representation of rational functions as sums of squares.

Theorem (Artin). Let $f \in |\mathcal{B}|$, $f \not\equiv 0$. Suppose that $f(a_1, \ldots, a_n) \geq 0$ for all $a_1, \ldots, a_n \in |\overline{\mathcal{F}}|$ such that $f(a_1, \ldots, a_n)$ is defined. Then there exists a positive integer k such that

$$f \;=\; \sum_{i=1}^{k} c_i g_i^2$$

where $c_i \in |\mathcal{F}|$, $c_i > 0$, $g_i \in |\mathcal{B}|$, $g_i \not\equiv 0$.

Proof. Let $I \subset |\mathcal{B}|$ be the set of all f of this form. We have

(i) $g^2 \in I$ for all $g \in |\mathcal{B}|$, $g \not\equiv 0$

(ii) $g, h \in I$ implies $g + h, g \cdot h \in I$

(iii) $0 \notin I$.

A set $I \subset |\mathcal{B}|$ satisfying (i), (ii) and (iii) is called an *order ideal.*

The following lemma holds for any order ideal in any field $\mathcal{B}$.

Lemma. Let I be an order ideal, $f \notin I$, $f \not\equiv 0$. Then there exists an order ideal $J \supseteq I$ such that $-f \in J$.

Proof of lemma. Assume $f \neq 0$. Let J be the set of elements of $|\mathcal{B}|$ of the form

$$p(-f) \quad = \quad (-f)^{m_1} \cdot g_1 + \cdots + (-f)^{m_l} \cdot g_l$$

where $l \geq 1$, $g_i \in I$, $m_i \geq 0$. Clearly $-f \in J$ and J satisfies (i) and (ii). Suppose (iii) fails for J, say $p(-f) = 0$. We then have

$$p(-f) \quad = \quad q(f^2) - f \cdot r(f^2) \quad = \quad 0$$

where $q(f^2)$ and $r(f^2) \in I$ or $= 0$. If $r(f^2) = 0$ then $q(f^2) = 0$ so both are 0. This is impossible since $l \geq 1$. Hence $r(f^2) \neq 0$. Hence

$$f \quad = \quad \frac{q(f^2)}{r(f^2)} \quad = \quad q(f^2) \cdot r(f^2) \cdot \left(\frac{1}{r(f^2)}\right)^2 \quad \in \quad I$$

since $f \neq 0$. This proves the lemma.

Now to prove the theorem, suppose $f \notin I$, $f \neq 0$. By the lemma, get an order ideal $J \supseteq I$ containing $-f$. By Zorn's lemma let $K \supseteq J$ be a maximal order ideal. Then K satisfies (i), (ii), (iii) and also

(iv) for all $g \neq 0$ either $g \in K$ or $-g \in K$.

This is immediate from the lemma, by maximality of K.

For $g, h \in |\mathcal{B}|$ define $g < h$ if and only if $h - g \in K$. It is easy to check that this definition makes $\mathcal{B}$ an ordered field. Furthermore, if $c \in |\mathcal{F}|$ is positive in $\mathcal{F}$, then it is positive in $\mathcal{B}$. So $\mathcal{B}$ is an ordered field extension of $\mathcal{F}$.

We see now that there exist $b_1 \ldots, b_n$ lying in an ordered field extension of $\mathcal{F}$, such that $f(b_1, \ldots, b_n) < 0$ (namely $b_i = x_i$ lying in $\mathcal{B}$). Hence by model completeness of RCOF (see also corollary 9.3.10), it follows that there exist $a_1, \ldots, a_n \in |\overline{\mathcal{F}}|$ such that $f(a_1, \ldots, a_n) < 0$. Q.E.D.

2.

Corollary (Artin's solution of Hilbert's 17th problem). Let $\mathcal{F}$ be either the real field $\mathbb{R}$ or the rational field $\mathbb{Q}$. Suppose $f(x_1, \ldots, x_n) \in \mathcal{F}(x_1, \ldots, x_n)$ is *definite over* $\mathcal{F}$, i.e. $f(a_1, \ldots, a_n) \geq 0$ for all $a_1, \ldots, a_n \in |\mathcal{F}|$ such that $f(a_1, \ldots, a_n)$ is defined. Then there exist $g_1, \ldots, g_k$ in $\mathcal{F}(x_1, \ldots, x_n)$ such that

$$f = \sum_{i=1}^{k} g_i^2 .$$

Proof. For the reals this is immediate from the previous theorem. For the rationals, it is immediate from the theorem plus the following two observations: (1) $\mathbb{Q}$ is dense in $\mathbb{R}$; (2) every positive element of $\mathbb{Q}$ is the sum of 4 squares in $\mathbb{Q}$.

3.

Remark. An ordered field $\mathcal{F}$ is said to be *Hilbertian* if (1) $\mathcal{F}$ is dense in $\overline{\mathcal{F}}$, (2) every positive element of $\mathcal{F}$ is a sum of squares in $\mathcal{F}$. For example, $\mathcal{F}$ = reals, $\mathcal{F}$ = rationals, $\mathcal{F}$ = any RCOF are Hilbertian. Clearly the above corollary holds for any Hilbertian ordered field. McKenna[1] has proved the converse: If the corollary holds for $\mathcal{F}$, then $\mathcal{F}$ is Hilbertian.

4.

We now derive a corollary concerning effective bounds, analogous to theorem 6.4.3.

Corollary. We can find a recursive function $K : \omega^2 \to \omega$ with the following property:

> Let $\mathcal{F}$ be an ordered field, and let $f \in \mathcal{F}(x_1, \ldots, x_n)$ be such that $f(a_1, \ldots, a_n) \geq 0$ for all $a_1, \ldots, a_n \in \overline{\mathcal{F}}$ such that this is defined. If f is of degree $\leq d$ then $f = \sum_{i=1}^{k} c_i g_i^2$ for some $k \leq K(n, d)$, $c_i \in \mathcal{F}$, $c_i \geq 0$, $g_i \in \mathcal{F}(x_1, \ldots, x_n)$ of degree $\leq K(n, d)$.

[1] SLNM 498, pp. 220-230.

Proof. Let T be the theory of real closed ordered fields with a distinguished subfield, $\mathcal{F}$. Let σ_{ndm} be the sentence asserting that the desired conclusion holds with $K(n,d)$ replaced by m. We have $\forall n\,\forall d\,\exists m\ \sigma_{ndm} \in T$ (otherwise we could use compactness to construct a counterexample to theorem 9.4.1). So put $K(n,d) =$ least m such that $\sigma_{ndm} \in T$. This is recursive since T is decidable.

5.

Remark. Ax (unpublished) and Pfister (1967)[2] have shown that for real closed $\mathcal{F}$ we can always get $k \leq 2^n$ (but apparently with no bound on the degrees). There are some open problems in this area.

6.

Exercise. Prove the following result which is known as the *Nullstellensatz for ordered fields.*

Theorem (Dubois[3]). Let $f, h_1, \ldots, h_m \in \mathcal{F}[x_1, \ldots, x_n]$ be polynomials in n variables over an ordered field $\mathcal{F}$. Suppose that every common zero of $h_1, \ldots, h_m$ in $\overline{\mathcal{F}}$, the real closure of $\mathcal{F}$, is a zero of f. Then there exist nonnegative integers r and k such that

$$f^{2r} + \sum_{i=1}^{k} c_i g_i^2 \;=\; \sum_{j=1}^{m} p_j h_j$$

where $c_i \in |\mathcal{F}|$, $c_i \geq 0$, $g_i, p_j \in \mathcal{F}[x_1, \ldots, x_n]$.

Deduce a version with effective bounds.

[2] Proc. Symp. Pure Math. 28, pp. 483-491.
[3] Ark. Mat. 8 (1969), pp. 111-114. See also Prestel's monograph.

Chapter 10

Prime models (countable case)

10.1 The omitting types theorem

1.

Definition. Let p be an n-type over a theory T. A model $\mathcal{A}$ of T is said to *omit* p if there is no n-tuple $a_1, \ldots, a_n \in |\mathcal{A}|$ which realizes p.

We have already seen how to use the compactness theorem to construct a model which realizes p. It is somewhat more difficult to construct a model which omits p. Indeed, such a model may not even exist. The omitting types theorem gives a sufficient condition for the existence of a model of T which omits p. (If T is complete, this sufficient condition is also necessary.)

2.

Definition. We say that p is *principal* over T if it is generated by a single formula, i.e. there exists $\varphi \in F_n(T)$ such that $p = \{\psi \in F_n(T) : T \models \forall v_1 \cdots v_n(\varphi(v_1, \ldots, v_n) \rightarrow \psi(v_1, \ldots, v_n))$. Such a φ is called a *generator* of p.

We say that p is *essentially nonprincipal* over T if p is not included in any principal n-type over T.

3.

Theorem (the omitting types theorem). Let T be a countable theory and let p be an n-type over T. Suppose that p is essentially nonprincipal over T.

Then there exists a model of T which omits p.

Proof. We go back to the Henkin proof of the completeness theorem. Let $C = \{\underline{c}_i : i \in \omega\}$ be a countable set of new constant symbols. Let $\{\sigma_s : s \in \omega\}$ be an enumeration of all sentences σ with $\mathrm{sig}(\sigma) \subseteq \mathrm{sig}(T) \cup C$. Let $\{\varphi_S(v) : s \in \omega\}$ be an enumeration of all formulas $\varphi(v)$ with no free variables other than v, such that $\mathrm{sig}(\varphi(v)) \subseteq \mathrm{sig}(T) \cup C$. Let $\{\langle \underline{c}_1^s, \ldots, \underline{c}_n^s \rangle : s \in \omega\}$ be an enumeration of all n-tuples of constant symbols from C.

Stage 0: $T_0 = T$.

Stage $3s + 1$: Put $T_{3s+1} = T_{3s} \cup \{\sigma_s\}$ if this is consistent, otherwise $T_{3s+1} = T_{3s} \cup \{\neg\sigma_s\}$.

Stage $3s + 2$: Let $h(s)$ be the least i such that $\underline{c}_i$ does not appear in $T_{3s+1} \cup \{\varphi_s(v)\}$. Put $T_{3s+2} = T_{3s+1} \cup \{(\exists v \varphi_s(v)) \to \varphi_s(\underline{c}_{h(s)})\}$.

Stage $3s + 3$: Let q be the set of all $\psi \in F_n(T)$ such that $T_{3s+2} \cup \{\neg\psi(\underline{c}_1^s, \ldots, \underline{c}_n^s)\}$ is inconsistent. Clearly q is a principal n-type over T. Hence q does not include p. Pick a formula $\varphi \in p \setminus q$ and put $T_{3s+3} = T_{3s+2} \cup \{\neg\varphi(\underline{c}_1^s, \ldots, \underline{c}_n^s)\}$.

Finally put $T_\infty = \bigcup_{s \in \omega} T_s$. Define a Henkin model $\mathcal{M}$ as in the proof of theorem 3.1.2. Each n-tuple of elements of $|\mathcal{M}|$ is denoted by an n-tuple of Henkin constants $\underline{c}_1^s, \ldots, \underline{c}_n^s$. Thus $\mathcal{M}$ omits p.

4.

Remark. The hypothesis that T is countable cannot be omitted from the omitting types theorem.

5.

We digress to present a typical application of the omitting types theorem.

Theorem. Let $\mathcal{A} = (|\mathcal{A}|, \in^{\mathcal{A}})$ be a countable model of ZF set theory. Then $\mathcal{A}$ has a proper elementary extension $\mathcal{B}$ such that $\omega^{\mathcal{A}} = \omega^{\mathcal{B}}$, i.e. $\mathcal{A}$ and $\mathcal{B}$ have the same natural numbers.

Proof. Let T be the theory generated by

$$(\text{elementary diagram of } \mathcal{A}) \cup \{\mathrm{rank}(\underline{c}) > \mathrm{rank}(\underline{a}) : a \in |\mathcal{A}|\}$$

where $\underline{c}$ is a new constant symbol. T is consistent by the compacttness theorem, and every model of T gives rise to a proper elementary extension of $\mathcal{A}$. Let p be the 1-type over T generated by $\{v_1 \in \underline{\omega}\} \cup \{v_1 \neq \underline{n} : n \in \omega^{\mathcal{A}}\}$. To prove the theorem, it suffices to find a model of T which omits p.

Suppose this were not possible. By the omitting types theorem, p is included in a principal 1-type over T. Let $\varphi \in F_n(T)$ be a generator of such a principal 1-type. In particular $T \models \neg\varphi(\underline{n})$ for all $n \in \omega^{\mathcal{A}}$. Let $\varphi(v) = \psi(\underline{c}, v)$ where $\mathrm{sig}(\psi) \subseteq \mathrm{sig}(\text{elementary diagram of } \mathcal{A})$. Then for each $n \in \omega^{\mathcal{A}}$ we have $T \models \neg\psi(\underline{c}, \underline{n})$. Hence by compactness, for each $n \in \omega^{\mathcal{A}}$ there exists $a \in |\mathcal{A}|$ such that the elementary diagram of $\mathcal{A}$ contains

$$\forall x \, (\mathrm{rank}(x) > \mathrm{rank}(\underline{a}) \rightarrow \neg\psi(x, \underline{n}))\,.$$

Hence

$$\mathcal{A} \models \forall n \in \omega \, \exists y \, \forall x \, (\mathrm{rank}(x) > \mathrm{rank}(y) \rightarrow \neg\psi(x, n))\,.$$

By the replacement axiom in $\mathcal{A}$ it follows that

$$\mathcal{A} \models \exists y \, \forall n \in \omega \, \forall x \, (\mathrm{rank}(x) > \mathrm{rank}(y) \rightarrow \neg\psi(x, n))\,.$$

Let $a \in |\mathcal{A}|$ be such that

$$\mathcal{A} \models \forall n \in \omega \, \forall x \, (\mathrm{rank}(x) > \mathrm{rank}(\underline{a}) \rightarrow \neg\psi(x, n))\,.$$

Since $T \models \mathrm{rank}(\underline{c}) > \mathrm{rank}(\underline{a})$ it follows that $T \models \forall n \in \omega \, \neg\psi(\underline{c}, \underline{n})$. In other words, $T \models \forall v_1 \, (\varphi(v_1) \rightarrow v_1 \notin \underline{\omega})$. This contradicts the assumption that the principal 1-type generated by φ includes p.

6.

The omitting types theorem can be generalized as follows:

Theorem. Let T be a countable theory. For each $i \in \omega$, let p_i be an essentially nonprincipal n_i-type over T. Then T has a model which omits p_i for each $i \in \omega$.

The proof is a straightforward generalization of the proof of the omitting types theorem.

7.

Exercise. Use the above generalization of the omitting types theorem to prove the following result, which generalizes theorem 10.1.4 above.

Theorem (Keisler and Morley). Let $\mathcal{A}$ be a countable model of ZF set theory. Then $\mathcal{A}$ has a proper elementary *end extension*, i.e. a proper elementary extension $\mathcal{B}$ such that $\text{rank}(a) < \text{rank}(b)$ for all $a \in |\mathcal{A}|$, $b \in |\mathcal{B}| \setminus |\mathcal{A}|$.

In the next section, the omitting types theorem will be used to study prime models.

10.2 Prime models

1.

Definition. Let T be a complete theory. A model $\mathcal{A}$ of T is said to be *prime* if every model of T has an elementary submodel which is isomorphic to $\mathcal{A}$.

2.

Examples.

1. Let $T = \mathsf{ACF}(0)$. This theory has a prime model $\overline{\mathcal{Q}}$ = the algebraic closure of the rational field $\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1)$.

2. Similarly for $T = \mathsf{ACF}(p)$, the theory of algebraically closed fields of prime characteristic p. The prime model is the algebraic closure of $\mathcal{F}_p = (\mathbb{F}_p, +, -, \cdot, 0, 1)$.

3. Generalizing examples 1 and 2, let $\mathcal{A}$ be any field and let $T = \mathsf{ACF} \cup$ (diagram of $\mathcal{A}$). We know that T is complete because ACF admits elimination of quantifiers. The prime model of T is $\overline{\mathcal{A}}$, the algebraic closure of $\mathcal{A}$.

4. $T = \mathsf{RCOF}$. The prime model is $\overline{\mathcal{Q}}$ = the real closure of the ordered field $\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1, <)$.

5. $T = \mathsf{RCOF} \cup$ (diagram of $\mathcal{F}$) where $\mathcal{F}$ is any ordered field. The prime model is $\overline{\mathcal{F}}$, the real closure of $\mathcal{F}$.

6. T = theory of dense linear ordering without end points. The prime model is $(\mathbb{Q}, <)$.

7. We give an example of a complete countable theory with no prime model. Let T have 1-place relations $S_i(x)$, $i \in \omega$, and axioms saying that the S_i's are *independent*, i.e. any nontrivial Boolean combination of them is nonempty. The completeness of T can be proved by quantifier elimination. It is easy to see that no model of T is prime.

3.

Remark. It is clear from the definition of prime model that any elementary submodel of a prime model is prime. Example 6 shows that a prime model may have proper elementary submodels.

4.

The purpose of this section is to establish necessary and sufficient conditions for a complete theory T to have a prime model, and to establish the uniqueness of prime models when they exist. Throughout this chapter we deal with countable theories T; for the uncountable case, see chapter 13.

5.

The following lemma is just a restatement of the omitting types theorem in the special case when T and p are complete.

Lemma. Let T be a countable complete theory, and let $p \in S_n(T)$ be a complete n-type over T. The following are equivalent.

(i) p is principal over T.

(ii) p is realized in every model of T.

Proof. Assume that p is principal, say generated by $\varphi \in F_n(T)$. In particular $T \cup \{\exists v_1 \cdots v_n\, \varphi(v_1, \ldots, v_n)\}$ is consistent. Since we are assuming that T is complete, it follows that $T \models \exists v_1 \cdots v_n\, \varphi(v_1, \ldots, v_n)$. Let $\mathcal{A}$ be any model of T. Since $\mathcal{A} \models \exists v_1 \cdots v_n\, \varphi(v_1, \ldots, v_n)$, there exist $a_1, \ldots, a_n \in |\mathcal{A}|$ such that $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$. Then clearly the n-tuple $a_1, \ldots, a_n$ realizes p.

Conversely, suppose that p is nonprincipal over T. Since p is complete, it follows that p is essentially nonprincipal over T. Hence by the omitting types theorem, there exists a model of T which omits p.

6.

Definition. A structure $\mathcal{A}$ is *atomic* if for every $a_1, \ldots, a_n \in |\mathcal{A}|$ the complete n-type realized by $a_1, \ldots, a_n$ is principal over $\mathrm{Th}(\mathcal{A})$.

7.

Theorem (Vaught). Let T be a countable complete theory. The following are equivalent for $\mathcal{A} \in \mathrm{Mod}(T)$.

(i) $\mathcal{A}$ is prime.

(ii) $\mathcal{A}$ is countable and atomic.

Proof. (i) $\Rightarrow$ (ii). Assume $\mathcal{A}$ is prime. Since T has a countable model $\mathcal{B}$ and $\mathcal{A}$ is isomorphic to an elementary submodel of $\mathcal{B}$, it follows that $\mathcal{A}$ is countable. To show that $\mathcal{A}$ is atomic, let $a_1, \ldots, a_n \in |\mathcal{A}|$ and let p be the complete n-type over T realized by $a_1, \ldots, a_n$. Since $\mathcal{A}$ is prime, p is realized in every model of T. Hence by the previous lemma it follows that p is principal. Thus $\mathcal{A}$ is atomic.

(ii) $\Rightarrow$ (i). Assume that $\mathcal{A}$ is countable and atomic. Let $|\mathcal{A}| = \{a_n : n \in \omega\}$ be an enumeration of $|\mathcal{A}|$. Let $\mathcal{B}$ be a model of T. We want to construct an elementary embedding of $\mathcal{A}$ into $\mathcal{B}$. Assume inductively that we have chosen $b_0, \ldots, b_{n-1} \in |\mathcal{B}|$ so that $(\mathcal{A}, a_i)_{i<n} \equiv (\mathcal{B}, b_i)_{i<n}$. Since $\mathcal{A}$ is atomic, so is $(\mathcal{A}, a_i)_{i<n}$. [For, let $c_1, \ldots, c_k \in |\mathcal{A}|$. The complete $(n+k)$-type realized by $\langle a_0, \ldots, a_{n-1}, c_1, \ldots, c_k\rangle$ is principal over T, say generated by $\theta(u_0, \ldots, u_{n-1}, w_1, \ldots, w_k)$. Put $\psi(w_1, \ldots, w_k) \equiv \theta(\underline{a}_0, \ldots, \underline{a}_{n-1}, w_1, \ldots, w_k)$. Then ψ generates the complete k-type realized by $c_1, \ldots, c_k$ over $\mathrm{Th}((\mathcal{A}, a_i)_{i<n})$.]

Let $p_n \in S_1((\mathcal{A}, a_i)_{i<n})$ be the complete 1-type realized by a_n. Since p_n is principal, it follows by the previous lemma that p_n is realized in $(\mathcal{B}, b_i)_{i<n}$, say by $b_n \in |\mathcal{B}|$. Then $(\mathcal{A}, a_i)_{i\leq n} \equiv (\mathcal{B}, b_i)_{i\leq n}$. Finally $(\mathcal{A}, a_i)_{i\in\omega} \equiv (\mathcal{B}, b_i)_{i\in\omega}$. Thus $a_i \mapsto b_i$ gives an elementary embedding of $\mathcal{A}$ into $\mathcal{B}$.

8.

Theorem (Vaught). Let T be a countable complete theory. If T has a prime model, it is unique up to isomorphism.

Proof. By the previous theorem it suffices to show that any two countable atomic models of T are isomorphic. We use a back-and-forth argument. The inductive step is as in the proof of the previous theorem. The details are routine.

9.

Theorem (Vaught). Let T be a countable complete theory. The following are equivalent.

(i) T has a prime model.

(ii) T has an atomic model.

(iii) Every principal n-type over T is included in a complete principal n-type over T.

10.

Remark. A Boolean algebra B is said to be *atomic* if for all $b \in B$, $b \neq 0$, there exists $a \leq b$ such that a is an *atom*, i.e. $a \neq 0$ and there is no a_1 such that $0 < a_1 < a$. The condition (iii) in the theorem can be restated as follows: for each $n \in \omega$, the Boolean algebra $B_n(T)$ is atomic.

11.

Proof of theorem 9.

(i) $\Rightarrow$ (ii). Let $\mathcal{A}$ be a prime model of T. By 10.2.7 above, $\mathcal{A}$ is atomic.

(ii) $\Rightarrow$ (iii). Let $\mathcal{A}$ be an atomic model. Let q be a principal n-type over T, say generated by $\varphi \in F_n(T)$. Since T is complete, we have $T \models \exists v_1 \ldots v_n \, \varphi(v_1, \ldots, v_n)$. Hence q is realized in $\mathcal{A}$, say by $a_1, \ldots, a_n$. Let $p \in S_n(T)$ be the complete n-type realized by $a_1, \ldots, a_n$. Then $p \supseteq q$ and p is principal.

(iii) $\Rightarrow$ (i). Assuming (iii) we construct a countable atomic model of T. We employ a Henkin construction as in the proof of the omitting types theorem 10.1.3. Let $\{\langle \underline{c}^s_1, \ldots, \underline{c}^s_{n_s}\rangle : s \in \omega\}$ be an enumeration of all finite sequences of Henkin constants. Stages 0, $3s+1$, and $3s+2$ are as in the proof of theorem 10.1.3.

Stage $3s+3$: Put $n = n_s$ and let q be the set of $\psi \in F_n(T)$ such that $T_{3s+2} \cup \{\neg\psi(\underline{c}^s_1, \ldots, \underline{c}^s_n)\}$ is inconsistent. Let $p \supseteq q$ be a principal complete n-type over T. Let $\varphi \in F_n(T)$ be a generator of p. Put $T_{3s+3} = T_{3s+2} \cup \{\varphi(\underline{c}^s_1, \ldots, \underline{c}^s_n)\}$. Clearly T_{3s+3} is consistent.

At the end of the construction we obtain a countable atomic model of T. This is prime by theorem 10.2.7.

12.

Corollary (Vaught). Let T be a countable complete theory. Suppose $S_n(T)$ is countable for each $n \in \omega$. Then T has a prime model.

Proof. Assuming that $S_n(T)$ is countable, we shall show that $B_n(T)$ is atomic, i.e. every principal n-type q over T is included in a complete principal n-type over T. Suppose not, i.e. any principal n-type which includes q is incomplete. We use a splitting argument. For each finite sequence s of 0's and 1's we define a principal n-type q_s which includes q. Begin with $q_\emptyset = q$. Given q_s, we know that q_s is incomplete, so let $\varphi_s \in F_n(T)$ be a formula such that $q_s \cup \{\varphi_s\}$ and $q_s \cup \{\neg\varphi_s\}$ are both consistent. Let q_{s1} be the n-type generated by $q_s \cup \{\varphi_s\}$ and let q_{s0} be the n-type generated by $q_s \cup \{\neg\varphi_s\}$.

For each $f \in 2^\omega$ let $q_f = \bigcup\{q_{f\upharpoonright n} : n \in \omega\}$. Then q_f is an n-type over T, and $f \neq g$ implies $q_f \cup q_g$ is inconsistent. Let p_f be a complete n-type over T which includes q_f. Then $f \neq g$ implies $p_f \neq p_g$. Hence $|S_n(T)| = 2^{\aleph_0}$ contradicting our assumption that $S_n(T)$ is countable.

13.

Remark. In terms of Boolean algebras, the previous argument shows that if B is a countable Boolean algebra with $S(B)$ countable, then B is atomic.

10.3 The number of countable models

1.

Let T be a complete theory and let κ be a cardinal. We write $\nu(T, \kappa) =$ the number of nonisomorphic models of T of cardinality κ. This function is a central object of study in pure model theory.

In this section, we study $\nu(T, \aleph_0)$ where T is complete and countable. We begin with the case when $\nu(T, \aleph_0) = 1$, i.e. T is $\aleph_0$-categorical.

2.

Theorem (Ryll-Nardzewski). Let T be a complete countable theory with no finite models. The following are equivalent:

(1) T is $\aleph_0$-categorical.

(2) $S_n(T)$ is finite for each $n \in \omega$.

(3) Every countable model of T is prime.

(4) Every countable model of T is saturated.

(The equivalence of (1) and (2) is due to Ryll-Nardzewski.)

Proof. (1) $\Rightarrow$ (2). Suppose that $S_n(T)$ is infinite. We claim that there exists a nonprincipal $q \in S_n(T)$. To see this, for each principal $p \in S_n(T)$ let $\varphi_p \in F_n(T)$ be a generator of p. Put $Y = \{\neg\varphi_p : p \in S_n(T), p \text{ non-principal}\}$. Since $S_n(T)$ is infinite, it follows by compactness that Y is consistent over T. Hence Y can be extended to a complete n-type $q \in S_n(T)$. Clearly q is nonprincipal.

By lemma 10.2.5 there exists a countable model $\mathcal{A} \in \mathrm{Mod}(T)$ such that $\mathcal{A}$ omits q. Also there exists a countable $\mathcal{B} \in \mathrm{Mod}(T)$ such that $\mathcal{B}$ realizes q. $\mathcal{A}$ and $\mathcal{B}$ are nonisomorphic, so T is not $\aleph_0$-categorical.

(2) $\Rightarrow$ (3). Since $S_n(T)$ is finite, every $p \in S_n(T)$ is principal. Hence every model of T is atomic. Hence by theorem 10.2.7 it follows that every countable model of T is prime.

(2) $\Rightarrow$ (4). Let $\mathcal{A}$ be a countable model of T., and let $X \subseteq |\mathcal{A}|$ be finite. From (2) it follows that $S_1(X)$ is finite. hence each $p \in S_1(X)$ is principal, hence realized in $\mathcal{A}_X$. This shows that $\mathcal{A}$ is saturated.

(3) $\vee$ (4) $\Rightarrow$ (1). This is because of uniqueness of prime models (10.2.8) or saturated models (7.2.8) respectively.

3.

Remark. In terms of Boolean algebras, the argument for (1) $\Rightarrow$ (2) above shows that if B is a Boolean algebra with $S(B)$ infinite (equivalently B is infinite), then B has a nonprincipal ultrafilter.

4.

Next we consider the situation where $1 < \nu(T, \aleph_0) < \aleph_0$. We say that a countable model $\mathcal{A}$ of T is *weakly saturated* if every $p \in S_n(T)$, $n \in \omega$ is realized in $\mathcal{A}$.

5.

Theorem (Rosenstein). Suppose $1 < \nu(T, \aleph_0) < \aleph_0$ where T is a countable complete theory. Then T has a weakly saturated countable model which is not saturated.

Proof. Suppose not. Let $\mathcal{A}_1, \ldots, \mathcal{A}_k$ be the countable models of T which are not saturated. By assumption these models are not weakly saturated, so let $p_i \in S_{n_i}(T)$ be omitted in $\mathcal{A}_i$, $1 \leq i \leq k$.

Introduce new constant symbols $\underline{c}^i_1, \ldots, \underline{c}^i_{n_i}$, $1 \leq i \leq k$, and let T^* be a countable complete theory which includes $T \cup \{\varphi(\underline{c}^i_1, \ldots, \underline{c}^i_{n_i}) : \varphi \in p_i, 1 \leq i \leq k\}$. Let $\mathcal{A}^* = (|\mathcal{A}^*|, \Phi^*)$ be any countable model of T^*. Then $\mathcal{A} = (|\mathcal{A}^*|, \Phi^* \restriction \mathrm{sig}(T))$ is a countable model of T. Since each p_i is realized in $\mathcal{A}$, $1 \leq i \leq k$, it follows that $\mathcal{A}$ is saturated. Hence $\mathcal{A}^*$ is saturated. We have now shown that any countable model of T^* is saturated. Hence T^* is $\aleph_0$-categorical. Hence by Ryll-Nardzewski's theorem 10.3.2, it follows that $S_n(T^*)$ is finite for all $n \in \omega$. Hence $S_n(T)$ is finite for all $n \in \omega$. Hence by 10.3.2 again it follows that T is $\aleph_0$-categorical, i.e. $\nu(T, \aleph_0) = 1$, contradicting the hypothesis of the theorem.

6.

Theorem (Vaught). Let T be a countable complete theory. Then $\nu(T, \aleph_0) \neq 2$.

Proof. Suppose $1 < \nu(T, \aleph_0) < \aleph_0$. From $\nu(T, \aleph_0) \leq \aleph_0$ it follows that $S_n(T)$ is countable for each $n \in \omega$. Hence among the countable models of T are a prime one, a saturated one, and a weakly saturated one (corollary 10.2.12 and theorems 7.3.1 and 10.3.5). These three countable models are nonisomorphic (theorems 10.2.7 and 10.3.5).

7.

Examples. (1) We give examples of countable complete theories T_k with $\nu(T, \aleph_0) = k$, $1 \leq k < \aleph_0$. We may take T_1 to be the theory of dense linear ordering without endpoints ($\aleph_0$-categorical by 2.3.1). By the previous theorem, T_2 does not exist.

We may take T_3 to be the theory whose axioms are $T_1 \cup \{\underline{c}_n < \underline{c}_{n+1} : n \in \omega\}$. Thus a model of T_3 consists of a dense linear ordering without endpoints together with an increasing sequence of distinguished elements c_n, $n \in \omega$. In the three countable models, the c_n's are unbounded, bounded with no limit, or bounded with a limit. (These three models are respectively prime, saturated, and weakly saturated.)

For $m \geq 1$ we take axioms for T_{3+m} to consist of

$$\begin{aligned} T_3 \ \cup & \ \{\underline{P}_0(\underline{c}_n) : n \in \omega\} \\ \cup & \ \{\forall x\, (\underline{P}_i(x) \text{ for exactly one } i,\ 0 \leq i \leq m)\} \\ \cup & \ \{\forall x\, \forall y\, (x < y \rightarrow \exists z(\underline{P}_i(z) \wedge x < z < y)) : 0 \leq i \leq m\}\,. \end{aligned}$$

Thus a model of T_{3+m} is a model of T_3 in which the points are colored with $m+1$ densely shuffled colors P_i, $0 \leq i \leq m$, and the distinguished points c_n have color P_0. In a given model of T_{3+m} we may have c_n's unbounded, bounded with no limit, or bounded with a limit in P_i, $0 \leq i \leq m$. This gives $m+3$ distinct countable models. (All except the prime one are weakly saturated.)

(2) There are plenty of examples of countable complete theories T with $\nu(T, \aleph_0) = \aleph_0$ or $2^{\aleph_0}$. For instance $\nu(\mathsf{ACF}(0), \aleph_0) = \aleph_0$ and $\nu(\mathsf{RCOF}, \aleph_0) = 2^{\aleph_0}$.

8.

We omit the proof of the following theorem.

Theorem (Morley). Let T be a countable complete theory. If $\nu(T, \aleph_0)$ is infinite, then it is equal to either $\aleph_0$, $\aleph_1$, or $2^{\aleph_0}$.

9.

Remark. Vaught has conjectured that for all countable complete theories T, if $\nu(T, \aleph_0) > \aleph_0$ then $\nu(T, \aleph_0) = 2^{\aleph_0}$. No counterexample to Vaught's conjecture is known. Steel[1] has shown that Vaught's conjecture holds for theories T whose models are trees. (A *tree* is a partially ordered set such that the predecessors of any element are linearly ordered.) Shelah has shown that Vaught's conjecture holds for superstable theories. Also, Lachlan[2] has shown that for superstable theories, $\nu(T, \aleph_0) < \aleph_0$ implies $\nu(T, \aleph_0) = 1$.

10.4 Decidable prime models

1.

Let $\mathcal{A}$ be a countable model such that $\mathrm{sig}(\mathcal{A})$ is finite (or recursive). Recall that $\mathcal{A}$ is said to be *decidable* if there exists an enumeration $|\mathcal{A}| = \{a_n : 1 \leq n < \omega\}$ of the elements of $\mathcal{A}$, such that

$$\left\{\varphi \in \bigcup_{n\in\omega} F_n(T) : \mathcal{A} \models \varphi(a_1, \ldots, a_n)\right\}$$

is recursive. (This is easily seen to be equivalent to definition 4.3.1.)

2.

Recall theorem 4.3.4 which said that if T is a decidable theory, then T has a decidable model. However, it is not in general true that if T is a complete decidable theory which has a prime model, then T has a decidable prime model. We now present a theorem which gives a necessary and sufficient condition for T to have a decidable prime model.

[1] SLNM 689
[2] FM 81

3.

Theorem (Harrington). Let T be a complete decidable theory. The following are equivalent.

(a) T has a decidable prime model.

(b) Given $n \geq 1$ and $\varphi \in F_n(T)$ such that φ is consistent with T, we can recursively enumerate a complete n-type $p^\varphi \in S_n(T)$, such that $\varphi \in p^\varphi$ and p^φ is principal.

Note. We do not assume in (b) that it is possible to pass recursively from φ to a generator of p^φ.

Proof. (a) $\Rightarrow$ (b). Let $\mathcal{A}$ be a decidable prime model of T. Given $\varphi \in F_n(T)$ consistent with T, since T is complete we have $T \models \exists v_1 \cdots v_n\, \varphi$ so we can recursively find an n-tuple $a_1, \ldots, a_n \in |\mathcal{A}|$ such that $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$. We can then enumerate the complete n-type realized by $a_1, \ldots, a_n$ and define p^φ to be this type. Clearly $\varphi \in p^\varphi$, and p^φ is principal by theorem 10.2.7.

(b) $\Rightarrow$ (a). Let T be a complete decidable theory such that (b) holds. Introduce a recursive set of new constant symbols $C = \{\underline{c}_n : 1 \leq n < \omega\}$ and let $T_0 = T \cup \{H_n(\underline{c}_1, \ldots, \underline{c}_n) : 1 \leq n < \omega\}$ where $H_n(\underline{c}_1, \ldots, \underline{c}_n)$ is the nth Henkin sentence $(\exists x\, \psi_n(x)) \rightarrow \psi_n(\underline{c}_n)$. Here we assume that $\{\psi_n(x) : 1 \leq n < \omega\}$ is an enumeration of all formulas $\psi(x)$ with one free variable x such that $\mathrm{sig}(\psi(x)) \subseteq \mathrm{sig}(T) \cup C$, and we assume that $\mathrm{sig}(\psi_n(x)) \subseteq \mathrm{sig}(T) \cup \{\underline{c}_1, \ldots, \underline{c}_n\}$.

We shall construct a complete decidable theory T_∞ extending T_0 such that, for each $m \geq 1$, there will exist a formula $\varphi_m \in F_m(T)$ such that $\underline{c}_1, \ldots, \underline{c}_m$ realizes p^{φ_m}. This clearly implies that the Henkin model associated with T (as in the proof of theorem 4.3.4) is decidable and prime.

If it were possible to pass recursively from φ to a generator of p^φ, we could construct T_∞ by straightforwardly combining the proofs of theorem 4.3.4 and 10.2.9. It would then be the case that φ_m is a recursive function of m. Unfortunately this is not possible. We shall have to define φ_m as the limit of a recursive sequence of approximations φ_m^s. We shall use a priority argument to show that the approximations converge.

We shall construct T_∞ recursively in stages $T_0 \subseteq T_1 \subseteq \cdots \subseteq T_s \subseteq \cdots$. At each stage $s \geq 0$, T_s will consist of T_0 plus finitely many sentences.

Let $\tau^s(\underline{c}_1, \ldots, \underline{c}_{N_s})$ be the conjunction of this finite set of sentences, where $\tau^s \in F_{N_s}(T)$. Put

$$\tau^s_m(v_1, \ldots, v_m) \equiv \exists v_{m+1} \cdots v_n \, [\, H_1(v_1) \wedge \cdots \wedge H_n(v_1, \ldots, v_n) \wedge \tau^s(v_1, \ldots, v_{N_s}) \,]$$

where $n = \max(m, N_s)$. Thus τ^s_m expresses all the assertions about $\underline{c}_1, \ldots, \underline{c}_m$ to which we are committed at the end of stage s. Note that $\tau^s_m \in F_m(T)$.

At each stage s, there will be a recursive sequence of formulas $\varphi^s_m \in F_m(T)$, $1 \leq m < \omega$. Here φ^s_m is the stage s approximation of the desired formula φ_m.

We now present the construction.

Stage 0. Let T_0 be as above. Put $\varphi^0_m(v_1, \ldots, v_m) = \tau^0_m(v_1, \ldots, v_m)$.

Stage $s+1$. Let $s = (n-1, k, t)$ in some fixed recursive enumeration of $\omega \times \omega \times \omega$. Let $\theta(v_1, \ldots, v_n)$ be the kth formula in some fixed recursive enumeration of $F_n(T)$. At this stage we shall consider whether to add $\theta(\underline{c}_1, \ldots, \underline{c}_n)$ to our theory.

If $\theta \in p^{\varphi^s_n}$ and $T \not\models \tau^s_n \to \theta$ and $T \models \tau^s_m \to \exists v_{m+1} \cdots v_n \, [\, \tau^s_n \wedge \theta \,]$ for each m, $1 \leq m < n$, then put

$$T_{s+1} \;=\; T_s \cup \{\theta(\underline{c}_1, \ldots, \underline{c}_n)\}$$

and, for all $m \geq 1$,

$$\varphi^{s+1}_m \;=\; \begin{cases} \varphi^s_m & \text{if } 1 \leq m \leq n\,, \\ \tau^{s+1}_m & \text{if } m > n\,. \end{cases}$$

Otherwise do nothing, i.e. $T_{s+1} = T_s$ and $\varphi^{s+1}_m = \varphi^s_m$ for all $m \geq 1$.

It is clear that the construction is recursive and hence produces a recursively axiomatizable theory T_∞. We shall now prove a sequence of claims leading to the conclusion that this theory is complete and gives rise to an atomic Henkin model.

Claim 1: $T \models \tau^s_m \to \varphi^s_m$.

This is clear by induction on s since $\varphi^0_m = \tau^0_m$ and $\varphi^{s+1}_m =$ either φ^s_m or τ^{s+1}_m.

Claim 2: $\tau^s_m \in p^{\varphi^s_m}$.

We prove this by induction on s. For $s = 0$ it is trivial since $\tau^0_m = \varphi^0_m \in p^{\varphi^0_m}$. If no action is taken at stage $s+1$, the induction step is trivial. Consider a stage $s+1$ at which action is taken. Let n and θ be as in the construction. For $m < n$ we have $\tau^{s+1}_m \equiv \exists v_1 \cdots v_n \, [\, \tau^s_n \wedge \theta \,]$ so by hypothesis

$T \models \tau_m^s \to \tau_m^{s+1}$. Also $\varphi_m^{s+1} = \varphi_m^s$ so by induction $\tau_m^{s+1} \in p^{\varphi_m^s} = p^{\varphi_m^{s+1}}$. For $m = n$, we have $\tau_n^{s+1} = \tau_n^s \wedge \theta$ and $\tau_n^s \in p^{\varphi_n^s}$ by induction, and $\theta \in p^{\varphi_n^s}$ by hypothesis. Hence $\tau_n^{s+1} \in p^{\varphi_n^s} = p^{\varphi_n^{s+1}}$ since $\varphi_n^{s+1} = \varphi_n^s$. Finally, for $m > n$, we have $\tau_m^{s+1} = \varphi_m^{s+1} \in p^{\varphi_m^{s+1}}$.

Claim 3: $\forall n \,\exists s \,\forall s' \geq s \,\varphi_n^{s'} = \varphi_n^s$. (We then define $\varphi_n = \lim_s \varphi_n^s$.)

Claim 4: $\forall n \,\exists s \,(\tau_n^s$ is a generator of $p^{\varphi_n})$.

We prove claims 3 and 4 simultaneously by induction on n. Assume that 3 nd 4 hold for all $m < n$. Let σ be a stage such that for all $m < n$ and all $s \geq \sigma$, $\varphi_m^s = \varphi_m^\sigma$ and τ_m^s generates $p^{\varphi_m^s}$.

We first prove 3 for n. Given $s \geq \sigma$, let $s = (m-1, k, t)$ and let θ be the kth formula in $F_m(T)$. If $m \geq n$, then $\varphi_n^{s+1} = \varphi_n^s$. If $m < n$, then $\theta \in p^{\varphi_m^s}$ implies $T \models \tau_m^s \to \theta$, so nothing was done at stage s. Hence again $\varphi_n^{s+1} = \varphi_n^s$. Thus by induction we have $\forall s \geq \sigma \,\varphi_n^s = \varphi_n^\sigma$.

Next we prove claim 4 for n. Let θ be a generator of $p^{\varphi_n} = p^{\varphi_n^\sigma}$. Look at what happens at a stage $s \geq \sigma$ such that $s = (n-1, k, t)$ where θ is the kth formula in $F_n(T)$. We have $\theta \in p^{\varphi_n^s}$ and, by claim 2, $\tau_n^s \in p^{\varphi_n^s}$. Hence $\tau_n^s \wedge \theta$ is consistent with T. Hence, for each $m < n$, since τ_m^s generates $p^{\varphi_m^s}$, it follows that τ_m^s is an atom in $B_m(T)$, so

$$T \quad \models \quad \tau_m^s \to \exists v_{m+1} \cdots v_n \,[\, \tau_n^s \wedge \theta \,]\,.$$

Thus, unless $T \models \tau_n^s \to \theta$ already, we have $T_{s+1} = T_s \cup \{\theta(\underline{c}_1, \ldots, \underline{c}_n)\}$. In any case $T \models \tau_n^{s+1} \to \theta$ so τ_n^{s+1} is a generator of p^{φ_n}. This completes the proof of claims 3 and 4.

From claim 4 it follows that $T_\infty = \bigcup_{s \in \omega} T_s$ generates a complete decidable extension of T_0 whose asociated Henkin model is atomic. This completes the proof of the theorem.

4.

Remark. In chapter 11 we shall apply the above theorem to show that the differential closure of a computable differential field of characteristic 0 is computable. This was Harrington's original application.

5.

Exercise. Prove the following theorem of Morley[3]:

[3]Israel J. Math. vol. 25 pp. 233-240

Theorem (Morley). A complete decidable theory T has a decidable saturated model if and only if there is a uniform recursive enumeration of $\bigcup_{n\in\omega} S_n(T)$.

For more information on decidable models, see e.g. Millar, Annals of Math. Logic vol. 13, pp. 45-72.

6.

Exercise. Use theorem 10.4.3 to prove that the algebraic closure of a computable field is computable, and the real closure of a computable ordered field is computable.

Chapter 11

Differentially closed fields of characteristic 0

11.1 Simple extensions

1.

The theory of differential fields, DF, has the following set of axioms:

(a) field axioms

(b) $\forall x \forall y \ (x+y)' = x' + y'$

(c) $\forall x \forall y \ (x \cdot y)' = x' \cdot y + x \cdot y'$

Here $'$ is a 1-place operation symbol.

A *differential field* is a model of these axioms. A *differential domain* is a model of the axioms for domains plus (b), (c). A *differential ring* is defined similarly. The *characteristic* of a diffrential field is the characteristic of the underlying field.

2.

Examples. (1) Let $\mathcal{A}$ be any field. The polynomial ring $\mathcal{A}[x]$ can be made into a differential ring by interpreting $'$ as the formal derivative, i.e. if $f = a_n x^n + \cdots + a_1 x + a_0$, $a_i \in |\mathcal{A}|$, then $f' = n a_n x^{n-1} + \cdots + a_1$.

We can extend the derivative to the field of rational functions $\mathcal{A}(x)$ = field of quotients of $\mathcal{A}[x]$ in the usual way:

$$\left(\frac{f}{g}\right)' = \frac{f' \cdot g - f \cdot g'}{g^2}.$$

Then $\mathcal{A}(x)$ becomes a differential field.

(2) Any field $\mathcal{A}$ can be made into a differential field in a trivial way by putting $a' = 0$ for all $a \in |\mathcal{A}|$.

Note that this is the *only* way to make the rational field $\mathcal{Q} = (\mathbb{Q}, +, -, \cdot, 0, 1)$ into a differential field. Thus the differential field $\mathcal{Q}$ is embedded in every differential field of characteristic 0.

(3) Let D be a connected open set in the complex z-plane (or more generally any Riemann surface). Let $\mathcal{M}$ be the field of meromorphic functions on D, i.e. the field of quotients of the ring of holomorphic functions on D. This can be regarded as a differential field in the obvious way: f' = derivative of $f = df/dz$.

(4) Let $\mathcal{N}$ be the set of all functions which are meromorphic in some open set containing 0 in the z-plane. (Identify two such functions if they are equal on some open set containing 0.) This is again a differential field in an obvious way.

3.

Remark. Ritt and Seidenberg[1] have shown that every differential field which is finitely generated over the reational field is isomorphic to a differential subfield of $\mathcal{N}$. Thus everything we shall do in this chapter is meaningful for analysis.

4.

We now study simple extensions of differential fields of characteristic 0.

Lemma. Let $\mathcal{D}$ be a differential domain and let $\mathcal{F}$ be its field of quotients, $\mathcal{F} = \{c/d : c, d \in |\mathcal{D}|, d \neq 0\}$. Then there is one and only one way to make $\mathcal{F}$ into a differential field extension of $\mathcal{D}$.

[1]Proc. AMS 9 (1958) 159-164; 23 (1969) 689-691.

Proof. We must have

$$c' = \left(\frac{c}{d} \cdot d\right)' = \left(\frac{c}{d}\right)' \cdot d + \frac{c}{d} \cdot d'$$

so

$$\left(\frac{c}{d}\right)' = \frac{c' \cdot d - c \cdot d'}{d^2}\,.$$

It is easy to check that this makes $\mathcal{F}$ a differential field, etc.

5.

Lemma. Let $\mathcal{A}$ be a differential field of characteristic 0. Let $\mathcal{B} = \mathcal{A}(b_0, \ldots, b_{n-1}, b_n)$ be a finitely generated field extension of $\mathcal{A}$ such that

(i) for each $i < n$, b_i is transcendental over $\mathcal{A}(b_0, \ldots, b_{i-1})$,

(ii) b_n is algebraic over $\mathcal{A}(b_0, \ldots, b_{n-1})$.

Then there is one and only one way to make $\mathcal{B}$ into a differential field extension of $\mathcal{A}$ subject to the conditions $b_i' = b_{i+1}$, $i < n$.

6.

Before proving this lemma, we give an example. Let b_0 be transcendental over $\mathcal{A}$ and let b_1 satisfy $b_1^2 - b_0 = 0$. Then $\mathcal{A}(b_0, b_1)$ can be made into a differential field extension of $\mathcal{A}$ in which b_0 is a solution of the first order differential equation $(y')^2 - y = 0$ and is not the solution of any algebraic equation over $\mathcal{A}$.

7.

Proof of lemma 5. Consider the ring $\mathcal{A}[\vec{y}] = \mathcal{A}[y_0, \ldots, y_{n-1}, y_n]$. A typical element g of $\mathcal{A}[y_0, \ldots, y_{n-1}, y_n]$ looks like

$$g(\vec{y}) = \sum_j a_j \vec{y}^j$$

where $a_j = a_{j_0 j_1 \ldots j_n} \in |\mathcal{A}|$ and $\vec{y}^j = y_0^{j_0} \cdots y_{n-1}^{j_{n-1}} y_n^{j_n}$. Let us write

$$\widetilde{g}(\vec{y}) = \sum_j a_j' \vec{y}^j$$

and $\partial g/\partial y_i$ = the formal partial derivative of $g(y_0,\ldots,y_{n-1},y_n)$ with respect to y_i. Given $g(b_0,\ldots,b_{n-1},b_n) \in \mathcal{A}[b_0,\ldots,b_{n-1},b_n]$ we are forced to define

$$g(b_0,\ldots,b_{n-1},b_n)' = \widetilde{g}(b_0,\ldots,b_{n-1},b_n) + \sum_{i=0}^{n} \frac{\partial g}{\partial y_i}(b_0,\ldots,b_{n-1},b_n) \cdot b_{i+1}$$

where however b_{n+1} remains to be determined.

Let $f(b_0,b_1,\ldots,b_{n-1},y_n) \in \mathcal{A}(b_0,\ldots,b_{n-1})[y_n]$ be irreducible such that $\mathcal{B} \cong \mathcal{A}(b_0,\ldots,b_{n-1})[y_n]/f$ (by theorem 6.1.12). Multiplying f by a suitable element of $\mathcal{A}[b_0,\ldots,b_{n-1}]$, we may safely assume that $f(y_0,\ldots,y_{n-1},y_n) \in \mathcal{A}[y_0,\ldots,y_{n-1},y_n]$. As in the previous paragraph, we want

$$0 = f(b_0,b_1,\ldots,b_{n-1},b_n)' = \widetilde{f}(b_0,\ldots,b_{n-1},b_n) + \sum_{i=0}^{n} \frac{\partial f}{\partial y_i}(b_0,\ldots,b_{n-1},b_n) \cdot b_{i+1}\,.$$

Since $\partial f/\partial y_n$ is of lower degree than f in y_n, we must have

$$\frac{\partial f}{\partial y_n}(b_0,\ldots,b_{n-1},b_n) \neq 0\,.$$

Hence we are forced to define

$$b_{n+1} = \frac{\widetilde{f}(b_0,\ldots,b_{n-1},b_n) + \sum_{i=0}^{n-1} \frac{\partial f}{\partial y_i}(b_0,\ldots,b_{n-1},b_n) \cdot b_{i+1}}{-\frac{\partial f}{\partial y_n}(b_0,\ldots,b_{n-1},b_n)}\,.$$

We must check that $'$ is well defined. We have $f(b_0,\ldots,b_{n-1},b_n) = 0$. Also $f(b_0,\ldots,b_{n-1},b_n)' = 0$ by choice of b_{n+1}. Suppose $g_1, g_2 \in \mathcal{A}[y_0,\ldots,y_{n-1},y_n]$ are such that $g_1(b_0,\ldots,b_{n-1},b_n) = g_2(b_0,\ldots,b_{n-1},b_n)$. Then

$$g_1(b_0,\ldots,b_{n-1},y_n) \equiv g_2(b_0,\ldots,b_{n-1},y_n) \mod f(b_0,\ldots,b_{n-1},y_n)$$

in the polynomial ring $\mathcal{A}(b_0,\ldots,b_{n-1})[y_n]$. Hence $g_1 - g_2 = f \cdot h$ where $h \in \mathcal{A}(y_0,\ldots,y_{n-1})[y_n]$. Letting $h = p/q$, we have

$$q \cdot (g_1 - g_2) = f \cdot p$$

where $p \in \mathcal{A}[y_0,\ldots,y_{n-1},y_n]$ $q \in \mathcal{A}[y_0,\ldots,y_{n-1}]$, $q \not\equiv 0$. Substituting b_i for y_i and differentiating both sides we have

$$\begin{aligned} q(\vec{b}) \cdot (g_1(\vec{b})' - g_2(\vec{b})') &= f(\vec{b})' \cdot p(\vec{b}) + f(\vec{b} \cdot p(\vec{b})' \\ &= 0 \end{aligned}$$

and $q(\vec{b}) \neq 0$, hence $g_1(b_0, \ldots, b_{n-1}, b_n)' = g_2(b_0, \ldots, b_{n-1}, b_n)'$. This shows that $'$ is well defined on the ring $\mathcal{A}[b_0, \ldots, b_{n-1}, b_n]$. It is easy to check the sum and product rules, so we now know that there is one and only one way to make $\mathcal{A}[b_0, \ldots, b_{n-1}, b_n]$ into a differential ring extension of $\mathcal{A}$. By lemma 11.1.4 we get the same conclusion for the field of quotients $\mathcal{B} = \mathcal{A}(b_0, \ldots, b_{n-1}, b_n)$. This completes the proof.

8.

Remark. Let $\mathcal{A}$ be a differential field. We have considered four different kinds of "simple extensions" of $\mathcal{A}$, namely:

$$\mathcal{A}[b] = \text{ring generated by } |\mathcal{A}| \cup \{b\},$$

$$\mathcal{A}(b) = \text{field generated by } |\mathcal{A}| \cup \{b\},$$

$$\mathcal{A}\{b\} = \text{differential ring generated by } |\mathcal{A}| \cup \{b\},$$

$$\mathcal{A}\langle b\rangle = \text{differential field generated by } |\mathcal{A}| \cup \{b\}.$$

Our goal is to classify the simple differential field extensions $\mathcal{A}\langle b\rangle$.

9.

Definition. Let $\mathcal{A}$ be a differential field. A *differential polynomial* $f(y) = f(y, y', \ldots, y^{(n)})$ is an element of the differential ring

$$\mathcal{A}\{y\} = \mathcal{A}[y, y', \ldots, y^{(n)}, \ldots]$$

where of course $y^{(n)} = y\overbrace{{}''\cdots{}'}^{}_{n\,\text{times}}$ and y is a differential indeterminate. (Alternatively, we could define a differential polynomial to be an equivalence class of terms $t(y)$ in the language of the diagram of $\mathcal{A}$, as in 6.1.4.)

Note that $\mathcal{A}\{y\}$ is a differential domain and $\mathcal{A}\langle y\rangle$, its field of quotients, is a differential field by lemma 11.1.4.

10.

Definition. If $f(y) \in \mathcal{A}\{y\}$ is a nonzero differential polynomial over $\mathcal{A}$, we define the *order* of f to be the largest n such that $y^{(n)}$ occurs in f, and the *degree* of f to be the largest k such that $(y^{(n)})^k$ occurs in f, where n is the order of f.

For example, the differential polynomial $(y^{(3)})^4 + (y')^5 \cdot y^{(3)} - y$ has order 3 and degree 4.

11.

We now survey all simple differential field extensions $\mathcal{A}\langle b\rangle$ of a differential field $\mathcal{A}$ of characteristic 0.

Case 1: $f(b) = 0$ for some nonzero $f \in \mathcal{A}\{y\}$. Choose such an f of lowest possible order n, and then of lowest possible degree k for that order. Note that $\partial f/\partial y^{(n)}$ has lower order or else the same order and lower degree. Hence

$$\frac{\partial f}{\partial y^{(n)}}(b) \neq 0\,.$$

Hence, as in 11.1.7, we have

$$b^{(n+1)} = \frac{\widetilde{f}(b) + \sum_{i=0}^{n-1} \frac{\partial f}{\partial y^{(i)}}(b) \cdot b^{(i+1)}}{-\frac{\partial f}{\partial y^{(n)}}(b)}$$

so in particular $b^{(n+1)} \in \mathcal{A}(b, \ldots, b^{(n-1)}, b^{(n)})$. From this it follows that $b^{(j)} \in \mathcal{A}(b, \ldots, b^{(n-1)}, b^{(n)})$ for all j, and hence the field $\mathcal{A}(b, \ldots, b^{(n-1)}, b^{(n)})$ is closed under $'$. Hence $\mathcal{A}\langle b\rangle = \mathcal{A}(b, \ldots, b^{(n-1)}, b^{(n)})$ and its structure is as in lemma 11.1.5. In this case we say that b is *differentially algebraic* over $\mathcal{A}$.

Case 2: negation of case 1. In this case it is easy to see that $\mathcal{A}\{b\} \cong \mathcal{A}\{y\}$, hence the quotient fields $\mathcal{A}\langle b\rangle \cong \mathcal{A}\langle y\rangle$ by lemma 11.1.4. In this case we say that b is *differentially transcendental* over $\mathcal{A}$.

This completes our survey of the simple differential field extensions of $\mathcal{A}$.

12.

From the above survey of simple extensions we immediately obtain the following result.

Theorem. Let $\mathcal{A}$ be a differential field of characteristic 0. Up to isomorphism over $\mathcal{A}$, there are exactly $\|\mathcal{A}\|$ distinct simple differential field extensions $\mathcal{A}\langle b\rangle$.

11.2 Differentially closed fields

1.

Definition (Blum). Let $\mathcal{B}$ be a differential field of characteristic 0. We say that $\mathcal{B}$ is *differentially closed* if $\mathcal{B}$ is algebraically closed and, for all $f, g \in \mathcal{B}\{y\}$ with $0 \leq \text{order}(g) < \text{order}(f)$, there exists $b \in |\mathcal{B}|$ with $f(b) = 0$, $g(b) \neq 0$.

2.

Lemma. Let $\mathcal{A}$ be a differential field of characteristic 0. Then there exists a differential field $\mathcal{B} \supseteq \mathcal{A}$ such that (i) $\mathcal{B}$ is differentially closed, (ii) $\mathcal{B}$ is differentially algebraic over $\mathcal{A}$, i.e. every $b \in |\mathcal{B}|$ is differentially algebraic over $\mathcal{A}$.

Proof. By transfinite induction, it suffices to show:

3.

Sublemma. Given $f \in \mathcal{A}\{y\}$ of order > 0, there exists a simple differential field extension $\mathcal{A}\langle b \rangle$ such that $f(b) = 0$ and $g(b) \neq 0$ for *all* $g \in \mathcal{A}\{y\}$ with $0 \leq \text{order}(g) < \text{order}(f)$.

(In this event b is called a *generic solution* of the differential equation $f(y) = 0$ over $\mathcal{A}$.)

To prove the sublemma, let n be the order of f and let $\mathcal{A}^* = \mathcal{A}(b_0, \ldots, b_{n-1})$ be a field extension of $\mathcal{A}$ in which each b_i, $i < n$, is transcendental over $\mathcal{A}(b_0, \ldots, b_{i-1})$. Define a polynomial $p(x) \in \mathcal{A}^*[x]$ by $p(x) = f(b_0, \ldots, b_{n-1}, x)$ where $f = f(y, \ldots, y^{(n-1)}, y^{(n)}) \in \mathcal{A}[y, \ldots, y^{(n-1)}, y^{(n)}]$. Let $\mathcal{A}^*(c)$ be a simple algebraic field extension of $\mathcal{A}^*$ such that $p(c) = 0$. Put $b_n = c$. By lemma 11.1.5 we can make $\mathcal{A}^*(c) = \mathcal{A}(b_0, \ldots, b_{n-1}, b_n)$ into a simple differential field extension $\mathcal{A}\langle b \rangle$ of $\mathcal{A}$ where $b = b_0$, $b_i' = b_{i+1}$ for $i < n$. We then have $f(b) = f(b_0, \ldots, b_{n-1}, b_n) = p(b_n) = p(c) = 0$ and by construction $g(b_0, \ldots, b_{n-1}) \neq 0$ for all nonzero $g \in \mathcal{A}[y_0, \ldots, y_{n-1}]$, in other words $g(b) \neq 0$ for all $g \in \mathcal{A}\{y\}$ of order m, $0 \leq m < n$. This completes the proof.

4.

Remark. In the statement of the lemma $\mathcal{B}$ is not unique. See 11.3.4 below.

5.

Theorem (Robinson, Blum). The theory of differentially closed fields of characteristic 0 ($\mathsf{DCF}(0)$) is the model completion of the theory of differential fields of characteristic 0 ($\mathsf{DF}(0)$).

Proof. By the previous lemma plus theorem 8.3.2 it suffices to show the following:

> Let $\mathcal{A} \subseteq \mathcal{B}$ where $\mathcal{A}$ is a $\mathsf{DF}(0)$, $\mathcal{B}$ a $\mathsf{DCF}(0)$, $\mathcal{B}$ κ^+-saturated where $\kappa = \|\mathcal{A}\|$. Then for any simple differential field extension $\mathcal{A}\langle c\rangle$ of $\mathcal{A}$, we can find $b \in |\mathcal{B}|$ such that $\mathcal{A}\langle b\rangle \cong \mathcal{A}\langle c\rangle$ over $\mathcal{A}$.

To prove this, we consider two cases.

Case 1: c is differential algebraic over $\mathcal{A}$, i.e. $f(c) = 0$ for some nonzero $f \in \mathcal{A}\{y\}$. Take f to be of least order, and of least degree for that order. Let p be the 1-type over $\mathcal{A}$ generated by the differential equation $f(y) = 0$ and the inequations $g(y) \neq 0$, $g \in \mathcal{A}\{y\}$, $0 \leq \text{order}(g) < \text{order}(f)$. Since $\mathcal{B}$ is differentially closed, every finite subset of p is realized in $\mathcal{B}$. Hence by saturation p is realized in $\mathcal{B}$, say by $b \in |\mathcal{B}|$. We claim that $\mathcal{A}\langle b\rangle \cong \mathcal{A}\langle c\rangle$ over $\mathcal{A}$. This is clear from 11.1.5 and 11.1.11.

Case 2: c is differentially transcendental over $\mathcal{A}$. Consider the 1-type p over $\mathcal{A}$ generated by $g(y) \neq 0$ for all $g \in \mathcal{A}\{y\}$, $0 \leq \text{order}(g)$. Again, since $\mathcal{B}$ is differentially closed, each finite subset of p is realized in $\mathcal{B}$. Hence p is realized in $\mathcal{B}$, say by $b \in |\mathcal{B}|$. Then b is differentially transcendental over $\mathcal{A}$, so $\mathcal{A}\langle b\rangle \cong \mathcal{A}\langle c\rangle$. This completes the proof.

6.

Corollary. $\mathsf{DCF}(0)$ is complete and decidable.

Proof. By the previous theorem $\mathsf{DCF}(0) \cup$ (diagram of $\mathbb{Q}$) is complete, but any $\mathsf{DCF}(0)$ has $\mathbb{Q}$ as a uniquely embedded differential subfield, so $\mathsf{DCF}(0)$ is complete. Decidability follows by theorem 4.2.3.

7.

Corollary (Seidenberg). Let S be a system of finitely many algebraic differential equations in n unknown functions $y_1, \ldots, y_n$ and their derivatives, with rational coefficients. Then we can recursively decide whether S has a solution in some differential field of characteristic 0 (equivalently by 11.1.3 in some field of meromorphic functions).

8.

Corollary. $\mathsf{DCF}(0)$ admits elimination of quantifiers.

Proof. From theorem 11.2.5 and lemma 11.1.4 it follows that $\mathsf{DCF}(0)$ is the model completion of the theory of differential domains of characteristic 0. The latter theory is universal so quantifier elimination follows by theorem 8.2.3.

9.

Corollary (Seidenberg). There exists an elimination theory for finite systems of algebraic differential equations and inequations (cf. 8.2.6).

11.3 Differential closure (countable case)

1.

Definition. Let $\mathcal{A}$ be a $\mathsf{DF}(0)$. A *differential closure* of $\mathcal{A}$ is a differential field $\mathcal{B} \supseteq \mathcal{A}$ such that (i) $\mathcal{B}$ is differentially closed, and (ii) if $\mathcal{B}_1$ is any differentially closed field containing $\mathcal{A}$, then $\mathcal{B}$ is embeddable into $\mathcal{B}_1$ over $\mathcal{A}$.

In view of theorem 11.2.5 this is equivalent to saying that $\mathcal{B}$ is a prime model of the complete theory generated by $\mathsf{DCF}(0) \cup$ (diagram of $\mathcal{A}$).

2.

Theorem (Blum). Let $\mathcal{A}$ be a countable $\mathsf{DF}(0)$. Then $\mathcal{A}$ has a differential closure, denoted $\overline{\mathcal{A}}$. Furthermore $\overline{\mathcal{A}}$ is unique up to isomorphism over $\mathcal{A}$.

(We shall see in chapters 12 and 13 that the hypothesis of countability can be dropped.)

Proof. Let T be the complete countable theory generated by $\mathsf{DCF}(0) \cup$ (diagram of $\mathcal{A}$). By quantifier elimination, a complete n-type over T is essentially the same thing as an isomorphism type of an n-fold simple differential field extension $\mathcal{A}\langle b_1, \ldots, b_n\rangle$ over $\mathcal{A}$. Hence by theorem 11.1.12 it follows that $S_n(T)$ is countable for all n. Hence by Vaught's theorem 10.2.12 it follows that T has a prime model $\mathcal{B} = \overline{\mathcal{A}}$. Furthermore, the uniqueness theorem for prime models 10.2.8 implies that $\overline{\mathcal{A}}$ is unique up to isomorphism over $\mathcal{A}$.

3.

Remark. By lemma 11.2.2 it follows that the differential closure $\overline{\mathcal{A}}$ is differentially algebraic over $\mathcal{A}$. However, this fact alone does not suffice to characterize $\overline{\mathcal{A}}$ among all differentially closed fields $\mathcal{B} \supseteq \mathcal{A}$.

For example, let $\mathcal{A}\langle b\rangle$ be a simple extension in which b is a generic solution of $y' = 0$. In other words, b is transcendental over $\mathcal{A}$ but $b' = 0$. Put $\mathcal{B} = \overline{\mathcal{A}\langle b\rangle}$. Then $\mathcal{B}$ is differentially closed, and differentially algebraic over $\mathcal{A}$, but we claim that $\mathcal{B}$ is not isomorphic to $\overline{\mathcal{A}}$ over $\mathcal{A}$. To see this, let p be the complete 1-type over $T = \mathsf{DCF}(0) \cup$ (diagram of $\mathcal{A}$) realized by $b \in |\mathcal{B}|$. Clearly p is nonprincipal. Hence by theorem 10.2.7 it follows that p is not realized in $\overline{\mathcal{A}}$.

4.

In order to prove our next theorem, we need the following algebraic result which we state without proof.

Lemma. Let $\mathcal{A}$ be a differential field of characteristic 0. Let S be an infinite set of algebraic differential equations in finitely many unknowns $y_1, \ldots, y_n$ with coefficients in $|\mathcal{A}|$. Then there exists a finite subset S_0 of S such that every solution of S_0 (in some differential field extension of $\mathcal{A}$) is a solution of S.

This is an immediate corollary of the Ritt basis theorem, i.e. theorem 11.4.4 below. For a proof, see Kaplansky, *Introduction to Differential Algebra.*

5.

Theorem (Harrington). Let $\mathcal{A}$ be a computable differential field of characteristic 0. Then $\overline{\mathcal{A}}$, the differential closure of $\mathcal{A}$, is computable.

Proof. $\overline{\mathcal{A}}$ is by definition the prime model of the complete theory T whose axioms are $\mathsf{DCF}(0) \cup$ (diagram of $\mathcal{A}$). Since $\mathcal{A}$ is computable, T is recursively axiomatizable, and hence decidable. But theorem 10.4.3 gives a necessary and sufficient condition for a complete decidable theory to have a decidable prime model. Thus, in order to show that $\overline{\mathcal{A}}$ is computable, it suffices to show that T verifies condition (b) of that theorem.

Given $\varphi \in F_n(T)$ consistent with T, by quantifier elimination we may safely assume that $\varphi(y_1, \dots, y_n)$ consists of finitely many algebraic diferential equations and inequations in n unknowns $y_1, \dots, y_n$ with coefficients in $|\mathcal{A}|$. Let $\{\psi_k(y_1, \dots, y_n) : k \in \omega\}$ be a recursive enumeration of all algebraic differential equations in n unknowns $y_1, \dots, y_n$ with coefficients in $|\mathcal{A}|$. Let p^φ be the n-type over T generated by $\{\varphi_k : k \in \omega\}$ where $\varphi_0 = \varphi$, $\varphi_{k+1} = \varphi_k \wedge \psi_k$ if this is consistent with T, $\varphi_{k+1} = \varphi_k$ otherwise. The passage from φ to p^φ is clearly recursive, and by quantifier elimination, p^φ is complete.

It remains to show that p^φ is principal. Let S be the set of all algebraic differential equations in p^φ. By lemma 11.3.4 there exists a finite set $S_0 \subseteq S$ such that S_0 is equivalent to S over T. Let $\psi_{k_1}, \dots, \psi_{k_m}$ be the elements of S_0. Then clearly $\varphi \wedge \psi_{k_1} \wedge \cdots \wedge \psi_{k_m}$ is a generator of p^φ. This completes the proof.

6.

Corollary. The differential closure of the rational field $\mathbb{Q}$ is computable.

7.

Remark. An explicit presentation of the differential closure of $\mathbb{Q}$ is lacking. It is unknown whether the following decision problem has a positive solution:

> Given a finite system of differential equations and inequations in n unknowns $y_1, \dots, y_n$ with rational coefficients, to decide whether the principal n-type generated by the system is complete.

In the absence of such a decision procedure, the only known construction of the differential closure of $\mathbb{Q}$ is by means of Harrington's priority construction in the proof of theorem 10.4.3.

11.4 Ritt's Nullstellensatz

1.

Let $\mathcal{A}$ be a differential field of characteristic 0. Let $\mathcal{R} = \mathcal{A}\{y_1, \ldots, y_n\}$ be the ring of differential polynomials in n indeterminates $y_1, \ldots, y_n$ over $\mathcal{A}$. Let $\overline{\mathcal{A}}$ be the differential closure of $\mathcal{A}$.

2.

Theorem (Ritt's Nullstellensatz). Let $f, g_1, \ldots, g_m \in \mathcal{R}$ be such that every common zero of $g_1, \ldots, g_m$ in $\overline{\mathcal{A}}$ is a zero of f. Then there exist positive integers k and l such that

$$f^k = \sum_{i=1}^{m} \sum_{j=1}^{l} p_{ij} g_i^{(j)}$$

where $p_{ij} \in \mathcal{R}$ and $g_i^{(j)}$ is the jth derivative of g_i.

Proof. Let I be the set of all $f \in \mathcal{R}$ for which the conclusion holds. Suppose $f \notin I$. Then we have

(i) $g, h \in I \Rightarrow g + h \in I$

(ii) $g \in I, h \in \mathcal{R} \Rightarrow g \cdot h \in I$

(iii) $0 \in I, 1 \notin I$

(iv) $g \in I \Rightarrow g' \in I$

i.e. I is a *differential ideal*. Here (i), (ii), (iii) are proved as in the proof of Hilbert's Nullstellensatz 6.4.2. For (iv), suppose $g \in I$, say $g^k \equiv 0 \mod g_1, \ldots, g_m$. Differentiate and multiply by $1/k$ to get

$$g^{k-1} g' \equiv 0 \mod g_1, \ldots, g_m \,.$$

This is the case $i = 1$ of the statement

$$g^{k-i}(g')^{2i-1} \equiv 0 \mod g_1, \ldots, g_m$$

which we prove by induction on i, $1 \le i \le k$. Differentiating and multiplying by g' we get

$$(k-i)g^{k-i-1}(g')^{2i+1} + \underbrace{(2i-1)g^{k-i}(g')^{2i-1}g''}_{\equiv 0} \equiv 0\,.$$

Multiplying by $1/(k-i)$ we get

$$g^{k-i-1}(g')^{2i+1} \equiv 0 \mod g_1, \ldots, g_m$$

so the induction step is proved. Finally for $i = k$ we get

$$(g')^{2k-1} \equiv 0 \mod g_1, \ldots, g_m$$

so (iv) is proved.

It is also clear that I is a *radical differential ideal*, i.e. I satisfies

(v) $g^k \in I \Rightarrow g \in I$,

and $f \notin I$. By Zorn's lemma let $J \supseteq I$ be a radical differential ideal such that $f \notin J$ and maximal with this property.

We claim that J is *prime*, i.e. $g \notin J, h \notin J$ imply $g \cdot h \notin J$. Suppose not. Note that $g \cdot h \in J$ implies $(g' \cdot h + g \cdot h') \cdot g' \cdot h \in J$ whence $(g' \cdot h)^2 \in J$, whence $g' \cdot h \in J$. Similarly $g^{(i)} \cdot h^{(j)} \in J$ for all $i, j \ge 0$. If $g \notin J$ then the radical differential ideal generated by J and g must contain f, say

$$f^k = s + \sum_i p_i g^{(i)}$$

where $s \in J$, $p_i \in \mathcal{R}$. Similarly if $h \notin J$ then

$$f^l = t + \sum_j q_j h^{(j)}$$

where $t \in J$, $q_j \in \mathcal{R}$. Hence

$$f^{k+l} = \underbrace{s \cdot f^l + t \cdot f^k + \sum_i \sum_j p_i q_j g^{(i)} h^{(j)}}_{\in J}$$

contradicting the fact that J is radical and $f \notin J$. This proves the claim.

Now let $\mathcal{R}_1 = \mathcal{R}/J$ = the quotient ring of $\mathcal{R}$ mod J. Clearly $\mathcal{R}_1$ is a differential domain. Let $\mathcal{B}$ be the field of quotients of $\mathcal{R}_1$. Then $\mathcal{B}$ is a differential field extension of $\mathcal{A}$. Furthermore, there exist $b_1, \ldots, b_n \in |\mathcal{B}|$ (corresponding to $y_1, \ldots, y_n \in |\mathcal{R}|$) such that for all $g \in \mathcal{R}$, $g(b_1, \ldots, b_n) = 0$ if and only if $g \in J$. In particular $f(b_1, \ldots, b_n) \neq 0$ while $g_1(b_1, \ldots, b_n) = \cdots = g_m(b_1, \ldots, b_n) = 0$. Hence by model completeness of $\mathsf{DCF}(0)$ we can find $a_1, \ldots, a_n \in |\overline{\mathcal{A}}|$ such that $f(a_1, \ldots, a_n) \neq 0$ and $g_1(a_1, \ldots, a_n) = \cdots = g_m(a_1, \ldots, a_n) = 0$. This completes the proof of the theorem.

3.

Remark. As in the case of Hilbert's Nullstellensatz, we can deduce a version of Ritt's Nullstellensatz with effective bounds. We leave to the reader the formulation and proof of this result. (Compare 6.4.3.)

4.

An interesting theorem related to Ritt's Nullstellensatz is the following.

Theorem (Ritt's basis theorem). Let $\mathcal{R}$ be as above and let I be a radical differential ideal in $\mathcal{R}$. Then I has a *finite basis*, i.e. there exists a finite set $g_1, \ldots, g_r \in I$ such that I is the radical differential ideal generated by $g_1, \ldots, g_r$, i.e. the set of all $f \in \mathcal{R}$ such that

$$f^k = \sum_{i=0}^{r} \sum_{j=0}^{l} p_{ij} g_i^{(j)}$$

for some $k, l \in \omega$ and $p_{ij} \in \mathcal{R}$.

Proof. See Kaplansky, *Introduction to Differential Algebra*.

5.

Remark. The Ritt Basis Theorem tends to clarify several aspects of the proof of the Ritt Nullstellensatz. In the first place, the Ritt Basis Theorem shows that the use of Zorn's lemma could have been eliminated. In the second place, since the radical differential ideal J has a finite basis $g_1, \ldots, g_r$, it follows by quantifier elimination that the complete n-type of $b_1, \ldots, b_n$ over $|\mathcal{A}|$ is generated by the differential equations $g_1(y_1, \ldots, y_n) = \cdots = g_r(y_1, \ldots, y_n) = 0$ and the single inequation $f(y_1, \ldots, y_n) \neq 0$. In particular this n-type is principal[2] and hence already realized in $\overline{\mathcal{A}}$. In other words, the differential field $\mathcal{B} = \mathcal{A}\langle b_1, \ldots, b_n\rangle$ is already embeddable into $\overline{\mathcal{A}}$. (We could have used Hilbert's Basis Theorem to make similar remarks about the proof of the Hilbert Nullstellensatz.)

[2]The algebraists express this fact by saying that J is *constrained*. See Wood, Israel J. Math. vol. 25, p. 331.

Chapter 12

Totally transcendental theories

12.1 Stability

1.

12.2 Rank of an element type

12.3 Indiscernibles

12.4 Existence of saturated models

Chapter 13

Prime models (uncountable case)

13.1 Strongly atomic models

1.

13.2 Normal sets

13.3 Uniqueness and characterization of prime models