CHAPTER 5
Model Theory

5.1 Introduction

Why should students of mathematics want to know something about model theory? Here is one answer: model theory helps one to understand what it takes to specify a mathematical structure uniquely.

Let us consider what is undoubtedly the most well-known problem of this sort, the problem of defining or characterizing the field of real numbers. (These are really two separate problems; we are more concerned with characterization.) This was a major concern in the eighteenth and nineteenth centuries. One of the main reasons it took nearly two centuries after the invention of calculus to put the subject on a firm theoretical foundation is that there was no satisfactory characterization of the reals.

So, how do we characterize the reals? The first part of the answer is that we expect the reals, unlike more "meager" structures such as the integers, to allow subtraction and division as well as addition and multiplication. That is, we expect $\mathbb{R}$ to be a field. We also believe that the reals correspond to points on a line; this requires $\mathbb{R}$ to be an ordered field.

The property of being an ordered field does not specify the structure of $\mathbb{R}$ uniquely, because $\mathbb{Q}$ is also an ordered field and doesn't even have the same cardinality as $\mathbb{R}$. What are some additional properties that we might require of $\mathbb{R}$ to set it apart from $\mathbb{Q}$? In terms of algebraic properties, we also expect $\mathbb{R}$ to be a **real-closed** ordered field, that is, an ordered field in which every nonnegative number has a square root and every polynomial of odd degree has a zero. Another nice way of characterizing real-closed ordered fields is as ordered fields satisfying the intermediate value theorem for polynomials.

The defining properties of a real-closed ordered field are easy enough to state in the standard first-order language of an ordered ring with unity (with symbols $+$, $\cdot$, $-$, 0, 1, and $<$), although the part about polynomials requires an infinite list of axioms. In Section 5.5 we will see that this theory is complete, and therefore becomes inconsistent if any more sentences are added to it. Does this perhaps suggest that the property of being a real-closed ordered field characterizes $\mathbb{R}$ uniquely? It turns out that there is no way this could be so. One of the most important results in model theory, the **Löwenheim–Skolem–Tarski (LST) theorem**, guarantees that if a countable set of first-order sentences is satisfied by some infinite structure, then there are structures of every infinite cardinality that satisfy these sentences. So no countable list of first-order properties can "pinpoint" $\mathbb{R}$.

Yet, we know that it is possible to characterize $\mathbb{R}$ uniquely. The usual way is to add to the ordered field axioms the **completeness property** of Dedekind, that every nonempty set of reals with an upper bound has a least upper bound. But the completeness property is second-order: it refers to arbitrary *sets* of real numbers as well as individual numbers. Thus the LST theorem no longer applies. It is worthwhile to realize that the completeness property of the reals is a fundamentally more complex statement than the properties that define real-closed ordered fields, and cannot be replaced by any list of first-order properties. This situation, in which important properties cannot be stated accurately in first-order logic, is reminiscent of the first two examples in Section 1.5.

Another good reason to know something about model theory is that it has many applications outside of foundations, notably in algebra. Model theory is the study of **structures** for first-order languages, and all of the objects studied in abstract algebra—groups, rings, fields, modules, vector spaces, etc.—are structures in this technical sense. Furthermore, several of the main techniques within contemporary model theory, such as the analysis of the **definable** sets of structures, generalize methods that appear in many guises throughout algebra. Thus, model theory creates a means for illuminating and generalizing a variety of concepts from different parts of abstract algebra, and also provides some powerful tools for obtaining results in algebra as well as other branches of mathematics. We will see examples of algebraic applications of model theory in sections 5.4 through 5.7.

For a more thorough treatment of model theory, see [Hod] or [CK]. The “bible” of the subject is [CK], while [Hod] is a somewhat lighter introduction. [Mar] and [HPS] are good sources for model theory and its applications to algebra and other branches of mathematics.

5.2 Basic concepts of model theory

In this chapter, P and Q always denote formulas and T denotes a theory, in some first-order language $\mathcal{L}$ with equality. Here is the most basic concept of first-order model theory:

Definition. A **structure** $\mathfrak{A}$ for a language $\mathcal{L}$ (**$\mathcal{L}$-structure** for short) consists of:

(1) a nonempty set A or $|\mathfrak{A}|$, called the **universe** of $\mathfrak{A}$, to be used as the domain of the variables of $\mathcal{L}$.

(2) a subset of A^n assigned to each n-ary relation symbol in $\mathcal{L}$.

(3) a function from A^n to A assigned to each n-ary function symbol in $\mathcal{L}$.

(4) an element of A assigned to each constant symbol in $\mathcal{L}$.

In this chapter, $\mathfrak{A}$ and $\mathfrak{B}$ always denote structures, while $\mathcal{C}$ denotes a class of structures. We must often refer to a class of structures because many important collections of structures are not sets. For example, the collection of all groups is a proper class, as is the collection of all finite groups.

Example 1. The simplest first-order language is the language of pure identity, whose only atomic formulas are equations between variables. A structure for this language is just a nonempty set.

Example 2. Let $\mathcal{L}$ be the first-order language with no function or constant symbols and a single binary relation symbol R. This language is appropriate for many theories, including orderings (reflexive and irreflexive, partial and total), equivalence relations, and set theory. Of course, we would use a different abbreviation for atomic formulas $R(v_i, v_j)$ in each of these theories (probably $v_i \leq v_j$ or $v_i < v_j$ for an ordering, $v_i \sim v_j$ or $v_i \equiv v_j$ for an equivalence relation, and $v_i \in v_j$ for set theory). An $\mathcal{L}$-structure is simply any pair (A, S), where $A \neq \emptyset$ and $S \subseteq A \times A$.

An $\mathcal{L}$-structure provides a realization or interpretation of every relation symbol, function symbol, and constant symbol in $\mathcal{L}$, as well as for every bound variable appearing in an $\mathcal{L}$-formula. So it should be possible to define what it means for a given $\mathcal{L}$-formula to be true or false in a given $\mathcal{L}$-structure, as long as we also provide values for the free variables of the formula. Accordingly:

Definition. An **assignment** in $\mathfrak{A}$ is a function $g : V \to A$, where V is the set of variables of $\mathcal{L}$.

If g is an assignment, we write g^i_x to denote the assignment whose value on v_i is x, and which is otherwise identical to g.

Lemma 5.1. *If g is an assignment in an $\mathcal{L}$-structure $\mathfrak{A}$, there is a unique function $\hat{g} : \mathcal{T} \to A$, where $\mathcal{T}$ is the set of all terms of $\mathcal{L}$, satisfying:*

(a) *If t is a variable of $\mathcal{L}$, then $\hat{g}(t) = g(t)$.*

(b) *If t is a constant symbol of $\mathcal{L}$, then $\hat{g}(t)$ is the element of A corresponding to that constant symbol.*

(c) *If t is a term of the form $f(t_1, t_2, \ldots, t_n)$ then $\hat{g}(t)$ is*

$$F(\hat{g}(t_1), \hat{g}(t_2), \ldots, \hat{g}(t_n)),$$

where F is the function in $\mathfrak{A}$ that corresponds to f.

We will refer to $\hat{g}(t)$ as the **interpretation** of the term t based on the assignment g. We omit the proof of this lemma, which is a straightforward induction on the number of function symbols appearing in the term t. In fact, it would be reasonable just to *define* $\hat{g}$ without mentioning that anything needs to be proved. We will take this approach in the next definition.

Example 3. Let $\mathcal{L}$ be the language of a ring, $\mathfrak{A}$ the ring of real numbers, t the term $v_2 \cdot (v_5 + v_1)$, and g any assignment in which $g(v_1) = 5$, $g(v_2) = 3$, and $g(v_5) = 12$. Then $\hat{g}(t) = 3(12 + 5) = 51$.

Definition. Let g be an assignment in an $\mathcal{L}$-structure $\mathfrak{A}$, and let P be any formula of $\mathcal{L}$. Then we can define what it means for P to be **true** in $\mathfrak{A}$ (or **satisfied** by $\mathfrak{A}$) under the assignment g, using a straightforward inductive definition on the structure of P:

(1) If P is an equation $t_1 = t_2$, then P is true in $\mathfrak{A}$ under g iff $\hat{g}(t_1) = \hat{g}(t_2)$.

(2) If P is of the form $R(t_1, t_2, \ldots, t_n)$, where R is an n-ary relation symbol of $\mathcal{L}$, then P is true in $\mathfrak{A}$ under g iff

$$(\hat{g}(t_1), \hat{g}(t_2), \ldots, \hat{g}(t_n)) \in \hat{R},$$

where $\hat{R}$ is the subset of A^n that corresponds to R.

(3) If P is of the form $\sim$ Q, then P is true in $\mathfrak{A}$ under g iff Q is *not* true in $\mathfrak{A}$ under g.

(4) If P is of the form $\mathrm{Q}_1 \wedge \mathrm{Q}_2$, then P is true in $\mathfrak{A}$ under g iff Q_1 and Q_2 are both true in $\mathfrak{A}$ under g.

(5) If P is of the form $\forall v_i \mathrm{Q}$, then P is true in $\mathfrak{A}$ under g iff Q is true in $\mathfrak{A}$ under the assignment g_x^i, for every x in A.

Recall that we don't need to include clauses for the other connectives and $\exists$, because these are all definable from $\sim$, $\wedge$, and $\forall$.

Notation. We write $\mathfrak{A} \models \mathrm{P}[g]$ to mean that P is true in $\mathfrak{A}$ under the assignment g. We also write $\mathfrak{A} \models T[g]$ to mean that $\mathfrak{A} \models \mathrm{P}[g]$ for every $\mathrm{P} \in T$. In this notation, the symbol $\models$ is usually read "satisfies."

We say that P (or T) is **satisfiable** if it is satisfied by *some* $\mathfrak{A}$ and g, and **valid** if it is satisfied by *every* $\mathfrak{A}$ and g.

A straightforward induction on the structure of P shows that the truth of $\mathfrak{A} \models \mathrm{P}[g]$ depends not on all values of g, but only on the values of g on the free variables of P. In particular, if P is a sentence, then g is irrelevant and we simply write $\mathfrak{A} \models \mathrm{P}$, and say that $\mathfrak{A}$ is a **model** of P (or $\mathfrak{A}$ satisfies P, or P is true in $\mathfrak{A}$). In the same vein, if T is a set of sentences we write $\mathfrak{A} \models T$ instead of $\mathfrak{A} \models T[g]$, and we say that $\mathfrak{A}$ is a **model** of T.

Notation. We write $T \models \mathrm{P}$ to mean that whenever $\mathfrak{A} \models T[g]$, we also have $\mathfrak{A} \models \mathrm{P}[g]$.

Two ways to read this notation are to say that P is a **consequence** of T, or T **entails** P, since it says that whenever all the formulas in T are true, so is P. If P and all the formulas in T are sentences, $T \models \mathrm{P}$ says that every model of T is also a model of P. Even though the symbol $\models$ appears in both $\mathfrak{A} \models \mathrm{P}$ and $T \models \mathrm{P}$, these notations have very different meanings. $\mathfrak{A} \models \mathrm{P}$ simply says that P is true in the structure $\mathfrak{A}$. $T \models \mathrm{P}$ is a more complex statement involving all possible structures.

Example 4. Let $\mathcal{L}$ be the first-order language with no relation or constant symbols and a single binary function symbol $\cdot$. Then an $\mathcal{L}$-structure $\mathfrak{A}$ is simply a nonempty set A together with a function from A^2 to A.

Now let T be the first-order theory of a group, that is, the usual list of defining properties or axioms of a group (as in Appendix D). Then $\mathfrak{A} \models T$ says that every sentence of T is true in $\mathfrak{A}$. In other words, it says that $\mathfrak{A}$ is a group. Similarly (in a different language), a model of ring theory is simply a ring, a model of field theory is a field, etc. So there's nothing mysterious about the notion of a model of a theory.

Let P be the commutative law. Then it is clear that $T \not\models \mathrm{P}$, because not every group is abelian. Similarly, $T \not\models \sim \mathrm{P}$.

Finally, let $\mathrm{Q}(v_0)$ be the formula $\forall v_1(v_0 \cdot v_1 = v_1 \cdot v_0)$, and let g be an assignment. Then $\mathfrak{A} \models \mathrm{Q}[g]$ if and only if $g(v_0)$ commutes with every element of A under the binary operation of $\mathfrak{A}$. If $\mathfrak{A}$ is a group, this says that $g(v_0)$ is in the **center** of $\mathfrak{A}$.

Example 5. The **standard model of arithmetic** is the structure $\mathfrak{N} = (\mathbb{N}, +, \cdot, S, 0)$. Here, a symbol like + denotes an actual function on $\mathbb{N}$; it is not a symbol of a formal language. (This sort of ambiguous usage is quite common, and at times one must be very careful to avoid difficulties that could occur from it.) $\mathfrak{N}$ is of course a structure for the language of Peano arithmetic. It is not trivial to prove that $\mathfrak{N} \models \mathrm{PA}$, but it is not very difficult.

The next example and the following two exercises examine whether certain axioms are true in structures that are not "intended" to be models of those axioms. This is good practice because it forces us to carefully examine the satisfaction relation without preconceived notions.

Example 6. Let $\mathcal{L}$ be the first-order language of set theory and let $\mathfrak{A}$ be the $\mathcal{L}$-structure $(\mathbb{R}, <)$. If P is the extensionality axiom of ZF, then the interpretation of P in $\mathfrak{A}$ is

$$\forall x, y \in \mathbb{R}[x = y \leftrightarrow \forall u \in \mathbb{R}(u < x \leftrightarrow u < y)].$$

This is clearly true, so $\mathfrak{A} \models \mathrm{P}$. If Q is the pairing axiom, the interpretation of Q in $\mathfrak{A}$ is

$$\forall x, y \in \mathbb{R} \exists z \in \mathbb{R} \forall u \in \mathbb{R}(u < z \leftrightarrow u = x \vee u = y).$$

This is clearly false, so $\mathfrak{A} \not\models \mathrm{Q}$. In fact, if the quantifiers on x and y were removed from Q, it would be false in $\mathfrak{A}$ under every assignment.

Exercise 1. Let $\mathfrak{A}$ be $(\mathbb{R}, <)$, as in Example 6. Show that the union axiom is true in $\mathfrak{A}$, but the empty set and power set axioms are not.

Exercise 2. Which proper axioms of ZFC are true in the structure $(\mathbb{N}, <)$?

Exercise 3. Show that all the proper axioms of ZFC except the axiom of infinity are true in the structure $(V_\omega, \in)$. Don't try to do this very rigorously, as it could get quite tedious. You may use Proposition 2.24 and other results from the last part of Chapter 2.

5.3 The main theorems of model theory

In this section we present the three theorems that form the foundation of first-order model theory: the completeness theorem, the compactness theorem, and the Löwenheim–Skolem–Tarski theorem.

In Chapter 1 we defined $T \vdash \mathrm{P}$ to mean that P can be deduced from T in first-order logic. This syntactic "single turnstile" relation, based on the notion of a formal proof, is conceptually very different from the semantic "double turnstile" relation $T \models \mathrm{P}$, which is based on the much more abstract concept of truth in structures. One of the most appealing features of first-order logic is that the two "turnstiles," which are the two reasonable notions of logical consequence, actually coincide:

Theorem 5.2 (Gödel's Completeness Theorem). *A set of formulas is satisfiable if and only if it is consistent.*

Proof. The forward direction is straightforward and is often stated separately as the **soundness theorem** for first-order logic. To prove it, assume that T is satisfiable, so that $\mathfrak{A} \models T[g]$ for some $\mathfrak{A}$ and g. Now consider any proof $\mathrm{Q}_1, \mathrm{Q}_2, \ldots, \mathrm{Q}_n$ from T. For any Q_i that is in T, we have $\mathfrak{A} \models \mathrm{Q}_i[g]$ by assumption. Next, it is straightforward to show that every axiom of first-order logic is true in every structure, under every assignment. So $\mathfrak{A} \models \mathrm{Q}_i[g]$ for any Q_i that is a logical axiom. Finally, whenever both P and $\mathrm{P} \to \mathrm{Q}$ are true, then Q must also be true, so modus ponens preserves truth in structures. Thus, by what is essentially induction on the length of a proof, every theorem of T is true in $\mathfrak{A}$ under g. But a contradiction can never be true in a structure, by the definition of $\models$, so T is consistent.

The reverse direction of the proof is complicated, so we will provide only a very bare outline of it. The full proof of both directions can be found in almost any text on mathematical logic or model theory, such as [End]. So assume that T is a consistent theory. For simplicity we first outline the proof under the additional assumption that the language $\mathcal{L}$ of T is denumerable:

(1) Add a denumerable set of new constant symbols $\{c_1, c_2, c_3, \dots\}$ to $\mathcal{L}$ to form a new language $\mathcal{L}'$. These constants are often described as "witnesses" for the elements of the model of T to be constructed. They are essentially Skolem constants, in the sense of Section 1.6.

(2) Create a list $\{\mathrm{P}_1, \mathrm{P}_2, \mathrm{P}_3, \dots\}$ of all $\mathcal{L}'$-formulas whose only free variable is v_0.

(3) Inductively define another list of formulas $\{\mathrm{Q}_1, \mathrm{Q}_2, \mathrm{Q}_3, \dots\}$, where Q_n is $\exists v_0 \mathrm{P}_n \to \mathrm{P}_n(c_k)$. Here, c_k is the first of the new constants that does does not appear in P_n or in any Q_m with $m < n$. Also, $\mathrm{P}_n(c_k)$ means the result of replacing every free occurrence of v_0 in P_n by c_k. The Q_n's are called **Henkin axioms**.

(4) Let T' consist of T and all the Q_n's. It is not hard to prove that T' is still consistent: each Q_n says that if a property holds for at least one object, it holds for the object determined by some new constant symbol. Intuitively, there is no way that statements of this form can lead to a contradiction.

(5) Extend T' to a complete $\mathcal{L}'$-theory T''. To do this, simply list all the sentences of $\mathcal{L}'$ and, starting with T', inductively add each sentence to the theory as long as its negation can't be proved from the theory constructed so far.

(6) Now let $\mathcal{T}$ be the set of all variable-free terms of $\mathcal{L}'$. We would like to use $\mathcal{T}$ as the universe of a model of T'' (and therefore a model of T). In a sense, this is easy, since the completeness of T'' means that T'' determines exactly how all the relation, function, and constant symbols of $\mathcal{L}'$ must be interepreted on elements of $\mathcal{T}$.

(7) However, there is one problem with the previous step: there might be distinct terms t_1 and t_2 in $\mathcal{T}$ such that the equation $t_1 = t_2$ is provable in T''. Then t_1 and t_2 must be interpreted as the same object in any model of T''. To rectify this, consider the equivalence relation on $\mathcal{T}$ defined by $T'' \vdash (t_1 = t_2)$. Then define A to be the set of all equivalence classes. Because of the "substitution of equals" axiom of predicate logic, it is clear that the relations, functions, and constants of $\mathcal{L}'$ are well-defined on A. Thus, we obtain

an $\mathcal{L}'$-structure with universe A, which can be shown to be a model of T''. Hence T'' is satisfiable, and so is T.

The proof when T is uncountable is similar, except that transfinite induction must be used in step (3), and some form of the axiom of choice is needed in step (5) (and perhaps step (2)). ■

For future reference, note that the cardinality of the model constructed in this proof is no greater than the cardinality of $\mathcal{L}$, since $\mathcal{T}$ has the same cardinality as $\mathcal{L}$. (When we refer to the cardinality of a structure, we mean the cardinality of its universe.)

Corollary 5.3. *A formula is valid if and only if it is a law of logic.*

Exercise 4. Prove this corollary.

Corollary 5.4. *For any T and* P*, $T \vdash$* P *if and only if $T \models$* P.

Exercise 5. Prove this corollary. Use Exercise 7(b) of Section 1.4.

This corollary is also referred to as Gödel's completeness theorem. The two versions are easily shown to be equivalent. Note that the completeness theorem does not say that first-order logic, or any particular first-order theory, is complete in the sense of Section 1.4. It is certainly not true that every sentence or its negation is provable in first-order logic. We have instead the weaker result that every valid sentence is provable.

Definition. Two $\mathcal{L}$-structures $\mathfrak{A}$ and $\mathfrak{B}$ are **elementarily equivalent**, denoted $\mathfrak{A} \equiv \mathfrak{B}$, if they satisfy the same sentences of $\mathcal{L}$.

Elementary equivalence is a relatively weak condition, as we will soon see.

Corollary 5.5. *A set of sentences is complete if and only if it has models and all of its models are elementarily equivalent.*

Exercise 6. Prove this corollary.

Theorem 5.6 (Compactness Theorem).

(a) *If every finite subset of a theory T is satisfiable, then T is satisfiable.*

(b) *If $T \models \mathrm{P}$, then there is a finite subset T_0 of T such that $T_0 \models \mathrm{P}$.*

Proof.

(a) Suppose that every finite subset of T is satisfiable. So every finite subset of T is consistent, by the completeness theorem ("soundness"). But then T must be consistent, because a proof only has a finite number of steps. Thus, by the completeness theorem, T is satisfiable.

The proof of (b) is similar. ■

Convention. For the rest of this chapter it will be understood, unless stated otherwise, that a theory means a set of *sentences*.

Corollary 5.7. *If a first-order theory has arbitrarily large finite models, then it has an infinite model.*

Proof. Suppose that T has arbitrarily large finite models. Form a new theory T' by adding to T a denumerable set of new constant symbols $\{c_n \mid n \in \mathbb{N}\}$, and the axiom schema $\{c_m \neq c_n \mid m \neq n\}$. Clearly, every finite subset of T' has a model because we can interpret any finite set of the c_n's as distinct elements in a sufficiently large finite model of T. Therefore T' has a model by the compactness theorem, and any such model must be infinite because all the c_n's must be interpreted as distinct elements. ■

This corollary establishes that in first-order logic, there is no way to state precisely, even with an infinite number of axioms, that there are a finite number of elements in the domain. It follows that there is no single axiom that expresses that there are an infinite number of elements. Note the similar point made in the next to last paragraph of Example 17 of Section 1.5. We will expand on this theme in the first part of Section 5.7.

Corollary 5.8. *If P is true in every infinite model of T, then P is also true in all sufficiently large finite models of T.*

Exercise 7. Prove this corollary from the previous one.

So now we know that every first-order sentence that is true in all infinite groups (or rings, or fields) is also true in all sufficiently large finite ones. By similar reasoning, we can show that every sentence that is true in all fields of characteristic 0 is also true in all fields of sufficiently large finite characteristic.

The compactness theorem is a very powerful tool. One of its most important applications, which will be the subject of Chapter 7, is the following: let T be some first-order theory in which one can carry out calculus or real analysis. T might be ZFC set theory, or it could be something more limited. Form a new theory T' by adding a new constant symbol h to the language of T, and then add to T the new axioms $h \in \mathbb{R}$, $h > 0$, and the schema $\{h < 1/n \mid n \in \mathbb{Z}^+\}$. If T is consistent, it is clear that every finite subset of T' is satisfiable, because we can always interpret h as a positive real number less than any given finite set of positive fractions. Therefore T' is satisfiable. But in a model of T', h must be interpreted as a positive real number that is less than every number of the form $1/n$. In other words, h must be a positive **infinitesimal** in such a model. This rather simple idea led to the development of **nonstandard analysis**.

To close our discussion of the compactness theorem, here are a couple of combinatorial applications of it. First, let's keep the promise that was made at the end of Chapter 4:

Proposition 5.9. *The infinite Ramsey's theorem (4.20) implies the statement* P *of Theorem 4.21 (and hence the finite Ramsey's theorem).*

Proof. Assume Theorem 4.20. Let natural numbers k, m, and n be given; we treat them as fixed. Recall that PA has numerals, terms that denote specific natural numbers. We will define a theory T, in the language of PA with an additional n-ary function symbol g. The proper axioms of T are:

(1) The usual proper axioms of PA.

(2) The statement that the value of g does not depend on the order of its arguments. For instance, if $n = 2$, this becomes $\forall x, y, [g(x, y) = g(y, x)]$. For $n > 2$, the number of equations increases. This statement implies that g defines a function on *unordered* n-tuples. (We don't care about g's values on n-tuples with repeated values.)

(3) The following axiom schema: for each natural number r, the statement that g's values on n-tuples of distinct numbers less than r are all less than m. For each r, this axiom implies that g defines a partition of $r^{(n)}$ into m subsets.

Now, let Q be the statement that says that there is a relatively large set B with at least k elements such that g maps all of $B^{(n)}$ to a fixed number less than m. Q can be expressed in the language of T using the bijection B of Appendix C. We claim that $T \models$ Q. To see this, first note that the universe of any model of PA has a subset that can be identified with ω within the model, namely the elements corresponding to the terms $\overline{0}, \overline{1}, \overline{2}, \ldots$. Then axioms (2) and (3) of T guarantee that g defines a partition of $\omega^{(n)}$ into m subsets. Therefore, by the infinite Ramsey's theorem, there must be an infinite homogeneous subset of ω with respect to g. And any infinite subset S of ω must contain a relatively large set with at least k elements: just take the first $k + s$ elements of S, where s is the smallest number in S.

By compactness, there is a finite subtheory T' of T such that $T' \models$ Q. Let j be the largest number r such that T' includes the statement of axiom schema (3) for r. To see that this j works to prove the desired result, suppose f is any partition of $j^{(n)}$ into m subsets. If we take the standard model of arithmetic, use f to interpret g on all n-tuples of distinct numbers less than j, and let g's value equal m on all other n-tuples, we obtain a model of T'. Since Q holds in this model, there is a relatively large homogeneous set with at least k elements, with respect to f. ■

Exercise 8. Give a simpler proof than the one just given (still using compactness) of the regular finite Ramsey's theorem from the infinite one. Your theory T need not include any of the axioms of PA.

Definitions. A **(nondirected) graph** with domain G is a pair (G, R), where R is any subset of $G^{(2)}$. Think of $\{x, y\} \in R$ as meaning that the "vertices" x and y are "connected."

A **(finite) subgraph** of (G, R) is a pair of the form $(H, R \cap H^{(2)})$, where H is a (finite) subset of G.

A k-coloring of a graph (G, R), where $k \in \mathbb{Z}^+$, is a function $g : G \to \{1, 2, \ldots, k\}$ such that $g(x) \neq g(y)$ whenever $\{x, y\} \in R$.

Theorem 5.10 (de Bruijn). *If every finite subgraph of a graph has a k-coloring, then so does the whole graph.*

Exercise 9. Prove de Bruijn's theorem. You can use an argument similar to, but simpler than, the proof of the previous proposition. Note that here the compactness theorem allows us to go from a finite version of a result to an infinite one, whereas the previous argument works in the opposite direction.

When presented to nonmathematicians, the four-color theorem is usually stated in terms of maps, but it has an equivalent formulation in terms of graphs: every finite planar graph can be 4-colored. (A planar graph is one that can be drawn in a plane, with connections between vertices shown as nonintersecting continuous curves. For example, it's certainly possible for five points to be connected to each of the others in a graph, but this cannot occur in a planar graph.) By de Bruijn's theorem, it follows that every planar graph can be 4-colored, and so the four-color theorem also holds for maps with an infinite number of countries.

The Löwenheim–Skolem–Tarski theorem

We now present the final cornerstone of first-order model theory. Chronologically, it was actually the first—Löwenheim proved the simplest version of it in 1915, well over a decade before the completeness and compactness theorems were obtained.

Theorem 5.11 (Löwenheim–Skolem–Tarski (LST) Theorem). *Let T be a theory in a first-order language $\mathcal{L}$. If T has an infinite model,* *then it has a model of every cardinality equal to or greater than* $Card(\mathcal{L})$.

Proof. Once again, we just outline the proof. The first task is to prove that if T is any satisfiable theory in a first-order language $\mathcal{L}$, then T has a model of cardinality $Card(\mathcal{L})$ or smaller. This follows directly from the construction used to prove Gödel's completeness theorem—note the remark about cardinality following our outline of that construction. Löwenheim and Skolem's original argument was somewhat simpler, because they could construct a model of "Skolem terms" within a given model, instead of needing to build a model from scratch.

For clarity, this result is often called the *downward* Löwenheim–Skolem theorem. Löwenheim proved it for finite T, and Skolem extended it to denumerable theories a few years later. Uncountable theories were not considered at that time. In the next section we will state a stronger version of the downward theorem (Theorem 5.13).

The full LST theorem follows from the downward theorem by the following argument: assume that T is a theory in some language $\mathcal{L}$ and that T has an infinite model. Let κ be any cardinal with $\kappa \geq Card(\mathcal{L})$. We use a simple adaptation of the proof of Corollary 5.7: instead of a denumerable set of new constant symbols, introduce κ new constant symbols and the axiom schema that says they are all distinct. Call this expanded theory and language T' and $\mathcal{L}'$, respectively. Note that $Card(\mathcal{L}') = \kappa$. By the compactness theorem, T' is satisfiable. Therefore, by the downward Löwenheim–Skolem theorem, T' has a model of cardinality $\leq \kappa$. But any model of T' clearly has cardinality $\geq \kappa$. Therefore, T' has a model of cardinality κ. ■

Independent of the downward theorem, the argument in the last paragraph of this proof shows that if a first-order theory has an infinite model, then it has models of arbitrarily large cardinality. This is sometimes called the *upward* Löwenheim–Skolem theorem, but it was actually proved by Tarski and Robert Vaught many years later. In fact, Skolem didn't even accept this result, since he did not believe in the existence of uncountable sets.

Corollary 5.12. *Let κ be an infinite cardinal and let $\mathfrak{A}$ be an infinite structure for a countable language. Then there is a structure of cardinality κ that is elementarily equivalent to $\mathfrak{A}$.*

Proof. Apply the LST theorem with T being the set of all sentences that are true in $\mathfrak{A}$. ■

The completeness theorem shows the equivalence of the syntactic and semantic notions of consequence. By contrast, the compactness theorem and the LST theorem make no mention of $\vdash$, so they are purely semantic results. Our proof of compactness used completeness, but it is also possible to give a purely semantic proof.

As indicated in the introduction to this chapter, the LST theorem shows that first-order logic is quite limited in its ability to describe a specific structure. The above corollary makes this point more clear, and the next two examples illustrate it further:

Example 7. We would like Peano's axioms for arithmetic to define $\mathbb{N}$ (with the usual operations), and only that structure. But now we see that (first-order) PA must have uncountable models. What would an uncountable model $\mathfrak{A}$ of PA look like? In any model of PA, there must be distinct elements that are the interpretations of the terms $\overline{0}$, $S(\overline{0})$, $S(S(\overline{0}))$, etc. It is not hard to show that the addition and multiplication operations on these elements must correspond to the usual operations on $\mathbb{N}$. In other words, $\mathfrak{A}$ must have a "substructure" that looks exactly like the structure $\mathbb{N}$.

But $\mathfrak{A}$ must of course have many more elements, since it is uncountable. All of these elements must be greater than the elements corresponding to $\mathbb{N}$ in the ordering of $\mathfrak{A}$. Since it is provable in PA that every element except 0 has both an immediate sucessor and an immediate predecessor, the ordering on $\mathfrak{A}$ must look like "a copy of $\mathbb{N}$ followed by an uncountable number of copies of $\mathbb{Z}$." The addition and multiplication operations in $\mathfrak{A}$ must be defined in such a way that all the usual properties of arithmetic hold, including the Euclidean algorithm, the unboundedness of the set of prime numbers, etc., not to mention the entire first-order axiom schema of mathematical induction. It is difficult to picture such a model; in fact, it is difficult to picture any model of PA that does not look like the model $\mathbb{N}$.

Example 8. The LST theorem tells us that ZFC set theory must have countable models, if it's consistent. This seems impossible, since within any model of set theory there must be "uncountable" sets. This so-called **Skolem's paradox** disappears when one realizes that a set can satisfy the definition of being uncountable *within* a model without actually being uncountable.

More specifically, imagine a model of set theory in which the interpretation of $\mathbb{N}$ is itself (with its usual elements), and $\mathcal{P}(\mathbb{N})$ is interpreted as some set B. Within the model, B must be uncountable, but that merely means that there is no bijection between $\mathbb{N}$ and B *in the model*. B could really be countable; in fact, we know that the whole model could be countable. We will encounter several more examples of this type of phenomenon in this book. Because of this example and others like it, the downward Löwenheim–Skolem theorem was a very surprising development at the time.

5.4 Preservation theorems

In this section we present several results of first-order model theory that explain why various operations (unions, intersections, homomorphic images, etc.) on mathematical structures do or do not preserve various properties of those structures.

Recall that a function $f : A \to B$ automatically induces functions from A^n to B^n, for each $n \in \mathbb{N}$, as well as functions from $\mathcal{P}(A^n)$ to $\mathcal{P}(B^n)$. In the next three definitions, it is assumed that $\mathfrak{A}$ and $\mathfrak{B}$ are structures for the same first-order language $\mathcal{L}$.

Definition. An **isomorphism** f between $\mathfrak{A}$ and $\mathfrak{B}$ is, as usual, a bijection between A and B that preserves all structural components: for each constant symbol of $\mathcal{L}$, its interpretation in $\mathfrak{A}$ is mapped to its interpretation in $\mathfrak{B}$; and similarly for relation symbols and function symbols, using the functions induced by f.

We write $\mathfrak{A} \cong \mathfrak{B}$ to mean that $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic (that is, there is an isomorphism between them).

Example 9. It is very easy to show that isomorphic structures must be elementarily equivalent. Thus, for instance, the group of integers $(\mathbb{Z}, +)$ and the group of even integers $(2\mathbb{Z}, +)$, where $+$ denotes ordinary addition, are elementarily equivalent since they are isomorphic under the function $f(n) = 2n$.

On the other hand, the *rings* $(\mathbb{Z}, +, \cdot)$ and $(2\mathbb{Z}, +, \cdot)$ are not elementarily equivalent, since only the first one has a multiplicative identity, a property which can be stated in the first-order language of a ring. Therefore, these rings are not isomorphic.

Some standard examples of elementarily equivalent structures that are not isomorphic will be given in Example 12.

Example 10. Let $\mathfrak{A} = (\mathbb{Z}\times\mathbb{Z}, P)$, the usual direct product of the group of integers with itself. (So P is "componentwise addition.") Then $\mathfrak{A}$ is not isomorphic to the group of integers. One standard argument for this is that the group $(\mathbb{Z}, +)$ is cyclic, generated by a single element (1 or -1), while $\mathfrak{A}$ is not cyclic.

Exercise 10.

(a) Give a direct proof that there is no isomorphism between the groups $(\mathbb{Z}, +)$ and $\mathfrak{A}$ discussed in the previous example.

(b) Even though these groups are not isomorphic, the property that a group is cyclic cannot be stated in the first-order language of a group. So these groups might still be elementarily equivalent. Show that in fact they are not. (Hint: the fact that every integer is even or odd can be used to construct a first order sentence that is true in $(\mathbb{Z}, +)$ but not in $\mathfrak{A}$.)

Definition. We say that $\mathfrak{A}$ is a **submodel** or **substructure** of $\mathfrak{B}$, denoted $\mathfrak{A} \subseteq \mathfrak{B}$, if: (i) $A \subseteq B$; (ii) for each constant symbol of $\mathcal{L}$, its interpretations in $\mathfrak{A}$ and $\mathfrak{B}$ are the same; and (iii) for each relation symbol and function symbol of $\mathcal{L}$, its interpretation in $\mathfrak{A}$ is the restriction to the appropriate A^n of its interpretation in $\mathfrak{B}$. It is implicit in this definition that A must be closed under all the functions that are part of the structure $\mathfrak{B}$.

If $\mathfrak{A} \subseteq \mathfrak{B}$, we also say that $\mathfrak{B}$ is an **extension** of $\mathfrak{A}$.

Definitions. We say that $\mathfrak{A}$ is an **elementary submodel** of $\mathfrak{B}$ or $\mathfrak{B}$ is an **elementary extension** of $\mathfrak{A}$, denoted $\mathfrak{A} \preceq \mathfrak{B}$, if $\mathfrak{A} \subseteq \mathfrak{B}$ and, for every $\mathcal{L}$-formula P and every $\mathfrak{A}$-assignment g, $\mathfrak{A} \models \mathrm{P}[g]$ if and only if $\mathfrak{B} \models \mathrm{P}[g]$. (You may recall that we are also using the symbol $\preceq$ to compare the cardinality of sets. There is no connection between the two meanings of this symbol.)

An isomorphism between $\mathfrak{A}$ and an elementary submodel of $\mathfrak{B}$ is called an **elementary embedding** of $\mathfrak{A}$ in $\mathfrak{B}$. More concretely, an elementary embedding of $\mathfrak{A}$ in $\mathfrak{B}$ is a function f from A to B such that, for every $\mathcal{L}$-formula P and every $\mathfrak{A}$-assignment g, $\mathfrak{A} \models \mathrm{P}[g]$ if and only if $\mathfrak{B} \models \mathrm{P}[f \circ g]$.

The property $\mathfrak{A} \preceq \mathcal{B}$ is very strong, strictly stronger than the conjunction of $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{A} \equiv \mathfrak{B}$, because the definition of $\mathfrak{A} \preceq \mathcal{B}$ refers to formulas and assignments, whereas $\mathfrak{A} \equiv \mathfrak{B}$ pertains only to sentences. Even $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{A} \cong \mathfrak{B}$ together do not imply $\mathfrak{A} \preceq \mathfrak{B}$, as the next example shows.

Example 11. Let $\mathcal{L}$ be the first-order language of an ordering: besides equality, it has a single binary relation symbol. Consider the two $\mathcal{L}$-structures $\mathfrak{A} = (\mathbb{Z}, <)$ and $\mathfrak{B} = (2\mathbb{Z}, <)$, as defined in Appendix D. $\mathfrak{B} \subseteq \mathfrak{A}$ and $\mathfrak{B} \cong \mathfrak{A}$. But it is not the case that $\mathfrak{B} \preceq \mathfrak{A}$. For instance, the formula $\exists v_1(v_0 < v_1 < v_2)$ is false in $\mathfrak{B}$ but true in $\mathfrak{A}$, if v_0 is assigned the value 0 and v_2 is assigned the value 2.

The following result of Tarski and Vaught is a strengthened version of the downward Löwenheim–Skolem theorem, provable by essentially the same construction.

Theorem 5.13. *Let $\mathfrak{A}$ be any structure for a first-order language $\mathcal{L}$. Then $\mathfrak{A}$ has an elementary submodel whose cardinality is no greater than that of $\mathcal{L}$.*

Example 12. It can be shown that the denumerable field $\mathbb{A}$ of complex algebraic numbers (see Theorem 8.2) is an elementary submodel of the uncountable field $\mathbb{C}$. Similarly, $\mathbb{R} \cap \mathbb{A} \preceq \mathbb{R}$.

Preservation under submodels and intersections

In the remainder of this section, we will present several **preservation properties**, all of which (with full proofs and much more detail) can be found in [CK].

Definition. Let T be a theory. T is said to be **preserved under submodels** if $\mathfrak{B} \models T$ and $\mathfrak{A} \subseteq \mathfrak{B}$ imply $\mathfrak{A} \models T$.

Similarly, we can define what it means for T to be preserved under finite intersections, preserved under arbitrary intersections, preserved under homomorphic images, etc.

Theorem 5.14 (Łoś–Tarski). *A theory is preserved under submodels if and only if it is equivalent to a set of* Π_1 *sentences.*

Exercise 11. Prove the reverse direction of this theorem. (This is the easy direction. The forward direction requires substantial proof.)

We mention the obvious "companion result":

Corollary 5.15. *A theory is preserved under extensions if and only if it is equivalent to a set of* Σ_1 *sentences.*

Now let's see some applications of this theorem. For example, is every submodel of a group also a group? It depends on what we mean by "submodel." If we express the axioms of a group in the first-order language with multiplication only, then the identity and inverse axioms require existential quantifiers, so the theorem tells us that this theory is not preserved under submodels. And, clearly, a subset of a group may be closed under multiplication without being a subgroup. But if we express the axioms of a group in the first-order language with symbols for the identity and inverses as well as multiplication, then the natural axiomatization consists of Π_1 sentences, and in this context every submodel is a subgroup. This situation applies to many types of algebraic structures:

Corollary 5.16.

(a) *Every subset of a group that contains* 1 *and is closed under multiplication and inverses is a subgroup.*

(b) *Every subset of a ring that contains* 0 *and is closed under* $+$, $-$, *and* $\cdot$ *is a subring.*

(c) *Every subset of a field that contains* 0 *and* 1 *and is closed under* $+$, $-$, $\cdot$, *and* $\div$ *is a subfield.*

Proof. Immediate from Theorem 5.14. ∎

There are, of course, endless variations on this corollary. Also, note that the converses of the parts of this corollary are true by definition, so they could be stated as biconditionals: a subset of a group is a subgroup if and only if it contains 1 and is closed under multiplication and inverses, etc.

Corollary 5.17. *The intersection of any collection of subgroups of a group is also a subgroup (and similarly for rings and fields).*

Proof. It is clear that if each set in a collection is closed under a certain operation (e.g., multiplication), then so is the intersection of that collection. ∎

So the well-known fact that most types of algebraic structures behave well under intersections may be viewed as a model-theoretic result. It is not hard to see why these same types of structures are not preserved under even finite unions of submodels: if two subsets of a set are closed under some *unary* function, then so is their union. But this property fails for functions of more than one variable, including binary operators such as + and ·.

Corollaries 5.16 and 5.17 also help clarify the notion of the subgroup of a given group (or subfield of a given field, etc.) **generated** by an arbitrary subset. See Appendix D for the various equivalent definitions of this concept.

Preservation under unions of chains

While unions of algebraic structures are generally not algebraic structures of the same type, unions of *chains* of algebraic structures usually are:

Definition. Let I be a well-ordered set. A collection of structures $\{\mathfrak{A}_i \mid i \in I\}$ indexed by I is called a **chain** (respectively, **elementary chain**) if $\mathfrak{A}_i \subseteq \mathfrak{A}_j$ (respectively, $\mathfrak{A}_i \preceq \mathfrak{A}_j$) whenever $i < j$.

For many applications, it suffices to consider chains in which $I = \mathbb{N}$.

If $\{\mathfrak{A}_i \mid i \in I\}$ is a chain, then we can define the new structure $\mathfrak{B} = \bigcup_{i \in I} \mathfrak{A}_i$ in the obvious way, and it is obvious that $\mathfrak{A}_i \subseteq \mathfrak{B}$ for every $i \in I$. The analogous result for elementary chains, due to Tarski, is not much more difficult to prove:

Exercise 12. Prove that if $\{\mathfrak{A}_i \mid i \in I\}$ is an elementary chain and $\mathfrak{B} = \bigcup_{i \in I} \mathfrak{A}_i$, then $\mathfrak{A}_i \preceq \mathfrak{B}$ for each $i \in I$. (Hint: The fact to be proved involves a formula P, so prove it by induction on the structure of P. The only nontrivial step is the one involving $\exists$—there is no need for a separate step involving $\forall$.) Consequently, every first-order sentence is preserved under unions of elementary chains.

Theorem 5.18 (Chang–Łoś–Suzsko). *For any theory T, the following are equivalent:*

(a) *T is equivalent to a set of Π_2 sentences. (Such a theory is called **inductive**.)*

(b) *T is preserved under unions of chains.*

(c) *T is preserved under unions of chains indexed by $\mathbb{N}$ (with the usual ordering).*

Proof. We prove the two easier parts of this theorem:

(a) implies (b): Assume that T is equivalent to a set T' of Π_2 sentences, and let $\{\mathfrak{A}_i \mid i \in I\}$ be a chain of models of T, with union $\mathfrak{B}$. We want to show $\mathfrak{B} \models T$, for which it suffices to prove that $\mathfrak{B} \models T'$.

So let P $\in T'$. P has the form $\forall x_1, x_2, \ldots, x_m \exists y_1, y_2, \ldots, y_n Q$, where Q is quantifier free. Let $a_1, a_2, \ldots, a_m \in B$. Then there must be an $i \in I$ such that $a_1, a_2, \ldots, a_m \in A_i$. Since $\mathfrak{A}_i \models T$, it follows that $\mathfrak{A}_i \models T'$, and thus $\mathfrak{A}_i \models$ P. So there exist $b_1, b_2, \ldots, b_n \in A_i$ such that $\mathfrak{A}_i \models Q$ when each x_j is assigned to a_j and each y_k is assigned to b_k. But then $\mathfrak{B} \models Q$ under the same assignment. Therefore $\mathfrak{B} \models$ P, as desired.

Trivially, (b) implies (c). We omit the proof that (c) implies (a). Interestingly, this proof requires the use of elementary chains, even though the theorem refers only to ordinary chains. ■

It follows immediately from this theorem that the union of any chain of groups is a group, and similarly for rings, fields, and many other algebraic structures. However, there are some subtleties involved in this claim:

Exercise 13.

(a) Write out the usual axioms for a group in a first-order language with symbols for multiplication and the identity, both with and without a symbol for inverses. Note that these axioms are all Π_2, so the union of any chain of groups in this language is again a group.

(b) Now write out the usual axioms for a group in a language with a symbol for multiplication but no symbol for the identity or inverses. Note that these axioms are not all Π_2.

(c) Show that, in spite of part (b), the union of any chain of groups in this language is also a group. The core of the proof is to show that all the groups in the chain have the same identity element.

(d) Therefore, there must be a set of Π_2 axioms for a group in this language too. Find such a list of axioms.

Preservation under homomorphic images

Definition. A **homomorphism** from $\mathfrak{A}$ to $\mathfrak{B}$ is a function f from A to B such that for any atomic formula P and any $\mathfrak{A}$-assignment g, $\mathfrak{A} \models \text{P}[g]$ implies $\mathfrak{B} \models \text{P}[f \circ g]$.

We say that $\mathfrak{B}$ is a **homomorphic image** of $\mathfrak{A}$ if there is a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$ that is *onto* B.

Exercise 14. Consider a simple first-order language, such as the language with one binary function symbol $\cdot$ and one binary relation symbol $<$. This could be the language of an ordered group, without symbols for the identity and inverses. Now verify that the above definition coincides with the usual definition of homomorphism for structures of this language, namely: a function between their universes such that $f(x \cdot y) = f(x) \cdot f(y)$, and $x < y$ implies $f(x) < f(y)$, for every x and y in the domain of f.

Definition. A first-order formula is said to be **positive** if it does not contain the connectives $\sim$, $\rightarrow$, or $\leftrightarrow$. So it may contain $\wedge$ and $\vee$, as well as the quantifiers.

Clearly, a positive formula must be true whenever all of its atomic subformulas are true. (This fact still holds if the connectives $\rightarrow$ and $\leftrightarrow$ are also allowed.) Furthermore, if a positive formula is true under a certain interpretation, it must remain true if the values of one or more atomic subformulas change from false to true.

Theorem 5.19 (R. C. Lyndon). *Let T be a consistent theory. Then T is preserved under homomorphic images if and only if T is equivalent to a set of positive sentences.*

We omit the proof. The reverse direction is easy to prove—a simple induction on the structure of P shows that every true positive formula stays true in a homomorphic image. (In fact, the definition of a homomorphism says precisely that this holds for atomic formulas.) The proof of the forward direction is more difficult.

Example 13. The usual axioms for a group are positive, whether or not symbols for the identity and inverses are included in the language. Therefore, every homomorphic image of a group is a group. The same holds for rings.

Example 14. The usual axioms for a ring with unity are not all positive, since they include the statement $0 \neq 1$. A fortiori, the usual axioms for an integral domain or a field are not all positive. In fact, none of these theories is preserved under homomorphic images: the unique mapping from the ring $\mathbb{R}$ to the ring $\{0\}$ is a ring homomorphism, but the domain is a field while the range is not even a ring with unity.

The usual form of the multiplicative inverse axiom contains the connective $\to$, but it has a positive equivalent:

$$\forall x[x = 0 \vee \exists y(x \cdot y = 1)].$$

(It also has a positive equivalent if the language does not have the symbols 0 and 1.) So we get a theory that is preserved under homomorphic images if we omit $0 \neq 1$ from the field axioms. What results is the first-order theory of fields *and* one-element rings.

Preservation under direct products

One of the most important ways of combining mathematical structures is to form their direct product. One can form direct products of just about every type of algebraic structure (see Appendix D), as well as orderings, topological spaces, etc. If $\{\mathfrak{A}_i \mid i \in I\}$ is a collection of structures for the same first-order language $\mathcal{L}$, we use the usual notation $\prod_{i \in I} \mathfrak{A}_i$ for their direct product. Finite direct products such as $\mathfrak{A} \times \mathfrak{B}$ and $\mathfrak{A} \times \mathfrak{B} \times \mathfrak{C}$ may be viewed as special cases of the general notion.

It would be nice to have a clear preservation result about theories preserved under direct products, but unfortunately the situation here is not as clear as with submodels, chains or homomorphic images. For instance, it is well known that an arbitrary direct product of groups is a group, and similarly for rings. But the direct product of just two fields is never even an integral domain. This is reminiscent of the situation for homomorphic images, but it is much harder with direct products to identify a syntactic condition on sentences that is equivalent to preservation. (Such a syntactic condition exists, but it is very involved and we will not present it here.)

We begin with a fact that at least makes it easier to think about preservation under direct products:

Proposition 5.20. *If a sentence is preserved under direct products of two structures, then it is preserved under arbitrary direct products.*

But even with the simplification afforded by this fact, the preservation situation is quite complex. Here are the relevant syntactic notions:

Definitions. A **basic Horn formula** is a formula that is of one of the forms P_1, $\sim \mathrm{P}_1$, or $(\mathrm{P}_1 \wedge \mathrm{P}_2 \wedge \cdots \wedge \mathrm{P}_n) \rightarrow \mathrm{P}_{n+1}$, where all of the P_i's are atomic. A **Horn formula** is a formula built up from basic Horn formulas using quantifiers and the connective $\wedge$. A **Horn sentence** is a Horn formula with no free variables.

Theorem 5.21.

(a) *If a sentence is equivalent to a Horn sentence, then it is preserved under direct products.*

(b) *If a Π_2 sentence is preserved under direct products, then it is equivalent to a Π_2 Horn sentence.*

The restriction to Π_2 sentences in (b) cannot be removed— there are Σ_2 counterexamples. The restriction can be removed if we talk about **reduced products** (a generalization of the notion of direct products) instead of direct products. In other words, a sentence is preserved under reduced products if and only if it is equivalent to a Horn sentence. However, we will not define reduced products or discuss them further because they are quite specialized.

Example 15. The axioms of a partial ordering (as in Appendix B) are Horn sentences, so any direct product of partial orderings is a partial ordering.

The extra axiom of a total ordering, trichotomy, is not a Horn sentence and there is no obvious equivalent Horn sentence. In fact there can't be one, because this property is not preserved under direct products. For instance, the product ordering on $\mathbb{R} \times \mathbb{R}$ is not total.

Example 16. The axioms of a group are all Horn sentences, so every direct product of groups is a group. This is true even if the language does not include symbols for the identity and inverses. The same holds for rings, commutative rings, and rings with unity.

Exercise 15. Write out the inverse axiom for groups in the language without symbols for the identity and inverses, and convince yourself that it's a Horn sentence.

Example 17. An integral domain is a commutative ring with unity that also has no **zero-divisors**: $\forall x, y(x \neq 0 \wedge y \neq 0 \rightarrow x \cdot y \neq 0)$. This is not a Horn sentence.

A field is a commutative ring with unity that also satisfies the multiplicative inverse axiom, $\forall x[x \neq 0 \rightarrow \exists y(x \cdot y = 1)]$. This is also not a Horn sentence.

In fact, neither of these two properties is equivalent to a Horn sentence, because they are not preserved under direct products. For instance, the ring $\mathbb{R} \times \mathbb{R}$ is the direct product of two fields but it is not even an integral domain.

5.5 Saturation and complete theories

In Chapter 4 we saw that many important first-order theories cannot be complete. We will now use model theory to identify some complete theories. It is valuable to know that a theory is complete, because then we know it cannot be strengthened without passing to a more powerful language. Also, if a theory T is axiomatizable and complete, then it is decidable, by Exercise 1(b) of Chapter 4. This means there is an effective procedure that determines whether or not any given sentence is provable in T, a useful thing to know.

In this section we will discuss a powerful and versatile method for proving that theories are complete. Its main drawback is that it is set-theoretic and rather abstract. In the next section we will describe another method that applies to fewer situations but provides more precise information when it does work. For most of the completeness results to be discussed, we will outline proofs using both methods.

Definition. Let κ be a cardinal. A theory T is said to be **categorical in power** κ or simply κ**-categorical** if it has, up to isomorphism, exactly one model of cardinality κ. In other words, it has models of cardinality κ, and they are all isomorphic. (The word "power" is sometimes used to mean "cardinality.")

Proposition 5.22 (Łoś–Vaught Test). *If T has only infinite models and is κ-categorical for some $\kappa \geq Card(T)$, then T is complete.*

Proof. Assume T is κ-categorical, with $\kappa \geq Card(T)$. Let $\mathfrak{A}_1$ and $\mathfrak{A}_2$ be any models of T. By Corollary 5.12, there is a structure $\mathfrak{A}'_i$ of cardinality κ such that $\mathfrak{A}'_i \equiv \mathfrak{A}_i$, for $i = 1, 2$. By categoricity, $\mathfrak{A}'_1 \cong \mathfrak{A}'_2$. Therefore, $\mathfrak{A}_1 \equiv \mathfrak{A}_2$, so we are done by Corollary 5.5. ■

The condition that T has only infinite models cannot be removed. As a trivial example, let T be the empty theory in the language of pure identity. Then a model of T is simply a nonempty set with no structure, so T is categorical in every nonzero power. But T is not complete; for instance, the statement that there are at least two elements is independent of T.

Notation. Suppose $(L, <)$ is a linear ordering, $y \in L$, and $C, D \subseteq L$. We write $C < D$ (respectively, $C < y < D$) to mean that $c < d$ (respectively, $c < y < d$) for every $c \in C$ and $d \in D$.

Lemma 5.23. *For any linear ordering $(L, <)$, the following are equivalent:*

(a) *$(L, <)$ is a dense unbounded ordering, that is, a dense ordering with no greatest or least element.*

(b) *Whenever C and D are finite subsets of L such that $C < D$, there is a $y \in L$ such that $C < y < D$.*

Exercise 16. Prove this lemma. Be sure to consider the possibility that C and/or D is empty.

We are now in a position to prove one of the simplest and oldest completeness results:

Theorem 5.24. *The theory of a dense unbounded ordering is* $\aleph_0$*-categorical, and therefore complete.*

Proof. Let T be this theory. Clearly, T is consistent (because $\mathbb{Q}$ is a model) and has no finite models. So once we show that T is $\aleph_0$-categorical, completeness follows by the Łoś–Vaught test.

Let $\mathfrak{A}$ and $\mathfrak{B}$ be any two countable models of T. Say $A = \{a_n \mid n \in \mathbb{N}\}$ and $B = \{b_n \mid n \in \mathbb{N}\}$. We will use the symbol $<$ to denote the ordering in both $\mathfrak{A}$ and $\mathfrak{B}$. We now use an inductive process to define an isomorphism f between $\mathfrak{A}$ and $\mathfrak{B}$:

At stage 1, let $f(a_1) = b_1$. (This is the only stage where we carry out one step instead of two.)

At stage $n+1$, up to $2n$ values of f will already have been defined, in such a way that f is order-preserving so far. Let E be the set of members of A at which f has already been defined. If $a_{n+1} \in E$, skip the rest of this paragraph. Otherwise, let $E_1 = \{x \in E \mid x < a_{n+1}\}$ and $E_2 = \{x \in E \mid x > a_{n+1}\}$. We need to define $f(a_{n+1})$ and keep f order-preserving. Let $C = f(E_1)$ and $D = f(E_2)$. Clearly $C < D$, and so by the previous lemma, there is a $y \in B$ such that $C < y < D$. Let $f(a_{n+1})$ be any such y.

The second step of stage $n+1$ is to include b_{n+1} in the range of f (if it's not already there), in such a way that f stays order-preserving. This process is nearly identical to the definition of $f(a_{n+1})$, and so we omit the details.

It is clear that the function f defined by this process is an order-preserving function from the whole set A onto the set B. ■

Note that the construction of f in this proof requires defining two values of f at each stage. This is the classic example, due to Cantor, of a **back-and-forth argument**, one of the most powerful techniques in model theory. For an introduction to the more sophisticated uses of this method, see [Bar73].

Exercise 17. Prove that the following theories are also complete: the theory of a dense ordering with greatest and least elements; the theory of a dense ordering with least element but no greatest element; and the theory of a dense ordering with greatest element but no least element.

These four theories are not categorical in any uncountable power. For example, $\mathbb{R}$, "a copy of $\mathbb{R}$ followed by a copy of $\mathbb{Q}$," and "a copy of $\mathbb{Q}$ followed by a copy of $\mathbb{R}$" are all dense unbounded orderings of the same uncountable cardinality, but no two of them are isomorphic.

Exercise 18. Prove everything stated in the previous sentence.

Why doesn't the construction of f in the proof of Theorem 5.24 work for uncountable dense unbounded orderings? The first obvious difference is that we must use transfinite induction instead of ordinary induction when the sets are uncountable. Then, within the induction process, Lemma 5.23 no longer suffices. Instead, we would need a generalization of this lemma that replaces "finite" by "of smaller cardinality than L." And this generalization is false.

These considerations led Felix Hausdorff and others to consider those uncountable orderings for which Lemma 5.23 can be generalized. Here are some relevant notions:

Definitions. A collection of sets is said to have the **finite intersection property** if the intersection of any finite number of those sets is nonempty. We will call a collection $\mathcal{C}$ of sets **resilient** if every subcollection of $\mathcal{C}$ with the finite intersection property has nonempty intersection.

Example 18.

(a) Every finite collection of sets is resilient, trivially.

(b) One way of defining compactness of a topological space is that the collection of all of its closed subsets is resilient. By the well-known Heine–Borel theorem, every closed, bounded subset of $\mathbb{R}$ is compact. Therefore, the collection of all closed, bounded subsets of $\mathbb{R}$ is resilient. In particular, the collection of all bounded closed intervals in $\mathbb{R}$ is resilient.

(c) The collection of closed rays $\{[n, \infty) \mid n \in \mathbb{N}\}$ in $\mathbb{R}$ has the finite intersection property but its intersection is empty. So this collection is not resilient.

Exercise 19. Find a countable collection of bounded open intervals in $\mathbb{R}$ that is not resilient.

We can now establish the correct generalization of Lemma 5.23 for uncountable orderings:

Proposition 5.25. *Let $(L, <)$ be a total ordering and let κ be an uncountable cardinal. Then the following are equivalent:*

(a) *$(L, <)$ is dense and unbounded, and every collection of fewer than κ open intervals and open rays in L is resilient.*

(b) *Whenever C and D are subsets of L of cardinality less than κ such that $C < D$, there is a $y \in L$ such that $C < y < D$.*

Proof. For the forward direction, assume the givens, and let C and D be subsets of L of cardinality less than κ such that $C < D$. We consider three cases: If $D = \emptyset$, consider the collection of all open rays $\{x \mid x > c\}$, where $c \in C$. Within any finite subset of this collection, there is a greatest element u among the left-hand endpoints of these rays, and since $(L <)$ is unbounded, there is an element of L that is greater than u. So this collection of rays has the finite intersection property, and therefore has nonempty intersection. And if y is in the intersection of all these rays, we clearly have $C < y < D$.

The proof for the case that $C = \emptyset$ is nearly identical.

Finally, assume C and D are both nonempty. Then consider the collection of all open intervals of the form (c, d), where $c \in C$ and $d \in D$. By Proposition C.1(c) in Appendix C, the cardinality of this collection is less than κ. Also, within any finite subset of this collection, there is a greatest element u among the left-hand endpoints of these open intervals, and a least element v among the right-hand endpoints of these open intervals. Since $C < D$, $u < v$. Because $(L <)$ is dense, there is an element of L between u and v. So this collection of intervals has the finite intersection property, and therefore has nonempty intersection. And if y is in the intersection of all these intervals, we clearly have $C < y < D$.

The proof of the other direction is left for the following exercise.

■

Exercise 20. Prove the reverse direction of this proposition.

Hausdorff defined an η_α-set to be a linear ordering that satisfies the conditions of Proposition 5.25 with $\kappa = \aleph_\alpha$. Over time, model theorists devised the notion of **saturation** to generalize this notion to first-order structures other than linear orderings. Before we can define saturation, we need to define a more basic and important concept that we have been mentioning informally since Chapter 1:

Definition. Let $\mathfrak{A}$ be an $\mathcal{L}$-structure. A subset of A is called **(first-order) definable** (in $\mathfrak{A}$) if it is of the form $\{x \in A : \mathfrak{A} \models \mathrm{P}[g_x^0]\}$, for some $\mathcal{L}$-formula P and some assignment g. If P has no free variables other than v_0, we call the set **definable without parameters** or **$\emptyset$-definable** ("zero-definable") in $\mathfrak{A}$.

Similarly, one can define the notion of a definable (with or without parameters) n-ary relation in $\mathfrak{A}$. A definable function is one whose graph is definable.

Example 19.

(a) In the structure $(\mathbb{R}, <)$, every interval is definable. For instance, $(e, 5\pi)$ is $\{x \in \mathbb{R} : e < x \wedge x < 5\pi\}$. The numbers e and 5π act as parameters here. More precisely, let P be the formula $(v_1 < v_0 \wedge v_0 < v_2)$, and let g be any assignment such that $g(v_1) = e$ and $g(v_2) = 5\pi$. Of course, not every interval in this structure is $\emptyset$-definable, because its language is denumerable while the number of intervals is uncountable.

(b) All intervals are also definable in the field of real numbers, with no inequality symbol. For instance, $(e, 5\pi)$ is

$$\{x \in \mathbb{R} \mid \exists y, z \in \mathbb{R}(y \neq 0 \wedge z \neq 0 \wedge e + y^2 = x \wedge x + z^2 = 5\pi)\}.$$

(c) Every finite subset of the domain of a structure is definable, because any finite set of elements can be used as parameters in a single formula. Similarly, every **cofinite** subset (a subset whose complement is finite) is definable. As in (a) and (b), there is no reason to expect these sets to be $\emptyset$-definable, in general.

(d) Any set consisting of a single algebraic number is $\emptyset$-definable in the ordered field $\mathbb{R}$. For example, the formula

$$(0 < x < 1) \wedge (x^3 - 3x + 1 = 0)$$

defines a unique algebraic number. It follows that any finite set of algebraic numbers is $\emptyset$-definable, as is the complement of such a set.

(e) Julia Robinson proved that $\mathbb{Z}$ and $\mathbb{N}$ are $\emptyset$-definable in the field of rationals. The defining formulas for these subsets are ingenious and sophisticated.

(f) An important theorem of number theory asserts that every natural number can be written as the sum of at most four squares of integers. Therefore, $\mathbb{N}$ is $\emptyset$-definable in the ring of integers.

Exercise 21. Show that the set of positive rationals is $\emptyset$-definable in the field of rationals. Conclude that every interval whose endpoints are rational numbers or $\pm\infty$ is $\emptyset$-definable in this field. What about an interval like $(0, \sqrt{2}) \bigcap \mathbb{Q}$? Is every interval of the form $(a, b) \bigcap \mathbb{Q}$, where a and b are reals, definable (with parameters) in the field of rationals?

Exercise 22. Show that, in contrast to Example 19(d), no nonempty set of transcendental numbers is $\emptyset$-definable in the field $\mathbb{R}$. You may use Example 12 from Section 5.4.

Exercise 23. Prove that the set $\{i\}$ is not $\emptyset$-definable in the field $\mathbb{C}$. Use the fact that the function $f(z) = \overline{z}$ is an automorphism of $\mathbb{C}$, that is, an isomorphism from $\mathbb{C}$ to itself. Here, $\overline{z}$ is the complex conjugate of z, as usual.

Julia Robinson (1919–1985) had more than her share of adversity in her childhood. Her mother died when she was two, and at age nine Julia had serious episodes of scarlet fever and then rheumatic fever that kept her out of school for over two years. The family's savings were wiped out in the great depression, and her father committed suicide in 1937. However, Julia was able to persist in her studies, and eventually she earned her bachelors and masters degrees from U. C. Berkeley, and her PhD from Princeton, all in mathematics. Along the way she married Raphael Robinson, one of her professors at Berkeley. Thus, for some time, there were three logicians named Robinson in the United States: Julia, Raphael, and Abraham.

Robinson made several significant contributions to foundations, mostly in model theory. In her dissertation, supervised by Tarski, she proved the $\emptyset$-definability of $\mathbb{Z}$ in $\mathbb{Q}$, mentioned above. In Section 3.4 we mentioned her crucial work on Hilbert's tenth problem. She also did important research in algebraic model theory and recursion theory. Julia Robinson received many honors. She was the first woman to be elected to the National Academy of Sciences in 1976, and became the first woman president of the American Mathematical Society in 1982. However, she made it clear that she wished to be remembered for her achievements as a mathematician, independently of her gender.

Exercise 24. Prove that the set of definable sets (with or without parameters) in any structure $\mathfrak{A}$ forms an **algebra** of subsets of A, meaning a collection of subsets of A that is closed under **Boolean combinations** (finite unions and intersections, and complementation).

By De Morgan's laws (translated from logic to set operations), a collection of subsets of a set A that is closed under finite unions and complementation, or closed under finite intersections and complementation, must be an algebra.

We are now ready to define saturation. To keep the definition simple, we will not give it in full generality, which would allow κ to equal $\aleph_0$ and/or the language of $\mathfrak{A}$ to be uncountable.

Definition. Let κ be an uncountable cardinal, and let $\mathfrak{A}$ be a structure for a countable language. Then $\mathfrak{A}$ is **κ-saturated** if every collection of fewer than κ definable subsets of A is resilient.

If $\mathfrak{A}$ is $Card(A)$-saturated, then we simply say that $\mathfrak{A}$ is **saturated**. It is implicit in this definition that a saturated structure must be uncountable.

Let $\mathfrak{A}$ be any infinite structure. Consider the collection of all sets of the form $A - \{b\}$, where $b \in A$. By Example 19(c), these subsets of A are all definable. Also, this collection has the finite intersection property, but the intersection of the entire collection is empty. In other words, this collection is not resilient, so $\mathfrak{A}$ cannot be $Card(A)^+$-saturated. Therefore, a saturated structure is as saturated as it can possibly be. For this reason, some authors call saturation "full saturation."

In a linear ordering, it is clear that every open interval and every open ray is definable. So if a dense unbounded ordering is κ-saturated, then the conditions of Proposition 5.25 hold. (Conversely, the conditions of Proposition 5.25 imply κ-saturation, but this is less obvious. See Proposition 5.4.2 of [CK] for the proof.) Therefore, from the discussion following Theorem 5.24, we can deduce that any two saturated dense unbounded orderings of the same cardinality are isomorphic. This result has an important generalization to all types of first-order structures:

Theorem 5.26. *Any two saturated, elementary equivalent structures of the same cardinality are isomorphic.*

This theorem can be rephrased as a uniqueness result: a complete theory has, up to isomorphism, at most one saturated model of any given cardinality. The proof is a sophisticated back-and-forth argument that traces its lineage all the way back to the proof of Theorem 5.24.

However, the usefulness of this theorem is diminished by the fact that saturated structures are quite hard to come by. For instance, Example 18(c) and Exercise 19 tell us that the ordering $(\mathbb{R}, <)$ is not even $\aleph_1$-saturated. More generally, the existence of saturated structures cannot be proved in ZFC. The main positive result is that if κ is an infinite cardinal and T is a consistent theory with $Card(T) \leq \kappa$, then T has a κ^+-saturated model of cardinality 2^κ [CK, Lemma 5.1.4]. The continuum hypothesis says that $2^{\aleph_0} = \aleph_0{}^+$, so we can prove the existence of lots of saturated models (of cardinality $\aleph_1$) in ZFC $+$ CH.

To eliminate the need to assume CH for important results in model theory, the concept of a special structure was devised. For the record, here is the definition of this concept. However, I advise you not to get bogged down by the abstractness of this definition. Specialness should be thought of as a convenient, minor adaptation of the notion of saturation.

Definition. A structure $\mathfrak{A}$ is called **special** if it is the union of an elementary chain $\{\mathfrak{A}_\kappa \mid \kappa < Card(A)\}$, where the subscript κ is restricted to infinite cardinals, and each $\mathfrak{A}_\kappa$ is κ^+-saturated.

Every saturated structure is special, trivially—just let each $\mathfrak{A}_\kappa$ be $\mathfrak{A}$ itself. But it is easier for a structure to be special than saturated. More precisely, the following can be proved in ZFC:

Theorem 5.27.

(a) *Every theory with an infinite model has arbitrarily large special models.*

(b) *If a theory in a countable language has an infinite model, then it has a special model of cardinality $\beth_\omega$. ($\beth_\alpha$ is defined in Appendix C.)*

(c) *Theorem 5.26 holds with "saturated" replaced by "special."*

We now return to proving that specific theories are complete. We will present the two most famous results of this type.

Let RCOF denote the first-order theory of a real-closed ordered field, as described in the introduction to this chapter. By definition, a real-closed ordered field is a type of ordered ring. It is also possible to define and axiomatize the notion of a **real-closed field**, with no reference to an ordering. To do this, first define a field to be **formally real** if -1 cannot be written as a sum of squares. In a first-order theory, this requires an infinite axiom schema: for each n, the statement that -1 is not the sum of n squares. Essentially, this property replaces the ordering axioms. More precisely, a field is formally real if and only if it is **orderable**, meaning that an order can be defined on it that makes it into an ordered field. For instance, it's obvious that the fields $\mathbb{Q}$ and $\mathbb{R}$ are formally real, while $\mathbb{C}$ is not.

In addition to the axioms of a formally real field, include the schema that every polynomial of odd degree has a zero, as in RCOF. Finally, instead of saying that every nonnegative number has a square root, say that every number or its negative has a square root. Let RCF denote the resulting theory.

Exercise 25. Show that if x is any nonzero number in a real-closed field, exactly one of the numbers x and $-x$ has a square root.

For most purposes, the theories RCF and RCOF are interchangeable. In terms of structures, there is a simple correspondence between real-closed fields and real-closed ordered fields: every real-closed ordered field becomes a real-closed field simply by dropping the ordering relation, and every real-closed field can be turned into a real-closed ordered field in a unique way. To do this, define a number to be positive if it is nonzero and has a square root. So the unique, implicit ordering of any real-closed field is $\emptyset$-definable.

Exercise 26.

(a) Using the facts stated in the previous paragraph, show that RCOF is a **conservative extension** of RCF, meaning that it's an extension of RCF but proves no additional theorems in the (smaller) language of RCF.

(b) Show that RCF is complete if and only if RCOF is complete.

However, we will see in the next section that there are some notable differences between the theories RCF and RCOF.

Theorem 5.28 (Tarski). *The theories RCF and RCOF are complete.*

Proof. Unlike the theory of a dense unbounded ordering, these theories are not categorical in any infinite power. But the following weaker result holds: any two special real-closed fields of the same cardinality are isomorphic. This result, restricted to saturated real-closed fields, first appeared in [EGH, Theorem 2.1], using a classic back-and-forth argument. The proof for cardinality $\aleph_1$ is also given in Theorem 5.4.4 of [CK]. The generalization to special real-closed fields is not difficult. What makes this back-and-forth argument more complex than that of Theorem 5.24 is that every time we extend the isomorphism to one more member of the domain or range, we must also extend it to a real-closed field containing the new member.

Now, assuming that RCF is not complete, we can form two different complete extensions T_1 and T_2 of RCF. Then, since all models of T are infinite, T_i has a special model $\mathfrak{A}_i$ of cardinality $\beth_\omega$, for $i = 1, 2$. By the previous paragraph we would have $\mathfrak{A}_1 \cong \mathfrak{A}_2$, which would imply $T_1 = T_2$, a contradiction. ■

Exercise 27. Using this theorem and some results from Chapter 4, prove that the set of natural numbers is not $\emptyset$-definable in the ordered field of real numbers. (In the next section, we will see that $\mathbb{N}$ isn't even definable with parameters in this ordered field.)

The proof of Theorem 2.1 of [EGH] also yields the following:

Corollary 5.29. *In order for a real-closed (ordered) field to be κ-saturated, it suffices for it to be κ-saturated as a linear ordering.*

We now turn our attention to the other major category of "closed" fields:

Notation. Let ACF denote the first-order theory of an algebraically closed field. In addition to the field axioms, this theory requires an infinite axiom schema: for each positive integer n, an axiom that says every polynomial of degree n has a zero.

Also, let ACF_k denote the first-order theory of an algebraically closed field of characteristic k. Here, k may be 0 or any prime number. For example, ACF_3 is ACF plus the axiom $1 + 1 + 1 = 0$, which we may abbreviate as $3 = 0$ (although $\overline{3} = 0$ or even $\overline{3} = \overline{0}$ would be more correct). ACF_0 is ACF plus the axioms $p \neq 0$ for every prime p.

Theorem 5.30. *Each of the theories* ACF_k *is complete.*

Proof. This result is older and perhaps simpler than the previous theorem. The key lemma, due to Ernst Steinitz, is that an algebraically closed field is uniquely determined (up to isomorphism) by its characteristic and its **transcendence degree**—the cardinality of a maximal set of independent transcendental elements over its smallest subfield. But for uncountable fields, the transcendence degree is just the cardinality of the field itself. In other words, for each particular characteristic, the theory we are considering is categorical in every uncountable power. Therefore, by the Łoś–Vaught test, it is complete. ■

To demonstrate the power of this completeness result, here is an interesting fact of "ordinary mathematics" that was first proved by Ax, using Theorem 5.30. The proof requires somewhat more knowledge of algebraic concepts than we have been assuming.

Theorem 5.31. *Let* $n \in \mathbb{Z}^+$ *and let* K *be an algebraically closed field. Then every one-to-one polynomial function from* K^n *to* K^n *is onto.*

Proof. First we consider the special case where K is the algebraic closure of $\mathbb{Z}_p$ for some prime p. Assume f is such a polynomial, and let $\vec{y} = (y_1, y_2, \dots, y_n) \in K^n$. We must show that there exists $\vec{x} \in K^n$ such that $f(\vec{x}) = \vec{y}$. Now, if L is a finite field and a is algebraic over L, then the field $L(a)$ is also finite. (Here, $L(a)$ denotes the smallest extension field of L that contains a.) Therefore, since every member of K is algebraic over every subfield of K, it follows by induction that every finite subset of K is contained in a finite subfield of K. So let K_0 be a finite subfield of K that contains $y_1, y_2, \dots, y_n$ as well as all the coefficients of f. Then the restriction of f to ${K_0}^n$ is a one-to-one function from the finite set ${K_0}^n$ to itself. Therefore, there exists $\vec{x} \in {K_0}^n$ such that $f(\vec{x}) = \vec{y}$. So the theorem holds for fields of this type.

Now let a prime p and $n \in \mathbb{Z}^+$ be fixed. For each $m \in \mathbb{Z}^+$, there is a first-order sentence P_m that expresses "Every one-to-one polynomial with n input variables and n output variables, of degree $\leq m$, is onto." (Example: the polynomial $f(x, y) = (x^3 - 5xy, -2x + x^3y^2)$ has two input variables, two output variables, and degree 5.) We have just shown that each P_m is true in the algebraic closure of $\mathbb{Z}_p$, which is a model of ACF_p. So $\mathrm{ACF}_p \nvdash \sim \mathrm{P}_m$. Therefore, since ACF_p is complete, $\mathrm{ACF}_p \vdash \mathrm{P}_m$. Since this holds for every m, the theorem is established for algebraically closed fields of prime characteristic.

It remains to prove the characteristic 0 case. Again, let n be fixed. It will suffice to show that $\mathrm{ACF}_0 \vdash \mathrm{P}_m$ for every m. Assume this is false. Then $\mathrm{ACF}_0 \vdash \sim \mathrm{P}_m$ for some m, since ACF_0 is complete. So consider some proof of $\sim \mathrm{P}_m$ from ACF_0. Only a finite number of axioms of the form $p \neq 0$ are used in this proof, so we can find a prime q that is larger than all these primes p. Then all of these sentences of the form $p \neq 0$ are theorems of ACF_q. But this implies that $\sim \mathrm{P}_m$ is a theorem of ACF_q, contradicting the previous paragraph. ■

Corollary 5.32. *Let $n \in \mathbb{Z}^+$. Then every one-to-one polynomial function from $\mathbb{C}^n$ to $\mathbb{C}^n$ is onto.*

There are a few other interesting examples of complete theories of an algebraic nature. The theory of an abelian group in which every element has order p (for a fixed prime p) is categorical in every power in which it has a model. Therefore, the theory of an infinite group of this type is complete. Another example is the theory of a nontrivial divisible torsion-free abelian group, which is categorical in every uncountable power. Outside of algebra, there are few "substantial" complete theories. This is partly explained by Gödel's incompleteness theorem: an axiomatizable theory that includes Peano arithmetic cannot be complete.

We have now seen examples of various combinations of categoricity. The theory of a dense unbounded ordering is $\aleph_0$-categorical but not categorical in uncountable powers. For the theory of an algebraically closed field of fixed characteristic, it's the opposite. The theory of a real-closed field is not categorical in any infinite power, even though it is complete. And the previous paragraph mentioned a theory that is categorical in every power. In terms of infinite models of theories in countable languages, there are no other possibilities: a deep result known as **Morley's theorem** shows that if a theory in a countable language is categorical in one uncountable power, then it is categorical in every uncountable power.

5.6 Quantifier elimination

If a set (or relation or function) is definable in a structure, one important measure of the complexity of the set is based on the quantifier complexity of some formula that defines it. Thus, we can refer to sets that are **quantifier-free definable** (also called **Σ_0-definable** or **Π_0-definable**), **Σ_2-definable**, **Π_5-definable**, etc. For short, one simply refers to Σ_n sets and Π_n sets. Somewhat imprecisely, we will sometimes use these same symbols to denote the corresponding collections of definable sets. In other words, "$A \in \Sigma_n$" means the same thing as "A is Σ_n." A set that is both Σ_n and Π_n is called a Δ_n set. This "hierarchy" of complexity, devised by Kleene, provides one of the main approaches to understanding the structure of definable sets. There is also a similar but more limited hierarchy of $\emptyset$-definable sets. Here are a few simple facts about these hierarchies:

Proposition 5.33. *For any structure $\mathfrak{A}$:*

(a) *Σ_n (that is, the collection of Σ_n sets) is closed under finite unions and intersections. So is Π_n.*

(b) *The Δ_n sets form an algebra.*

(c) *Every Boolean combination of Σ_n sets is Δ_{n+1}.*

Proof. The proof of (a) is straightforward, using ideas mentioned in the proof of Theorem 1.5. (See the exercise below.) Part (b) follows immediately from (a). To prove (c), first note that the insertion of “dummy quantifiers” (also mentioned in Section 1.6) guarantees that every Σ_n set is Δ_{n+1}. But (b) says that the Δ_{n+1} sets are closed under Boolean combinations, so (c) follows. ■

Exercise 28. Prove this proposition in more detail, especially (a).

Thus the hierarchy of definable subsets of a structure looks like Figure 5.1, where each line indicates inclusion from left to right. Here, Γ_n denotes the Boolean combinations of Σ_n sets, but this notation is not standard. Clearly, every Π_n set is a Boolean combination of Σ_n sets (specifically, the complement of a Σ_n set), and vice-versa.

Example 20. The ordering on $\mathbb{R}$ is $\emptyset$-definable in the field of reals, since

$$x \leq y \leftrightarrow \exists z(x + z \cdot z = y)$$

is true in this structure, for any reals x and y. So, more specifically, the ordering relation is Σ_1, without parameters. We will soon see that $\leq$ and $<$ are not Σ_0, even with parameters.

Example 21. Sets (including relations and functions) that are definable in the standard model of arithmetic $\mathfrak{N}$ are called **arithmetical**. Since every natural number is defined by a term (numeral), definability and $\emptyset$-definability are the same in $\mathfrak{N}$. Given that $\mathfrak{N}$ is a model of PA, it is clear that every set that is representable in PA is arithmetical. Therefore, every recursive set is arithmetical.

But where do recursive sets fit in the **arithmetical hierarchy**? By the main lemma used to solve Hilbert’s tenth problem (discussed in Section 3.4), every RE set is Σ_1. (This can also be obtained by first

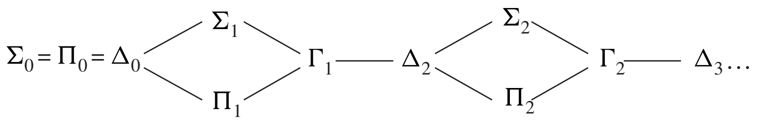

Figure 5.1. The hierarchy of definable subsets of a structure

showing that Kleene's T-predicate and the upshot function U are Σ_1, and then applying Theorems 3.5 and 3.10(e)). The converse of this is also true, since a Σ_0 set is obviously decidable. So a subset of $\mathbb{N}^k$ is RE if and only if it's Σ_1. Then, from Theorem 3.12, it follows that the recursive sets are precisely the Δ_1 ones. So recursive and RE sets form two of the lowest levels in the arithmetical hierarchy. (Some authors define the Σ_0 and Π_0 sets in $\mathfrak{N}$ to include all PR ones or even all recursive ones. This modification has no effect on the higher levels of the hierarchy, including Σ_1 and Δ_1.)

It is common to adjoin a superscript of 0 to the notation for the arithmetical hierarchy. Thus, recursive sets are Δ^0_1 (read "Delta-0-1"), and RE sets are Σ^0_1. This superscript indicates that the quantifiers allowed in the defining formulas are first-order, over *elements* of $\mathbb{N}$. (Logically, one could claim that this superscript is "one off," but that's just how it is.) In Section 6.5, we will define the analogous notation with a superscript of 1, as well as another approach to defining these hierarchies.

Definable sets of a structure are, for the most part, more "well behaved" than other sets. On the other hand, definable sets of high syntactic complexity are not usually easy to deal with. Therefore, model theorists are interested in theories whose models have particularly simple definable sets. The following definition describes the most desirable situation of this sort:

Definition. Let T be a theory in a language $\mathcal{L}$. We say that T has **quantifier elimination** if, for every $\mathcal{L}$-formula P, there is a quantifier-free $\mathcal{L}$-formula Q such that $T \vdash (\mathrm{P} \leftrightarrow \mathrm{Q})$.

Exercise 29. Prove that, in order for a theory to have quantifier elimination, it suffices for the definition to hold whenever P is a Σ_1 formula with just one quantifier.

Lemma 5.34. *Suppose T has quantifier elimination. Then, if* P *is not a sentence or the language of T has at least one constant symbol, the quantifier-free formula* Q *can be chosen to have only variables that are* *free in* P. *Otherwise,* Q *can be chosen to have v_0 (or any other specified v_i) as its only variable.*

Proof. Suppose T has quantifier elimination and P is a formula. So there's a quantifier-free formula Q such that $T \vdash (\mathrm{P} \leftrightarrow \mathrm{Q})$. Now, suppose some variable v_k is in Q but is not free in P. By the generalization theorem (Theorem 1.4), $T \vdash \forall v_k(\mathrm{P} \leftrightarrow \mathrm{Q})$, since T consists of sentences. Then, by the universal specification axiom of logic (as in Appendix A), $T \vdash (\mathrm{P} \leftrightarrow \mathrm{Q}')$, where Q' results from Q by replacing every occurrence of v_k by some free variable of P or a constant symbol, (or by v_0, if P is a sentence and there are no constant symbols). Since we can do this for each such v_k, the lemma is proved. ■

There is often a relationship between quantifier elimination and completeness for a theory, although neither of these properties implies the other in general. We will now illustrate this relationship by revisiting some of the complete theories discussed in the previous section.

Theorem 5.35. *The theory of a dense unbounded ordering has quantifier elimination.*

Proof. Let T be this theory. In its language $\mathcal{L}$, all atomic formulas are of the form $v_j = v_k$ or $v_j < v_k$. By an **arrangement** of a finite nonempty set of variables, we mean a conjunction of atomic formulas that precisely specifies the order of those variables in a consistent way. For example, the formula

$$(v_2 < v_7) \wedge (v_7 = v_{12}) \wedge (v_{12} < v_3)$$

is an arrangement of $\{v_2, v_3, v_7, v_{12}\}$.

The main lemma for this theorem, whose proof can be found in [CK] or [Mar], is that every quantifier-free formula Q of $\mathcal{L}$ is equivalent, in T (or even in the theory of a total ordering), to either a contradiction or a disjunction of arrangements of the free variables of Q. Using this lemma, we proceed as follows:

By Exercise 29, we can restrict our attention to a formula P of the form $\exists v_k \mathrm{R}$, where R is quantifier-free. If R is equivalent to a contradiction, then so is P, and we have $T \vdash (\mathrm{P} \leftrightarrow v_0 < v_0)$. On the other hand, if R is equivalent to a disjunction of arrangements $\mathrm{A}_1 \vee \mathrm{A}_2 \vee \ldots \vee \mathrm{A}_n$, then we obtain

$$T \vdash \mathrm{P} \leftrightarrow (\exists v_k \mathrm{A}_1 \vee \exists v_k \mathrm{A}_2 \vee \ldots \vee \exists v_k \mathrm{A}_n),$$

because the existential quantifier "distributes" over disjunctions in first-order logic. But in any dense unbounded ordering (and hence in T), a formula of the form $\exists v_k \mathrm{A}_i$ is equivalent to the new arrangement B_i obtained by deleting all conjuncts that mention v_k from A_i. Therefore P is equivalent, in T, to the quantifier-free formula $\mathrm{B}_1 \vee \mathrm{B}_2 \vee \ldots \vee \mathrm{B}_n$. ■

Corollary 5.36. *The theory of a dense unbounded ordering is complete.*

Proof. By the theorem and Lemma 5.34, every $\mathcal{L}$-sentence P is equivalent, in T, to a quantifier-free formula with only one variable v_0, which in turn must be equivalent to a contradiction or a disjunction of arrangements of v_0. But the only arrangement of v_0 is $v_0 = v_0$, and a disjunction formed from this equation is clearly equivalent to the same equation. So either $T \vdash (\mathrm{P} \leftrightarrow v_0 < v_0)$, which implies $T \vdash \sim \mathrm{P}$; or $T \vdash (\mathrm{P} \leftrightarrow v_0 = v_0)$, in which case $T \vdash \mathrm{P}$. ■

Thus we have an alternative proof of Theorem 5.24. The proof using quantifier elimination is more concrete; specifically, the quantifier elimination process is effective and therefore provides an actual decision procedure for the theory. This is the main advantage of quantifier elimination over more abstract methods such as saturation. In the words of [CK], "the method is extremely valuable when we want to beat a particular theory into the ground." (But this cute remark is certainly not intended as a criticism of this powerful technique.)

Exercise 17 in the previous section referred to three other complete theories involving dense orderings. Interestingly, none of these theories has quantifier elimination:

Exercise 30. Prove that the theory of a dense ordering with endpoints does not have quantifier elimination. Specifically, show that the formula $\exists v_1(v_1 < v_0)$, which says that v_0 is not the least element, has no quantifier-free equivalent. (Hint: Note that the "main lemma" referred to in the proof of Theorem 5.35 also applies to this theory.)

By contrast, the proof of Corollary 5.36 shows that in the theory of a dense unbounded ordering, you can't say anything nontrivial about a single element, such as the assertion that it's the least one.

In spite of this exercise, the method of quantifier elimination can still be used to prove these three theories are complete. For example, let T be the theory of a dense ordering with endpoints. Now consider the same theory, augmented by two constant symbols a and b and axioms "a is the least element" and "b is the greatest element." This theory T' can be shown to have quantifier elimination and to be complete, using arguments very similar to the proofs of Theorem 5.35 and Corollary 5.36. It is also clear that every model of T can be turned into a model of T' in a unique way, by interpreting a and b in the obvious way. Therefore, by the reasoning of Exercise 26, T' is a conservative extension of T, and T is also complete.

The argument just given is part of a general phenomenon: every theory has a conservative extension with quantifier elimination. However, the construction of this conservative extension is often complicated and gives little or no useful information about the original theory.

We've just seen three examples of complete theories without quantifier elimination. The next exercise and theorem illustrate the opposite phenomenon:

Exercise 31. Let $\mathcal{L}$ be the language with no relation or function symbols, and two constant symbols a and b. Let T be the $\mathcal{L}$-theory whose only proper axiom says that every element equals a or b.

(a) Show that T is not complete, by finding two models of T that are not elementarily equivalent.

(b) Show that T has exactly two complete extensions (up to equivalence of theories).

(c) Show that T has quantifier elimination. (Hint: Use Exercise 29, Proposition 1.1, and the fact that, in T, the negation of any equation involving a variable is equivalent to a positive combination of equations.)

Theorem 5.37. *The theory ACF of algebraically closed fields has quantifier elimination.*

The proof of this theorem is not terribly difficult, but it requires some nontrivial concepts from algebra and we will not give it. It can be found in [Mar]. From the previous section, we know that ACF is not complete. It has an infinite number of non-equivalent completions, the theories ACF_k. But ACF is "close to" being complete, in the sense that all that's required to complete it is to add some *quantifier-free* axioms (to specify the characteristic). Clearly, if it needed the addition of axioms with quantifiers to become complete, it couldn't have quantifier elimination.

Exercise 32. Use the previous theorem to give an alternative proof that the theories ACF_k are complete. (Hint: Use the facts about prime fields that are given in Appendix D. From them, prove that for each k, all models of ACF_k are elementarily equivalent.)

Corollary 5.38. *In any algebraically closed field, the definable sets are precisely the finite or cofinite subsets of the field.*

Proof. Let S be a definable subset of an algebraically closed field $\mathfrak{A}$. By the theorem, that means $S = \{x \in A : \mathfrak{A} \models \mathrm{Q}[g_x^0]\}$, for some quantifier-free formula Q and some assignment g. But in the language of field theory, an atomic formula must be an equation between two polynomials. These polynomials may have variables other than v_0, but under g these variables are replaced by particular members of A, which can be viewed as part of the coefficients of the polynomials. In other words, if Q is atomic, then $S = \{x \in A : p(x) = q(x)\}$, where p and q are polynomials in one variable with coefficients in A. By the fundamental theorem of algebra, such an equation either has a finite number of solutions or is true for all members of the domain. So since any quantifier-free Q is a Boolean combination of atomic formulas, S must be the result of taking (finite) unions, intersections, and complements of finite sets. Therefore, it is finite or cofinite. ■

In particular, this corollary tells us that very few sets are (first-order) definable in the field of complex numbers. The reals, the pure imaginary numbers, and the unit circle are examples of undefinable sets. However, all of these sets become definable if the complex conjugation function is added to the structure.

A theory is called **strongly minimal** if the only definable sets in its models are finite or cofinite. By Exercise 19(c), this means that its models have no more definable sets than absolutely necessary. So this corollary says that ACF is strongly minimal. The fact that ACF has quantifier elimination also gives useful information about the definable *relations* in an algebraically closed field: they are precisely the so-called **constructible** relations, an important notion in algebraic geometry. This is the sort of link that leads to applications of model theory outside of foundations.

We now turn our attention to real-closed fields. It is here that we will see an interesting distinction between the theories RCF and RCOF.

Proposition 5.39. *The theory RCF does not have quantifier elimination.*

Proof. The proof of Corollary 5.38 actually shows that in any field whatsoever, the Σ_0 sets (with parameters) are just the finite and cofinite sets. So, if RCF had quantifier elimination, these would be the only definable sets in the real-closed field $\mathbb{R}$. But the first example in this section shows that the positive reals are definable in the field of reals.

■

Theorem 5.40. *The theory RCOF has quantifier elimination.*

As with ACF, we will not provide the proof of this result because it requires a significant amount of algebraic machinery. This theorem is due to Tarski and was his original path to proving completeness:

Corollary 5.41. *The theories RCF and RCOF are complete, in their respective languages.*

Proof. Every real-closed ordered field has characteristic 0 and therefore has a subfield isomorphic to $\mathbb{Q}$. From this it can be shown, as in Exercise 32, that all models of RCOF are elementary equivalent, so RCOF is complete. The completeness of RCF then follows from Exercise 26(b). ■

So we have a situation here that fits the phenomenon described after Exercise 30. The theory RCF can't have quantifier elimination but, fortunately, it is not hard in this case to find a conservative extension that does. Then the completeness of both theories follows.

As with ACF, the fact that RCOF has quantifier elimination provides sharp information about the definable sets in $\mathbb{R}$ and other real-closed fields. Since the unique ordering of any real-closed field is definable in RCF, we will get the same definable sets whether or not the symbol $<$ is part of the language. But with this symbol included, the proof of Corollary 5.38 must be modified to allow atomic formulas of the form $p(x) < q(x)$, for polynomials p and q. And the solution set of such an inequality can be any finite union of intervals. Furthermore, when you take Boolean combinations of finite unions of intervals, what you get is finite unions of points and intervals. We have just more or less proved:

Corollary 5.42. *In any real-closed field, including $\mathbb{R}$, the definable sets are precisely the finite unions of points and intervals. In particular, the sets $\mathbb{N}$, $\mathbb{Z}$, and $\mathbb{Q}$ are not definable in the field of reals.*

So we see that RCF and RCOF are not strongly minimal, in contrast to ACF. In fact, this is obvious from the fact that their models are ordered (explicitly for RCOF, implicitly for RCF). If a theory T proves that some formula $\mathrm{P}(x, y)$ defines a total ordering, then every interval under that ordering must be definable in every model of T. Therefore, if T has any infinite models, it cannot be strongly minimal. Instead, we consider the following concept:

Definition. Let T be a theory in which it is provable that some binary relation symbol defines a total ordering. T is called **o-minimal** if the only definable sets in any model of T are finite unions of points and intervals.

So the previous corollary says that RCOF is o-minimal. A bit loosely, we can say the same for RCF. More generally, definable relations in a real-closed field are precisely the relations that are known as **semialgebraic**. As with algebraically closed fields, these model-theoretic results about real-closed fields are powerful tools for doing “ordinary mathematics.” For instance, the o-minimality of RCOF can be used to prove that every semialgebraic (first-order definable) function from $\mathbb{R}$ to itself is piecewise continuous.

5.7 Additional topics in model theory

This chapter has barely scratched the surface of the deep and rich subject of model theory. In this section we briefly touch on a few other aspects of the field.

Axiomatizable and nonaxiomatizable classes

The relation $\models$ provides a fruitful way of passing back and forth between sets of first-order sentences and classes of first-order structures.

Notation. If T is a set of $\mathcal{L}$-sentences, $Mod(T)$ denotes the class of all models of T. We also write $Mod(\mathrm{P})$, where P is a single sentence.

In the other direction, if $\mathcal{C}$ is a class of $\mathcal{L}$-structures, $Th(\mathcal{C})$ denotes the set of all $\mathcal{L}$-sentences that are true in every structure in $\mathcal{C}$. We also write $Th(\mathfrak{A})$, where $\mathfrak{A}$ is a single structure.

Note that $Th(\mathfrak{A})$ is automatically a complete theory. The converse, that every complete theory is of the form $Th(\mathfrak{A})$, is essentially a restatement of the harder direction of the completeness theorem.

Definition. A class of $\mathcal{L}$-structures is called **axiomatizable** (respectively, **finitely axiomatizable**, **recursively axiomatizable**) if it is $Mod(T)$ for some T (respectively, finite T, recursive T).

Clearly, a finitely axiomatizable class must be of the form $Mod(\mathrm{P})$, where P is the conjunction of all the sentences in the finite set T.

Recall, from Chapter 4, that for *theories* we are using "axiomatizable" as an abbreviation for "recursively axiomatizable." We do that because there is no other reasonable meaning for the phrase "axiomatizable theory." But for classes of structures, it makes sense to give these terms different meanings.

To illustrate this new terminology, let's use the compactness theorem to prove a fact that was mentioned in Section 1.5 (Example 17):

Proposition 5.43.

(a) *The class of all fields of finite characteristic is not axiomatizable.*

(b) *The class of all fields of characteristic zero is recursively axiomatizable but not finitely axiomatizable.*

Proof.

(a) Assume $Mod(T)$ consists of all fields of finite characteristic. Now let

$$T' = T \cup \{1+1 \neq 0, 1+1+1 \neq 0, 1+1+1+1+1 \neq 0, \dots\},$$

where there is an inequality for each prime number. Since the characteristic of every field is either 0 or a prime number, every finite subset of T' has a model but T' itself does not. This violates the compactness theorem.

(b) The standard axiomatization of this class of fields, consisting of the field axioms plus the schema shown in braces in the proof of part (a), is certainly recursive. Now assume this class of fields is finitely axiomatizable. Then, for some sentence P, $Mod(\mathrm{P})$ is the class of all fields of characteristic zero. But then if T consists of $\sim$ P and all the field axioms, $Mod(T)$ is the class of all fields of finite characteristic, in violation of part (a). ■

Exercise 33.

(a) Prove that the class of all finite groups is not axiomatizable. (Hint: Recall, from Section 1.3, that statements of the form "There are at least n elements" and "There are exactly n elements" can be formalized in first-order logic, for each positive integer n.)

(b) Prove that the class of all infinite groups is recursively axiomatizable but not finitely axiomatizable.

Similarly, the class of all finite fields and the class of all finite rings are not axiomatizable, while the class of all infinite fields and the class of all infinite rings are recursively axiomatizable but not finitely axiomatizable. There is a great variety of similar but more sophisticated examples involving algebraic theories:

Proposition 5.44. *The class of all divisible groups is recursively axiomatizable but not finitely axiomatizable.*

Proof. As in Appendix D, the class of divisible groups is axiomatized by the recursive set of axioms $T \cup \{\mathrm{P}_n \mid n \in \mathbb{Z}^+\}$, where T is the usual set of axioms for a group and P_n is the statement $\forall x\, \exists y(y^n = x)$. To prove that this class is not finitely axiomatizable, it will suffice to prove, as in the previous proposition, that the class of groups that are not divisible isn't even axiomatizable.

So assume that the class of non-divisible groups has an axiomatization T'. Then the theory $T' \cup \{\mathrm{P}_n \mid n \in \mathbb{Z}^+\}$ is inconsistent, since only a divisible group can satisfy all the P_n's. Therefore, $T' \cup \{\mathrm{P}_1, \mathrm{P}_2, \dots, \mathrm{P}_k\}$ is inconsistent for some k. But then let G_k be the set of all rational numbers whose denominators, in lowest terms, have no prime factors greater than k. It is routine to verify that G_k, under addition (so that y^n becomes ny), is a model of this theory. By this contradiction, we're done. (Note the tacit use of compactness here.) ■

Exercise 34.

(a) Complete this proof by by showing that $(G_k, +)$ is a group, is not divisible, and does satisfy the statements $\mathrm{P}_1, \mathrm{P}_2, \dots, \mathrm{P}_k$.

(b) Prove that the class of all torsion-free groups is recursively axiomatizable but not finitely axiomatizable. You will need to find a group in which every element except the identity has finite order greater than a given number k.

(c) Prove that the class of all torsion groups is not axiomatizable. (Hint: Consider the discussion of nonstandard analysis in Section 5.3, in which a new constant symbol is added to a language.)

Even when $\mathcal{C}$ is a nonaxiomatizable class, we call $\mathit{Th}(\mathcal{C})$ the first-order theory of $\mathcal{C}$. For instance, if we refer to the first-order theory of finite groups, we do not mean a particular list of axioms that define finite groups precisely, because there is no such list of axioms. Rather, we mean $\mathit{Th}(\mathcal{C})$, where $\mathcal{C}$ is the class of all finite groups. Exercise 37(c) below will highlight this point.

The study of theories of nonaxiomatizable classes has led to some significant research. For instance, the class of all finite groups and the class of all finite fields are both nonaxiomatizable, by identical reasoning. But while the first-order theory of finite groups is not decidable, a very deep result due to James Ax [Ax] shows that the first-order theory of finite fields is decidable. Ax's result is particularly striking because many simpler (axiomatizable or even finitely axiomatizable) theories, such as group theory, field theory, and Peano arithmetic, are not decidable.

Notation. Recall that $\mathit{Thm}(T)$ denotes the set of theorems of any theory T. We will also write $\overline{\mathit{Thm}}(T)$ for the set of *sentences* that are theorems of T.

The following exercises might look substantial, but for the most part they are quite trivial, once the new terminology is "unraveled."

Exercise 35. Let P be a sentence, and T a theory. Prove:

(a) $T \models \mathrm{P}$ if and only if $\mathrm{P} \in \mathit{Th}(\mathit{Mod}(T))$. Therefore, Gödel's completeness theorem (restricted to sentences) can be stated in the form $\overline{\mathit{Thm}}(T) = \mathit{Th}(\mathit{Mod}(T))$.

(b) $\mathit{Mod}(T) = \mathit{Mod}(\overline{\mathit{Thm}}(T))$.

(c) If $\mathcal{C}$ is any class of structures, $\mathit{Th}(\mathcal{C})$ must be closed under $\vdash$, that is, $\overline{\mathit{Thm}}(\mathit{Th}(\mathcal{C})) = \mathit{Th}(\mathcal{C})$.

Exercise 36. Prove that the following are equivalent, for any theories T_1 and T_2:

(a) T_1 and T_2 are equivalent, that is, $Thm(T_1) = Thm(T_2)$.

(b) $\overline{Thm}(T_1) = \overline{Thm}(T_2)$.

(c) $Mod(T_1) = Mod(T_2)$.

Exercise 37.

(a) Fill in the blank: if $\mathfrak{A}$ is any structure, then $Mod(Th(\mathfrak{A}))$ is the class of all structures that are ______ to $\mathfrak{A}$.

(b) Let $\mathcal{C}$ be an axiomatizable class of structures, such as the class of all groups or the class of all infinite rings. What is $Mod(Th(\mathcal{C}))$?

(c) Now let $\mathcal{C}$ be a nonaxiomatizable class of structures. What can be said about $Mod(Th(\mathcal{C}))$? For instance, suppose that $\mathcal{C}$ is the class of all finite fields. Does $Mod(Th(\mathcal{C}))$ include all finite fields? Does it include any infinite fields?

Thus we see that the operators *Mod* and *Th* are inverses to a limited extent: Exercise 35(a) tells us that $Th \circ Mod$ is the identity on sets of sentences that are closed under $\vdash$. Exercise 37(b) tells us that $Mod \circ Th$ is the identity on axiomatizable classes of structures.

By Exercise 37(c), there are **pseudofinite fields**, meaning infinite fields that are models of the first-order theory of finite fields. This seems odd, but actually it requires some thought to see why there are any infinite fields that are *not* pseudofinite. To establish this, we need to find a first-order sentence that is true in every finite field but not in every field. One key to this is the fact that any function from a finite set to itself is one-to-one if and only if it is onto. Of course, this is not true for infinite sets.

Example 22. Let's show that the field of real numbers is not pseudofinite. The polynomial $f(x) = x^3 - x$ (technically $x \cdot x \cdot x - x$, in the first-order language of a field), is a simple example of a function from $\mathbb{R}$ onto itself that is not one-to-one. We can express that f is onto in the usual way: $\forall y \exists x[y = f(x)]$. Similarly, the assertion that f is one-to-one can be formalized as $\forall x, y[f(x) = f(y) \rightarrow x = y]$.

Now let P be the formal statement that says that f is one-to-one if and only if it is onto. P is true in every finite field, but not in $\mathbb{R}$. So $\mathbb{R}$ is not a model of the first-order theory of finite fields.

Exercise 38. Show that the fields $\mathbb{Q}$ and $\mathbb{C}$ are not pseudofinite. You can use the same reasoning as in the previous exercise, but you might need to find different polynomials.

Unfortunately, it is not possible to give an example of a pseudofinite field without using more sophisticated methods, such as ultraproducts. By the way, readers of [Ax] must digest quasifinite fields, pseudofinite fields, and hyperfinite fields. But the effort is worthwhile: the main result about finite fields has some important algebraic consequences, such as a decision procedure for determing whether a system of Diophantine equations has, for all primes p, a solution modulo p (or a p-adic solution).

Stone spaces

Why is the compactness theorem called that? Its content certainly resembles a statement of topological compactness, and indeed we can make this precise: given a first-order language $\mathcal{L}$, let $\mathcal{S}$ be the set of all theories of the form $Th(\mathfrak{A})$, where $\mathfrak{A}$ is an $\mathcal{L}$-structure. Note that each member of $\mathcal{S}$ is therefore a complete $\mathcal{L}$-theory. Conversely, the completeness theorem says that each complete $\mathcal{L}$-theory that is closed under $\vdash$ is in $\mathcal{S}$.

Intuitively, $\mathcal{S}$ can also be thought of as the collection of all equivalence classes of $\mathcal{L}$-structures under $\equiv$. But these equivalence classes would be proper classes, and it is “iffy” to work with a collection of proper classes. This can be rectified by considering only $\mathcal{L}$-structures whose cardinality is no greater than that of $\mathcal{L}$.

For each sentence P of $\mathcal{L}$, let $U_{\mathrm{P}} = \{T \in \mathcal{S} \mid \mathrm{P} \in T\}$. The collection of all the U_{P}'s is closed under finite intersections, since $U_{\mathrm{P}} \cap U_{\mathrm{Q}} = U_{\mathrm{P} \wedge \mathrm{Q}}$. Therefore we can use the U_{P}'s as a basis for a topology on $\mathcal{S}$, called the **Stone space** of $\mathcal{L}$. So an open set in this space is any union of U_{P}'s.

Note that $U_{\mathrm{P}} = \mathcal{S} - U_{\sim \mathrm{P}}$, so each U_{P} is **clopen** (closed and open). Therefore the Stone space of $\mathcal{L}$ is **totally disconnected**, meaning that it has a basis of clopen sets. Recall that the only sets that must be clopen in a topological space are the whole space and $\emptyset$, and in **connected** spaces (such as $\mathbb{R}^n$), these are the only clopen sets. The Stone space is also Hausdorff (see the next exercise). Finally, it is not hard to show that the compactness theorem is equivalent to the compactness of the Stone space of first-order languages.

Exercise 39.

(a) Verify that the Stone space of $\mathcal{L}$ is Hausdorff. That is, any two distinct points are contained in a pair of disjoint open sets.

(b) Without assuming the completeness theorem, prove that the compactness theorem is equivalent to the compactness of the Stone space of all first-order languages.

In general topology, a Stone space is any compact, Hausdorff, totally disconnected space. The most well-known Stone space is the **Cantor set**, which has several homeomorphic (topologically equivalent) versions. It can be defined simply as the set of real numbers between 0 and $1/9$ (inclusive) that have a decimal expression using only 0's and 1's, under the relative topology induced by $\mathbb{R}$. A version that is easier to visualize is the set of real numbers between 0 and 1 that have a base 3 expression using only 0's and 2's. (See Figure 5.2. Here, the Cantor set is the complement in [0, 1] of the union of all the sets A_n.) The Cantor set is also the product space $\{0, 1\}^{\mathbb{N}}$, where $\{0, 1\}$ is given the discrete topology. Finally, any set of the form $\mathcal{P}(A)$, where A is denumerable, has a natural Cantor set topology induced by the standard bijection between $\{0, 1\}^{\mathbb{N}}$ and $\mathcal{P}(\mathbb{N})$. We will say more about the Cantor set in Section 6.5.

Tarski's undefinability theorem

The main syntactic version of Tarski's truth theorem was presented in Chapter 4 (Theorem 4.14). Here is a stronger semantic result that is also sometimes referred to as Tarski's truth theorem:

Theorem 5.45 (Tarski's Undefinability Theorem). *Let $\mathfrak{N}$ be the standard model of arithmetic. Then Th$(\mathfrak{N})$ (as a set of Gödel numbers) is not arithmetical.*

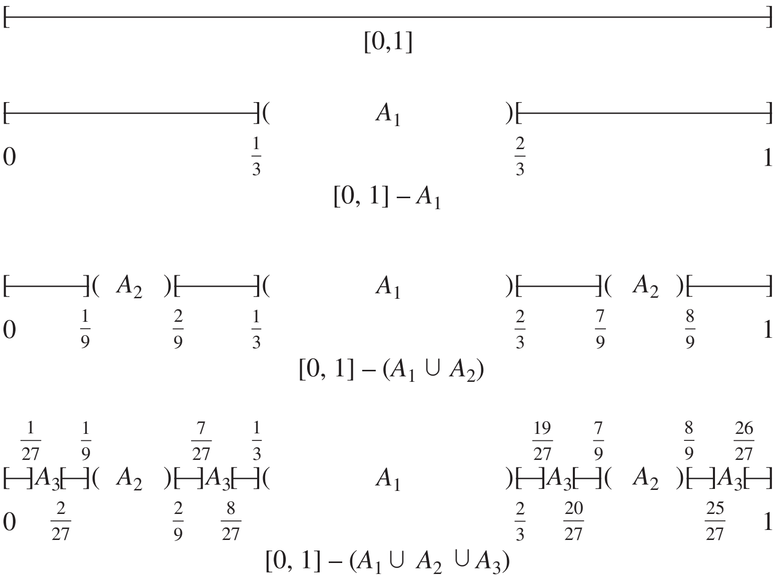

Figure 5.2. The first three stages in the construction of the Cantor set

Proof. Assume on the contrary that $Th(\mathfrak{N})$ is definable in $\mathfrak{N}$, so we have $Th(\mathfrak{N}) = \{n \in \mathbb{N} \mid \mathfrak{N} \models \mathrm{R}(n, b_1, \ldots, b_k)\}$, for some formula R and some $b_1, \ldots, b_k \in \mathbb{N}$. Replacing each b_i by its numeral changes R into a formula $\mathrm{P}(n)$ with one free variable. (In other words, definability in $\mathfrak{N}$ implies $\emptyset$-definability, as we mentioned earlier.) So, for any sentence Q, $\#\mathrm{Q} \in Th(\mathfrak{N})$ (that is, $\mathfrak{N} \models \mathrm{Q}$) if and only if $\mathfrak{N} \models \mathrm{P}(\overline{\#\mathrm{Q}})$.

Now apply the fixed point lemma (Lemma 4.7) to $\sim$ P, with T being PA, to obtain a sentence Q. So $\mathrm{PA} \vdash [\mathrm{Q} \leftrightarrow \sim \mathrm{P}(\overline{\#\mathrm{Q}})]$. Since $\mathfrak{N}$ is a model of PA, we have $\mathfrak{N} \models [\mathrm{Q} \leftrightarrow \sim \mathrm{P}(\overline{\#\mathrm{Q}})]$. This clearly contradicts the previous paragraph. ■

Exercise 40. Assuming Tarski's undefinability theorem (but not the fixed point lemma), prove Theorem 4.14 for any theory T that is satisfied by $\mathfrak{N}$.

Since we know, from the previous section, that all recursive relations are definable in $\mathfrak{N}$, Tarski's undefinability theorem is also a strengthening of Corollary 4.12.

Second-order model theory

In Section 1.6 we described second-order logic. Now that we know something about model theory, we can enter into a more meaningful discussion of this subject.

Let $\mathcal{L}$ be a first-order language. We will denote the second-order language based on $\mathcal{L}$ by $\mathcal{L}^+$. Recall that the essential new ingredients of $\mathcal{L}^+$ are an infinite list of n-ary relation variables for each $n \in \mathbb{Z}^+$, terms based on these relation variables, and quantification over relation variables.

Now let $\mathfrak{A}$ be any $\mathcal{L}$-structure. Then $\mathfrak{A}$ can also be viewed as an $\mathcal{L}^+$-structure, and indeed there is no need to make a distinction between $\mathcal{L}$-structures and $\mathcal{L}^+$-structures. When $\mathfrak{A}$ is viewed as an $\mathcal{L}^+$-structure, a quantified n-ary relation variable is naturally interpreted as ranging over *all* possible subsets of A^n. This is an essential requirement of second-order semantics.

The next step is to define the second-order versions of the two relations denoted by $\models$. Of course, the definition of an assignment g becomes more complex in the second-order setting: an assignment must not only map each individual variable of $\mathcal{L}^+$ to an element of A, but must also map each n-ary relation variable of $\mathcal{L}^+$ to a subset of A^n. With this modification, the definition of the second-order version of $\mathfrak{A} \models \mathrm{P}[g]$ becomes straightforward. Then we can define $T \models \mathrm{P}$, where T and P consist of $\mathcal{L}^+$-formulas, just as in Section 5.2.

Section 5.3 presented the three main results of first-order model theory: the completeness theorem, the compactness theorem, and the LST theorem. Let us now consider these theorems in the context of second-order logic.

One standard, concise version of the completeness theorem is that $T \vdash \mathrm{P}$ if and only if $T \models \mathrm{P}$. But model theorists tend to view this theorem in a different light. To a model theorist, the basic ingredients of "a logic" are a formal language (with connectives, perhaps quantifiers, and grammatical rules for the formation of formulas), and some reasonable definition of structures and models for that language. In other words, the semantic notions $\mathfrak{A} \models \text{P}$ and $T \models \text{P}$ are basic, necessary parts of a logic; the syntactic notion $T \vdash \text{P}$ is not.

In this context, the completeness theorem for a logic becomes the assertion that *there exists* a notion of deduction, based on some clear-cut, mechanical procedure for manipulating formulas of the language, such that $\vdash$ and $\models$ coincide. In this sense, the completeness theorem is false for second-order logic, and for most of the other important logics that extend first-order logic. In other words, there is no reasonable way to axiomatize second-order logic in such a way that the provable statements are precisely the valid ones. We will prove this shortly.

In contrast to the completeness theorem, the compactness theorem and the LST theorem make no mention of $\vdash$ in their statements.

Proposition 5.46. *The compactness theorem fails for second-order logic.*

Proof. Let T be the following second-order theory: for each n, T includes the (first-order) sentence that states that there are at least n distinct elements. Also, let R be a binary relation variable. T includes the statement that, for every R, if R defines a total ordering, then there is a least element and a greatest element in this ordering. This statement is true in every finite structure but in no infinite structures. Therefore, every finite subset of T is satisfiable, but T is not. ■

Corollary 5.47. *The completeness theorem fails for second-order logic.*

Proof. Recall that the compactness theorem for first-order logic is a simple corollary of the completeness theorem for first-order logic. If there were a notion of deduction for second-order logic (in which proofs are finite) making the completeness theorem hold, then the compactness theorem would hold as well. ■

Proposition 5.48. *The LST theorem fails for second-order logic.*

Proof. The defining properties of a complete ordered field can be stated (as a single sentence, if desired) in second-order logic. Every model of this sentence is isomorphic to the field $\mathbb{R}$ and so has cardinality $2^{\aleph_0}$.

■

An alternative proof of this proposition would be to consider second-order Peano arithmetic, as described in Section 1.5. Every model of this theory is isomorphic to the usual structure with domain $\mathbb{N}$, and so is denumerable.

So none of the major theorems of first-order logic holds for second-order logic. This situation highlights the specialness of first-order logic. In fact, a theorem due to Tom Lindström asserts that first-order logic is the only "reasonable" logic that satisfies both the compactness and LST theorems. Here, "reasonable" means that the relation $\models$ satisfies the obvious defining conditions with respect to the standard connectives and quantifiers. (For example, $\mathfrak{A} \models \mathrm{P} \wedge \mathrm{Q}$ should hold if and only if $\mathfrak{A} \models \mathrm{P}$ and $\mathfrak{A} \models \mathrm{Q}$.) Of course, this special status of first-order logic does not necessarily make first-order logic superior to other logics. On the contrary, the LST theorem demonstrates a serious limitation of first-order logic and is one of the primary reasons for considering stronger logics.