IV.23 Logic and Model Theory

David Marker

1 Languages and Theories

Mathematical logic is the study of formal languages that are used to describe mathematical structures and what these can tell us about the structures themselves. We can learn a lot about a formal language by investigating which of its sentences are true for the structure it describes, and we can learn a lot about the structure by investigating the subsets of it that can be defined using the language. In this article, we shall see several examples of languages and the structures that they are used to describe. We shall also see instances of the remarkable phenomenon that theorems in logic can sometimes be used to prove "purely mathematical" results that seem to have nothing to do with logic. This introductory section briefly introduces some of the basic ideas that will be needed to understand the later sections.

All the formal languages that we consider will be extensions of a basic logical language that we shall denote by $\mathcal{L}_0$. The statements, or *formulas*, of this language are made up of the following components: *variables*, which are denoted by letters of the alphabet such as x or y, or letters with subscripts such as $v_1, v_2, \dots$; the *parentheses* "(" and ")"; the *equality symbol* "="; the *logical connectives* $\wedge$, $\vee$, $\neg$, $\to$, $\leftrightarrow$, which we read as "and," "or," "not," "implies," and "if and only if"; and the *quantifiers* $\exists$ and $\forall$, which we read as "there exists" and "for all." (If these symbols are unfamiliar to you, then you should read THE LANGUAGE AND GRAMMAR OF MATHEMATICS [I.2] before attempting to read this article.) Here are a couple of formulas of $\mathcal{L}_0$:

(i) $\forall x\ \forall y\ \exists z\ (z \neq x \wedge z \neq y)$;
(ii) $\forall x\ (x = y \vee x = z)$.

The first of these says that if any object exists at all then there are at least three objects, and the second says that y and z are the only objects. There is an important difference between the two formulas: the variables x, y, and z that occur in the first formula are all *bound* variables, which means that they are all attached to quantifiers, whereas in the second formula, only the variable x is bound, while the variables y and z are *free*. This means that the first formula expresses a statement about some mathematical structure, while the second is a statement about not just a structure but also the particular elements y and z.

There are various rules that allow one to build larger formulas out of smaller ones. We will not give them all, but for example if ϕ and ψ are formulas, then $\neg\phi$, $\phi \vee \psi$, $\phi \wedge \psi$, $\phi \to \psi$, and $\phi \leftrightarrow \psi$ are all formulas. In general, if ϕ is built out of smaller formulas $\phi_1, \dots, \phi_n$ using logical connectives (and parentheses), then we call ϕ a *Boolean combination* of $\phi_1, \dots, \phi_n$. Another important way to modify a formula is quantification: if $\phi(x)$ is a formula involving a free variable x, then $\forall x \phi(x)$ and $\exists x \phi(x)$ are both formulas.

The formulas just discussed are "purely logical," which makes them not very useful for describing interesting mathematical structures. Suppose, for example, that we wanted to study real solutions to algebraic and exponential equations over the FIELD [I.3 §2.2] of real numbers. We can think of this as studying the "mathematical structure"

$$\mathbb{R}_{\exp} = (\mathbb{R}, +, \cdot, \exp, <, 0, 1),$$

where the right-hand side is a septuple that consists of the set $\mathbb{R}$ of real numbers, the binary operations of addition and multiplication, the EXPONENTIAL FUNCTION [III.25], the "less than" relation, and the real numbers 0 and 1.

The various components of this structure are of course related to each other in many ways, but we cannot express these relationships unless we are prepared to extend the basic language $\mathcal{L}_0$. For example, if we wanted to write, in a formal way, the statement that the exponential function turns addition into multiplication, then the obvious thing to write down would be

(i) $\forall x \forall y\ \exp(x) \cdot \exp(y) = \exp(x + y)$.

Here we have two quantifiers, two bound variables x and y, and the equals sign, but the rest of the formula involves extraneous elements such as "+", "$\cdot$", and "exp". Thus, to discuss the structure $\mathbb{R}_{\exp}$, we extend the language $\mathcal{L}_0$ to a language $\mathcal{L}_{\exp}$, by adding in the symbols "+", "$\cdot$", "exp", "<", "0", and "1". Of course, these come with various syntactic rules that reflect the fact that "+" is a binary operation, "exp" is a function, and so on. For instance, these rules would allow us to write $\exp(x + y) = z$ but would forbid us to write $\exp(x = y) + z$.

Here are three more $\mathcal{L}_{\exp}$-formulas:

(ii) $\forall x\ (x > 0 \to \exists y\ \exp(y) = x)$;
(iii) $\exists x\ x^2 = -1$;
(iv) $\exists y\ y^2 = x$.

We interpret these formulas as the assertions "for all positive x, there is a y such that $\mathrm{e}^y = x$," "-1 is a square," and "x is a square." The first three formulas above are declarative statements about the structure $\mathbb{R}_{\exp}$. Formulas (i) and (ii) are true in $\mathbb{R}_{\exp}$, while (iii) is false. Formula (iv) is different because x is a free variable: thus, it expresses a property of x. (For instance, it is true if $x = 8$, but false if $x = -7$.) A *sentence* is defined to be a formula with no free variables. If ϕ is an $\mathcal{L}_{\exp}$-sentence, then ϕ is either true or false in $\mathbb{R}_{\exp}$.

If ϕ is a formula with free variables $x_1, \dots, x_n$, and $a_1, \dots, a_n$ are real numbers, then we write $\mathbb{R}_{\exp} \models \phi(a_1, \dots, a_n)$ if the formula ϕ is true for the particular sequence $(a_1, \dots, a_n)$. We think of the formula as defining the set

$$\{(a_1, \dots, a_n) \in \mathbb{R}^n : \mathbb{R}_{\exp} \models \phi(a_1, \dots, a_n)\},$$

that is, the set of all sequences $(a_1, \dots, a_n)$ for which the formula is true when you set x_i to equal a_i for every i. For example, the formula

$$\exists z\ (x = z^2 + 1\ \wedge\ y = z \cdot \exp(\exp(z)))$$

defines the parametrized curve

$$\{(t^2 + 1, t\mathrm{e}^{\mathrm{e}^t}) : t \in \mathbb{R}\}.$$

For another example, one that illustrates an important point, let us consider the structure $(\mathbb{Z}, +, \cdot, 0, 1)$: that is, the integers, with addition, multiplication, 0, and 1. The language used to describe this structure is the *language of rings*, $\mathcal{L}_{\mathrm{rng}} = \mathcal{L}(+, \cdot, 0, 1)$. (The notation here lists the symbols that we add to the basic language $\mathcal{L}_0$.) The language $\mathcal{L}_{\mathrm{rng}}$ has no symbol for the usual ordering on $\mathbb{Z}$, but, surprisingly, this ordering can nevertheless be defined in terms of $\mathcal{L}_{\mathrm{rng}}$. (To appreciate the nonobviousness of this fact, the reader is encouraged to try to work out why it is true before reading on.)

The trick is to use a well-known theorem due to LAGRANGE [VI.22], which asserts that every nonnegative integer is a sum of four squares. It follows that the statement $x \geqslant 0$ can be defined by the formula

$$\exists y_1 \exists y_2 \exists y_3 \exists y_4 \quad x = y_1^2 + y_2^2 + y_3^2 + y_4^2.$$

(Of course, we are also using the fact that a negative integer cannot be written as a sum of four squares. Note too that a similar trick would work even if all one knew was that every nonnegative integer was a sum of a hundred squares.) Once one has a way of expressing the statement that x is nonnegative, it is easy to define the symbol "$<$". The interesting aspect of this is that the reformulation was not obvious—it depended on a genuine mathematical theorem.

It is important to understand that formulas are restricted in several ways, of which two stand out in particular.

- Formulas are finite. We do not allow formulas like
  $$\forall x > 0\ (x < 1 \vee x < 1 + 1 \vee x < 1 + 1 + 1 \vee \cdots),$$
  which would express the fact that $\mathbb{R}$ has the so-called Archimedean property. (If we did, then it would be much easier to define "$<$" above.)
- Quantifiers range over *elements* of the structure, and not subsets. This rules out a "second-order" formula such as
  $$\forall S \subseteq \mathbb{R} \quad (\text{if } S \text{ is bounded above, then } S \text{ has a least upper bound}),$$
  which would express the completeness of $\mathbb{R}$ by quantifying over all subsets S of $\mathbb{R}$. Since we look just at "first-order" formulas, what we are studying is often called *first-order logic*.

Now that we have seen some examples of languages, let us discuss them more generally. A *language* is basically something like $\mathcal{L}_{\exp}$ or $\mathcal{L}_{\mathrm{rng}}$ above: that is, a set of symbols (combined with the basic logical symbols) together with some rules concerning their use. If $\mathcal{L}$ is a language, then an *$\mathcal{L}$-structure* is a mathematical structure in which all the sentences of $\mathcal{L}$ can be interpreted. (This concept will become clearer in a moment, when we give a couple of examples.) An *$\mathcal{L}$-theory T* is just a set of $\mathcal{L}$-sentences, which one can think of as axioms that an $\mathcal{L}$-structure might or might not satisfy. A *model* of T is then an $\mathcal{L}$-structure $\mathcal{M}$ in which all the sentences of T, suitably interpreted, are true. For instance, the structure was a model for the formulas (i) and (ii) of the language $\mathcal{L}_{\exp}$ that we discussed earlier. (Another model for the same two formulas would be one in which we replaced the exponential function by the function 2^x and interpreted "exp" as referring to that function instead.)

The justification for the word "theory" is clearer in another example, the language of GROUPS [I.3 §2.1], $\mathcal{L}_{\mathrm{grp}} = \mathcal{L}(\circ, e)$. Here, $\circ$ is a binary operation symbol and e is a constant. We might look at the theory T_{grp} consisting of the sentences

(i) $\forall x \forall y \forall z\ x \circ (y \circ z) = (x \circ y) \circ z$;
(ii) $\forall x\ x \circ e = e \circ x = x$;
(iii) $\forall x \exists y\ x \circ y = y \circ x = e$;

which are the usual axioms for groups.

In order to interpret this language in some mathematical structure $\mathcal{M}$ we need $\mathcal{M}$ to consist of a set M, a binary operation $f : M^2 \to M$, and an element $a \in M$. We then interpret "$\circ$" as referring to f, "e" as referring to the element a, and quantification as being over the set M. Thus, for example, the interpretation of (iii) is that for every x in M there exists a y in M such that $f(x, y) = a$. Under this interpretation of the symbols of $\mathcal{L}_{\text{grp}}$, the structure $\mathcal{M}$ becomes an $\mathcal{L}_{\text{grp}}$-structure. This $\mathcal{L}_{\text{grp}}$-structure is a model of T_{grp} if in addition the sentences (i), (ii), and (iii) are all true. Since sentences (i)–(iii) are the axioms for groups, a model of T_{grp} is nothing other than a group.

We say that an $\mathcal{L}$-sentence ϕ is a *logical consequence* of a theory T, and write $T \vDash \phi$, if ϕ is true in every model of T. That is, $T \vDash \phi$ if ϕ is true in every structure in which all the sentences of T are true. Thus, the symbol "$\vDash$" has two different meanings, according to whether there is a structure or a theory on the left-hand side. However, these two meanings are closely related in that they are both concerned with truth in models: $\mathcal{M} \vDash \phi$ means that ϕ is true in the model $\mathcal{M}$, and $T \vDash \phi$, as we have just said, means that ϕ is true in every possible model of T. Either way, the symbol "$\vDash$" stands for a "semantic" notion of entailment.

Returning to the example of groups, if ϕ is a sentence in $\mathcal{L}_{\text{grp}}$, then $T_{\text{grp}} \vDash \phi$ if and only if ϕ is true for every group. So, for instance,

$$T_{\text{grp}} \vDash \forall x \forall y \forall z \ (xy \neq xz \vee y = z),$$

because if x, y, and z are elements of any group and $xy = xz$, then we can multiply both sides on the left by the inverse of x to deduce that $y = z$.

We can now describe some of the basic problems in logic.

(i) Given an $\mathcal{L}$-theory T, can we decide if a sentence ϕ is a logical consequence of T, and if so how?

(ii) Given an interesting mathematical structure, like $\mathbb{R}_{\text{exp}}$, or $(\mathbb{N}, +, \cdot, 0, 1)$, or the complex field, and a language $\mathcal{L}$ that describes the structure, can we determine which $\mathcal{L}$-sentences are true of the structure?

(iii) Given a structure described by a language, do the subsets of the structure that can be defined in the language have special properties? Are they in some sense "simple"? For example, earlier we saw how to use $\mathcal{L}_{\text{exp}}$ to define a certain curve in the plane. Now consider a very complicated set such as a CANTOR SET [III.17] or the MANDELBROT SET [IV.14 §2.8]. Is it possible to prove that these sets *cannot* be defined in $\mathcal{L}_{\text{exp}}$ because they are "too complex" in some sense?

2 Completeness and Incompleteness

Let T be an $\mathcal{L}$-theory and let ϕ be an $\mathcal{L}$-sentence. To show that $T \vDash \phi$, we must show that ϕ holds in every model of T. Checking all models of T sounds like a daunting task, but fortunately it is not necessary, since instead we can use a *proof*. One of the first tasks in mathematical logic is to say precisely what this means.

Suppose, then, that $\mathcal{L}$ is some language and that T is a set of sentences in $\mathcal{L}$, i.e., an $\mathcal{L}$-theory. Suppose also that ϕ is a formula of $\mathcal{L}$. Informally speaking, a proof of ϕ assumes the statements of T and ends up establishing ϕ. We express this idea formally as follows. A *proof of ϕ from T* is a finite sequence of $\mathcal{L}$-formulas $\psi_1, \dots, \psi_m$ (which one can think of as the lines of the proof) with the following properties:

(i) each ψ_i is either a logical axiom, or a sentence of T, or a formula that follows from the previous formulas $\psi_1, \dots, \psi_{i-1}$ by means of simple logical rules;

(ii) $\psi_m = \phi$.

We shall not say precisely what a "simple logical rule" is, but three examples are

- from ϕ and ψ it follows that $\phi \wedge \psi$;
- from $\phi \wedge \psi$ it follows that ϕ;
- from $\phi(x)$ it follows that $\exists v \ \phi(v)$.

The other possible rules are similarly elementary.

There are three points about proofs that need to be stressed. The first is that they are finite, which may seem too obvious to mention but is important because it has a number of consequences that are not obvious. The second is that proof systems have to be *sound*: if there is a proof of ϕ from T, then ϕ is true in every model of T. To put this more succinctly, let us introduce the notation $T \vdash \phi$ for the statement that there is a proof of ϕ from T. Then soundness is the assertion that if $T \vdash \phi$ then $T \vDash \phi$. This is why we can prove that ϕ is true in every model of T by finding a proof rather than by looking at all the models. The third point is that it is easy to check whether a sequence of sentences is a proof. More precisely, there is an algorithm that can look at a sequence $\psi_1, \dots, \psi_m$ and decide whether it really is a proof of ϕ from T.

It is not too surprising that if ϕ can be proved from T, then ϕ is true in all models of T. Much more remarkable is that the converse is also true: if ϕ cannot be proved from T, then there must be a model of T in which ϕ is false. This tells us that two very different notions—the finitistic, syntactic notion of "proof" and the semantic notion of "logical consequence," which concerns truth in models—always agree. This result is known as Gödel's completeness theorem. Here is its formal statement.

Theorem. *Let T be an $\mathcal{L}$-theory and let ϕ be an $\mathcal{L}$-sentence. Then $T \vDash \phi$ if and only if $T \vdash \phi$.*

Suppose that T is a simple theory like T_{grp}, where there is an algorithm to decide whether a sentence is in T. (In the case of T_{grp} this algorithm is particularly simple, but some theories might have infinitely many sentences.) We could write a computer program which, given a formula ϕ as its input, would systematically generate all possible proofs σ from T and check to see whether σ was a proof of ϕ. If such a program finds a proof of ϕ, then it halts and tells us that $T \vDash \phi$. We say that $\{\phi : T \vDash \phi\}$ is *recursively enumerable.*

However, one might hope for more. If $T \nvDash \phi$, our program above will go on searching forever, so it will never tell us that there is no proof of ϕ. We say that an $\mathcal{L}$-theory T is *decidable* if there is a computer program which, when given an $\mathcal{L}$-sentence ϕ as input, will always halt and tell us, one way or another, whether $T \vDash \phi$. Such a program would have to be cleverer than the one that just checks all possible proofs σ, and unfortunately such a program does not have to exist: as Gödel [VI.92] proved in his famous incompleteness theorem [V.15], many important theories are undecidable. Here is a first version of his theorem, concerning the *theory of the natural numbers* (or theory of $\mathbb{N}$ for short), which means the set of all sentences in the language $\mathcal{L}_{\mathrm{rng}}$ that are true of the structure $(\mathbb{N}, +, \cdot, 0, 1)$.

Theorem. *The theory of the natural numbers is undecidable.*

At first, this might seem rather strange: after all, if T is the theory of $\mathbb{N}$, then T contains all true sentences about $\mathbb{N}$. So a sentence ϕ is provable from T if and only if it has a one-line proof (the line being ϕ itself). However, this does not make ϕ decidable, because the theory T is very complicated and there is no algorithm for deciding whether ϕ belongs to T.

One approach to proving the incompleteness theorem is to associate a natural number with each computer program in such a way that statements about programs can be recast as statements about natural numbers. The theory of $\mathbb{N}$ then determines whether a program P halts on input x, thus solving what is known as the *halting problem.* Since the halting problem was shown by Turing [VI.94] to be undecidable (a sketch of the proof can be found in the insolubility of the halting problem [V.20]), it follows that the theory of $\mathbb{N}$ is undecidable.

How can we understand the theory of $\mathbb{N}$? One might hope to find a much smaller theory that yielded the same true sentences. That is, we could try to find a simple set of axioms about $\mathbb{N}$ that we know are true and hope that every true sentence follows from these axioms. A good candidate is *first-order Peano arithmetic*, or PA. This is a theory in the language $\mathcal{L}(+, \cdot, 0, 1)$ that involves a few simple axioms about addition and multiplication, such as

$$\forall x \forall y \; x \cdot (y + 1) = x \cdot y + x,$$

together with axioms for induction.

Why do we need more than one axiom of induction? The reason is that the obvious statement that expresses the principle of mathematical induction, namely

$$\forall A \; (0 \in A \wedge \forall x \; x \in A \to x + 1 \in A) \to \forall x \; x \in A,$$

is not a first-order sentence, because the quantifier is applied to all subsets A of $\mathbb{N}$. (It is also not a sentence in $\mathcal{L}_{\mathrm{rng}}$ since it uses the symbol "$\in$", but this is a less fundamental problem.) To get around this difficulty, one has a separate axiom of induction for each formula ϕ. It is the assertion that

$$[\phi(0) \wedge \forall x \; (\phi(x) \to \phi(x + 1))] \to \forall x \; \phi(x).$$

In words, this says that if $\phi(0)$ is true and $\phi(x + 1)$ is true whenever $\phi(x)$ is true, then $\phi(x)$ is true for every x in $\mathbb{N}$.

Most of number theory can be formalized in PA and one might hope that $\mathrm{PA} \vdash \phi$ for every ϕ that is true in $\mathbb{N}$. Sadly, this is not true. Here is a second version of Gödel's incompleteness theorem. Recall that the notation $\mathbb{N} \vDash \psi$ means simply that ψ is true in $\mathbb{N}$.

Theorem. *There is a sentence ψ such that $\mathbb{N} \vDash \psi$ but $\mathrm{PA} \nvdash \psi$.*

Another way to state this result is to say that there is a sentence ψ such that $\mathrm{PA} \nvdash \psi$ and $\mathrm{PA} \nvdash \neg\psi$. To see that this is an equivalent statement, let ψ be any sentence. Then precisely one of ψ and $\neg\psi$ is true. Therefore, if the theorem is false, then PA must prove either ψ or $\neg\psi$. But this means that we can decide which by simply going through all possible proofs in PA until we find a proof of ψ or a proof of $\neg\psi$.

Gödel's original example of a true but unprovable sentence was a self-referential sentence that effectively asserted

"I am not provable from PA."

More precisely, he found a sentence ψ for which he was able to show that ψ is true in $\mathbb{N}$ if and only if ψ is not provable from PA. With more work he showed that there is a sentence that asserts

"PA is consistent"

that is unprovable from PA. The somewhat artificial and metamathematical nature of these sentences might lead one to hope that all "mathematically interesting" sentences about $\mathbb{N}$ are settled by PA. However, more recent work has shown that even this is a forlorn hope, since there are undecidable statements related to RAMSEY'S THEOREM [IV.19 §2.2] in finite combinatorics.

Undecidability also appears in number theory in a very basic way. *Hilbert's tenth problem* asked if there is an algorithm to decide whether a polynomial $p(X_1, \dots, X_n)$ with integer coefficients has an integer zero. Davis, Matijasevic, Putnam, and Robinson showed that the answer is no.

Theorem. *For any recursively enumerable $S \subseteq \mathbb{N}$ there is $n > 0$ and $p(X, Y_1, \dots, Y_n) \in \mathbb{Z}[X, Y_1, \dots, Y_n]$ such that $m \in S$ if and only if $p(m, Y_1, \dots, Y_n)$ has an integer zero.*

Since the halting problem provides an undecidable recursively enumerable set, the answer to Hilbert's tenth problem is no. An important open question is whether there is an algorithm to decide if a polynomial with *rational* coefficients has a *rational* zero. Hilbert's tenth problem is also discussed in THE INSOLUBILITY OF THE HALTING PROBLEM [V.20], and other interesting examples of undecidability can be found in GEOMETRIC AND COMBINATORIAL GROUP THEORY [IV.10].

3 Compactness

A theory T is called *satisfiable* if there are structures that satisfy all of the sentences in T (that is, if T has a model), and we call T *consistent* if we cannot derive a contradiction from T. Since our proof system is sound, any satisfiable theory is consistent. On the other hand if T is not satisfiable, then every sentence ϕ is a logical consequence of T, for the trivial reason that there are no models of T in which ϕ is required to be true. But the completeness theorem then tells us that $T \vdash \phi$ for every ϕ. Choosing ϕ to be some contradictory statement, of the form $\psi \wedge \neg\psi$, for instance, we see that T is inconsistent. This way of reformulating the completeness theorem has the following simple consequence, called the *compactness theorem*, which turns out to be surprisingly important, as we shall see.

Theorem. *If every finite subset of T is satisfiable, then T is satisfiable.*

The reason this is true is that if T is not satisfiable then it is inconsistent (as we have just seen), which means that a contradiction can be proved from T. Since this proof, like all proofs, must be finite, it involves only finitely many sentences from T. Therefore, T has a finite subset that implies a contradiction, which contradicts our assumption that all finite subsets of T are satisfiable.

Although the compactness theorem is an easy consequence of the completeness theorem, it has many immediate intriguing consequences and lies at the heart of many constructions in model theory. Here are two simple applications that show that theories have many models that you might not expect. If $\mathcal{M}$ is some $\mathcal{L}$-structure, let us write $\mathrm{Th}(\mathcal{M})$ for *the theory of* $\mathcal{M}$: that is, for the set of all $\mathcal{L}$-sentences that are true in $\mathcal{M}$. We also extend our earlier notation $\mathcal{M} \models \phi$ from single formulas to collections of formulas, so if $\mathcal{M}$ is an $\mathcal{L}$-structure and T is an $\mathcal{L}$-theory, then $\mathcal{M} \models T$ means that every sentence of T is true in M, or in other words that $\mathcal{M}$ is a model of T.

Corollary. *There exists an $\mathcal{L}_{\mathrm{exp}}$-structure $\mathcal{M}$ containing an infinite element a (which means that $a > 1$, $a > 1+1$, $a > 1 + 1 + 1$, etc.), such that $\mathcal{M} \models \mathrm{Th}(\mathbb{R}_{\mathrm{exp}})$.*

That is, there is a structure $\mathcal{M}$ in which all the true first-order statements about the structure $\mathbb{R}_{\mathrm{exp}}$ are still true, but $\mathcal{M}$ is different from $\mathbb{R}_{\mathrm{exp}}$ because it contains an infinite element. To prove this, we add one more constant symbol c to our language and consider the theory T that consists of all the statements of $\mathrm{Th}(\mathbb{R}_{\mathrm{exp}})$ (that is, all true statements about $\mathbb{R}_{\mathrm{exp}}$), together with the infinite sequence of statements $c > 1$, $c > 1 + 1$, $c > 1 + 1 + 1$, and so on. If Δ is any finite subset of T, then we can make $\mathbb{R}$ a model of Δ simply by interpreting

c as a sufficiently large real number—large enough to satisfy all the statements of the form $c > 1 + 1 + \cdots + 1$ that belong to Δ. Since we can model every finite subset Δ of T, the compactness theorem tells us that we can model T itself. If $\mathcal{M} \vDash T$, then the element named by c must be infinite.

The element $1/a$ will be an *infinitesimal* element of $\mathcal{M}$ (which means that it satisfies statements that effectively say that it is smaller than $1/n$ for every positive integer n). This observation is the first step toward a rigorous development of calculus with infinitesimals.

For another example, let $\mathcal{L}_{\text{rng}} = \mathcal{L}(+, \cdot, 0, 1)$ be the language of rings. Let T be the set of $\mathcal{L}$-sentences that are true in every finite field. We call T the *theory of finite fields*. Recall that a field is said to have *characteristic p* if p is the smallest positive integer (which has to be prime) such that $1 + 1 + \cdots + 1 = 0$ in the field, where the number of 1s in the sum is p. If there is no such p, then the field is said to have *characteristic zero*. Thus, the fields $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$ all have characteristic zero.

Corollary. *There is a field F with characteristic zero such that* $F \vDash T$.

This result tells us that there is no possible set of axioms that characterizes the finite fields: given any set of statements that are true in all finite fields, there is an infinite field in which they are also all true. To prove it, we look at the theory T' that consists of T together with the statements $1 + 1 \neq 0$, $1 + 1 + 1 \neq 0$, and so on. Any finite set of statements in T' will be true of a finite field of sufficiently large characteristic, and thus satisfiable. By the compactness theorem T' is satisfiable, but a model of T clearly has to have characteristic zero.

The compactness theorem can sometimes be used to show the existence of interesting algebraic bounds. The next result allows us to deduce from HILBERT'S NULLSTELLENSATZ [V.17] a stronger "quantitative version." It is our first example of a statement that does not appear to be logical in nature but which can be proved using logic. Recall that a field is *algebraically closed* if every polynomial with coefficients in the field has a root in the field. (THE FUNDAMENTAL THEOREM OF ALGEBRA [V.13] is the assertion that $\mathbb{C}$ is an algebraically closed field.)

Proposition. *For any three positive integers n, m, d there is a positive integer l such that if K is an algebraically closed field and $f_1, \dots, f_m$ are polynomials in n variables with coefficients in K, degree at most d and no common zero, then there are polynomials $g_1, \dots, g_m$ of degree at most l such that* $\sum g_i f_i = 1$.

Hilbert's Nullstellensatz itself is the same statement but without the extra information about the degrees of the polynomials g_i.

To see how the proposition is proved, we will restrict our attention to the case $n = d = 2$. This is just for notational simplicity: the proof is almost identical in larger cases. For each i between 1 and m let

$$F_i = a_i X^2 + b_i Y^2 + c_i XY + d_i X + e_i Y + f_i.$$

For each k write down a formula ϕ_k that asserts that there are no polynomials $G_1, \dots, G_m$ with degree at most k such that $1 = \sum F_i G_i$. Let T be the theory of algebraically closed fields with the formulas $\phi_1, \phi_2, \dots$ and the assertion that the polynomials $F_1, \dots, F_m$ have no common zero. If there is no positive integer l satisfying the conclusion of the proposition, then every finite subset of T is satisfiable. Hence, by the compactness theorem, T is satisfiable. If $K \vDash T$, then $F_1, \dots, F_m$ are polynomials over an algebraically closed field with no common zero, but it is impossible to find polynomials $G_1, \dots, G_m$ such that $\sum G_i F_i = 1$. This contradicts Hilbert's Nullstellensatz.

Notice that in the above argument we did not say anything about the dependence of l on n, m, and d. This is because the proof does not actually find a bound: it merely shows that some sort of bound must exist. However, good explicit bounds were recently discovered—see ALGEBRAIC GEOMETRY [IV.4] for more details.

4 The Complex Field

A surprising counterpoint to Gödel's incompleteness theorem is a result of TARSKI [VI.87], which states that the theories of the fields of real and complex numbers *are* decidable. The key to these results is a method known as *quantifier elimination*. If we have a formula without quantifiers that concerns the natural numbers, then it is easy to decide whether it is true or false. The negative solution to Hilbert's tenth problem shows that as soon as we start adding existential quantifiers (as we do if, for example, we assert that a polynomial has a zero), then we leave the realm of decidability.

Thus, if we want to show that a formula is decidable, it will be very useful if we can find an equivalent formula that does not have quantifiers. And in some settings, this turns out to be possible. For example, let $\phi(a, b, c)$ be the formula

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

The usual rule for solving quadratics tells us that, as long as $a \neq 0$, this is true in $\mathbb{R}$ if and only if $b^2 \geqslant 4ac$. Therefore, $\mathbb{R} \vDash \phi(a, b, c)$ if and only if

$$[(a \neq 0 \wedge b^2 - 4ac \geqslant 0) \vee (a = 0 \wedge (b \neq 0 \vee c = 0))].$$

As for the complex numbers, it is easy to see that $\mathbb{C} \vDash \phi(a, b, c)$ if and only if

$$a \neq 0 \vee b \neq 0 \vee c = 0.$$

In either case, ϕ is equivalent to a formula with no quantifiers.

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

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

The formula $\phi(a, b, c, d)$ is the obvious way of asserting that the matrix $(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix})$ is invertible. However, by the DETERMINANT [III.15] test, we know that, for any field F, $F \vDash \phi(a, b, c, d)$ if and only if $ad - bc \neq 0$. Thus the existence of an inverse can be expressed by the quantifier-free formula $ad - bc \neq 0$.

Tarski proved that we can *always* eliminate quantifiers in algebraically closed fields.

Theorem. *For any $\mathcal{L}_{\text{rng}}$-formula ϕ there is a quantifier-free formula ψ such that ϕ is equivalent to ψ in every algebraically closed field.*

Furthermore, Tarski gave an explicit algorithm for eliminating the quantifiers.

The equivalent quantifier-free formulas above were both finite Boolean combinations of formulas of the form $p(v_1, \dots, v_n) = q(v_1, \dots, v_n)$, where p and q are polynomials in n variables with integer coefficients. It is not hard to see that this is true of any quantifier-free $\mathcal{L}_{\text{rng}}$-formula. It follows that a quantifier-free $\mathcal{L}_{\text{rng}}$-*sentence* is particularly simple: if no free variables are allowed and no quantifiers are allowed, then there cannot be any variables! Therefore, the polynomials p and q have to be constant, which means that a quantifier-free $\mathcal{L}_{\text{rng}}$-*sentence* is a finite Boolean combination of formulas of the form $k = l$ (where this should be regarded as an abbreviation for $1 + 1 + \cdots + 1 = 1 + 1 + \cdots + 1$, with k 1s on the left-hand side and l 1s on the right-hand side).

This leads to the decidability result. If we want to know whether $\mathbb{C} \vDash \phi$, then we use Tarski's algorithm to convert ϕ into an equivalent quantifier-free sentence. But the very simple form of such sentences makes their truth or falsity easy to decide.

In the remainder of this section, we shall discuss a number of other consequences of Tarski's theorem. The first is that sentences in the language $\mathcal{L}_{\text{rng}}$ cannot distinguish between different algebraically closed fields of the same characteristic. That is, if ϕ is any $\mathcal{L}_{\text{rng}}$-sentence that is true for some algebraically closed field of characteristic p (where p is allowed to be zero), then it is true in every algebraically closed field of characteristic p.

To see why this is true, let K and F be two algebraically closed fields of characteristic p, and suppose that $K \vDash \phi$ (or in other words that ϕ is true of K). Let k be the field $\mathbb{Q}$ if the characteristic is zero and the field with p elements otherwise. Tarski's theorem tells us that there is a quantifier-free sentence ψ that is equivalent to ϕ in all algebraically closed fields of characteristic p. However, the extremely simple nature of the quantifier-free sentences of $\mathcal{L}_{\text{rng}}$ means that their truth or falsity in any given field depends only on the elements 0, 1, 1 + 1, and so on. Therefore,

$$K \vDash \psi \ \Leftrightarrow \ k \vDash \psi \ \Leftrightarrow \ F \vDash \psi.$$

Since $K \vDash \phi$ and ϕ and ψ are equivalent in all algebraically closed fields of characteristic p, it follows that $F \vDash \phi$ as well.

A consequence of this theorem is that an $\mathcal{L}_{\text{rng}}$-sentence ϕ is true of the complex numbers if and only if it is true of the algebraic numbers $\mathbb{Q}^{\text{alg}}$. (Recall that these are all roots of polynomials with integer coefficients. As one would expect, the algebraic numbers form an algebraically closed field, though this is not a wholly obvious fact.) Thus, rather surprisingly, if we wish to prove something about $\mathbb{Q}^{\text{alg}}$, we have the option of working in $\mathbb{C}$ and using the methods of complex analysis; similarly, if we want to prove something about $\mathbb{C}$ we can, if it makes things easier, work in $\mathbb{Q}^{\text{alg}}$ and use number-theoretic methods.

Combining these ideas with the completeness theorem gives another useful tool. If ϕ is any $\mathcal{L}_{\text{rng}}$-sentence, then the following are equivalent:

(i) ϕ is true in every algebraically closed field of characteristic zero;
(ii) for some $m > 0$, ϕ is true in every algebraically closed field of characteristic $p > m$;
(iii) there are arbitrarily large p such that ϕ is true in some algebraically closed field of characteristic p.

Let us see why this is so. Suppose first that ϕ is true in every algebraically closed field of characteristic 0. The completeness theorem then implies that there is a *proof* of ϕ from the axioms for algebraically closed fields combined with the sentences $1 \neq 0$, $1 + 1 \neq 0$, $1+1+1 \neq 0$, and so on. Since proofs are finite sequences of formulas, there must be some m such that the proof used only the first m of these sentences (not necessarily all of them). If p is some prime bigger than m, then this proof shows that ϕ holds in algebraically closed fields of characteristic p, since all the sentences we used are true in such fields.

We have just shown that (i) implies (ii). It is obvious that (ii) implies (iii). To see that (iii) implies (i), let us suppose that (i) fails, so that there is an algebraically closed field of characteristic zero in which $\neg\phi$ is true. Then, by the principle we proved earlier, $\neg\phi$ is true in *every* algebraically closed field of characteristic zero. Thus, since (i) implies (ii), there is an m such that $\neg\phi$ is true in every algebraically closed field of characteristic $p > m$. Therefore (iii) fails.

An interesting application of this theorem was found by Ax. It is another example of a statement that has nothing to do with logic, but which can be proved using logical tools. It is perhaps more striking than the previous example because in this case one does not even feel with hindsight that the statement did after all have some logical content.

Theorem. *If a polynomial map from $\mathbb{C}^n$ to $\mathbb{C}^n$ is an injection, then it must also be a surjection.*

The basic thought behind the proof of this result is very simple indeed: what is remarkable is that it is of any help. It is the observation that if k is a finite field, then every injective polynomial map from k^n to k^n is a surjection. This is true because every injection from a finite set to itself is automatically a surjection.

How do we exploit this observation? Well, the previous results tell us that, in several situations, statements are true for one field if and only if they are true for another. We shall use these results to transfer our problem from $\mathbb{C}$, where it is hard, to a finite field k, where it is trivial. The first step is a routine exercise: one shows that for each positive integer d there is a sentence ϕ_d in $\mathcal{L}_{\text{rng}}$ that expresses the fact that every injective polynomial map from F^n to F^n, with the n polynomials all of degree at most d, is surjective. We would like to prove that all the sentences ϕ_d are true when $F = \mathbb{C}$.

The equivalences in the previous theorem imply that it is enough to prove that the sentences ϕ_d are true when F is the field $\mathbb{F}_p^{\text{alg}}$, the algebraic closure of the p-element field. (It can be shown that any field F is contained in an algebraically closed field. Roughly speaking, the *algebraic closure* of F is the smallest algebraically closed field that contains F.) Suppose, then, that some ϕ_d fails for $\mathbb{F}_p^{\text{alg}}$. Then there must be an injective polynomial map f from $(\mathbb{F}_p^{\text{alg}})^n$ to $(\mathbb{F}_p^{\text{alg}})^n$ that is not surjective. Since every finite subset of $\mathbb{F}_p^{\text{alg}}$ is contained in a finite subfield, there is a finite subfield k such that all the n polynomials used to define f have coefficients in k, from which it follows that f maps k^n to k^n. Moreover, by enlarging k if necessary, we can ensure that there is an element of k^n that is not in the image of f. But now we have succeeded in transferring ourselves to a finite field: this function $f : k^n \to k^n$ is an injection between finite sets that is not a surjection, which is a contradiction.

Quantifier elimination has other useful applications. Let F be a field, let K be a subfield of F, let $\psi(v_1, \dots, v_n)$ be a quantifier-free formula, and let $a_1, \dots, a_n$ be elements of K. Since, as we have already mentioned, quantifier-free formulas are just Boolean combinations of equalities between polynomials, the statement $\psi(a_1, \dots, a_n)$ involves just the elements of K, and is therefore true in K if and only if it is true in F. By quantifier elimination, if K and F are algebraically closed, then the same is true for *all* formulas ψ, and not just those that are quantifier free. From this observation we can prove the "weak version" of Hilbert's Nullstellensatz. (For the proof, we shall need to assume a certain degree of familiarity with the basics of RING THEORY [III.81]. We shall also write $K[\boldsymbol{X}]$ for the polynomial ring $K[X_1, \dots, X_n]$ and $\bar{v}$ for the n-tuple $(v_1, \dots, v_n)$.)

Proposition. *Suppose that K is an algebraically closed field, P is a prime ideal in $K[\boldsymbol{X}]$, and g is a polynomial in $K[\boldsymbol{X}]$ that does not belong to P. Then there is some $a = (a_1, \dots, a_n)$ in K^n such that $f(a) = 0$ for every f that belongs to P, and such that $g(a) \neq 0$.*

Proof. Let F be the algebraic closure of the fraction field of the integral domain $K[\boldsymbol{X}]/P$. We can view F as an extension field of K with a natural homomorphism $\eta : K[\boldsymbol{X}] \to F$. Let $b_i = \eta(X_i)$ and let $b \in F^n$ be the element $(b_1, \dots, b_n)$. Then $f(b) = 0$ for all $f \in P$ and $g(b) \neq 0$. We would like to find such an element in K. Since ideals in polynomial rings are finitely generated, we can find polynomials $f_1, \dots, f_m$ that generate P. The sentence

$$\exists v_1 \cdots \exists v_n (f_1(\bar{v}) = \cdots = f_m(\bar{v}) = 0 \ \wedge g(\bar{v}) \neq 0)$$

is true in F. Thus it is also true in K and we can find $a \in K^n$ such that each $f \in P$ vanishes at a but $g(a) \neq 0$. □

Notice that the above proof has the same basic structure as the result about polynomial maps on $\mathbb{C}^n$. The idea was to come up with a different field, in this case F, where the result was easy to prove, and use logical ideas to deduce the result for the field we were originally interested in, in this case K.

5 The Reals

Quantifier elimination in the language of rings does not work in the field of real numbers. For instance, the formula

$$\exists y\; x = y \cdot y,$$

which asserts "x is a square," is not equivalent to a quantifier-free formula in the language of rings. Of course, x is a square if and only if $x \geqslant 0$. So we *could* eliminate this quantifier if we were prepared to add a symbol for the ordering to our language. An amazing result of Tarski shows that this is the only obstruction to quantifier elimination.

Let $\mathcal{L}_{\text{or}}$ be the language of ordered rings, which is the language of rings with the addition of the symbol "<" for an ordering. Which $\mathcal{L}_{\text{or}}$-sentences are true in the real field? Some of the properties of $\mathbb{R}$ that we can formalize in $\mathcal{L}_{\text{or}}$ include:

(i) the axioms for ordered fields, such as the sentence
$$\forall x \forall y\; (x > 0 \wedge y > 0) \to x \cdot y > 0;$$

(ii) the intermediate-value property for polynomials, which states that if $p(x)$ is a polynomial and there exist a and b such that $a < b$ and $p(a) < 0 < p(b)$, then there exists a real number c such that $a < c < b$ and $p(c) = 0$.

The intermediate-value property is expressed not by just one sentence, but by the infinite sequence of sentences

$$\forall d_0 \cdots \forall d_n \forall a \forall b$$
$$\Big(\sum d_i a^i < 0 < \sum d_i b^i \to \exists c \sum d_i c^i = 0\Big),$$

one for each positive integer n.

An ordered field that satisfies the intermediate-value property is called a *real closed* field. It turns out that an equivalent way of axiomatizing real closed fields is as ordered fields for which every positive element is a square and every polynomial of odd degree has a zero. Tarski's theorem is the following statement.

Theorem. *For any $\mathcal{L}_{\text{or}}$-formula ϕ there is a quantifier-free $\mathcal{L}_{\text{or}}$-formula ψ such that ϕ and ψ are equivalent in every real closed field.*

What are the quantifier-free formulas of $\mathcal{L}_{\text{or}}$? It turns out (and is not hard to show) that they are finite Boolean combinations of formulas of the form $p(v_1, \dots, v_n) = q(v_1, \dots, v_n)$ and formulas of the form $p(v_1, \dots, v_n) < q(v_1, \dots, v_n)$, where, as in the case of $\mathcal{L}_{\text{rng}}$, p and q are polynomials in n and m variables, respectively, with integer coefficients. As for quantifier-free *sentences*, they are Boolean combinations of sentences of the form $k = l$ and sentences of the form $k < l$.

One consequence of quantifier elimination is the following result, which tells us that every $\mathcal{L}_{\text{or}}$ statement that is true in $\mathbb{R}$ can be proved from the real-closed-field axioms. One says that these axioms *completely axiomatize* the theory of the real field.

Corollary. *Let K be a real closed field and let ϕ be an $\mathcal{L}_{\text{or}}$-sentence. Then $K \vDash \phi$ if and only if $\mathbb{R} \vDash \phi$.*

To prove this, first use Tarski's theorem to find a quantifier-free sentence ψ such that ϕ and ψ are equivalent in any real closed field. Every ordered field has characteristic zero and contains the rational numbers as an ordered subfield. Therefore $\mathbb{Q}$ is a subfield of both K and $\mathbb{R}$. But the very simple nature of quantifier-free sentences in $\mathcal{L}_{\text{or}}$ means that

$$K \vDash \psi \;\Leftrightarrow\; \mathbb{Q} \vDash \psi \;\Leftrightarrow\; \mathbb{R} \vDash \psi.$$

Since ϕ and ψ are equivalent in all real closed fields, it follows that $K \vDash \phi$ if and only if $\mathbb{R} \vDash \phi$.

By the completeness theorem, ϕ is true in every real closed field if and only if we can prove ϕ from the axioms for real closed fields, and ϕ is false in every real closed field if and only if we can prove $\neg\phi$ from the axioms for real closed fields. It follows that the $\mathcal{L}_{\text{or}}$-theory of the real field is decidable. Indeed, if ϕ is true in $\mathbb{R}$, then by the corollary above, it is true in every real closed field, so it has a proof. If ϕ is false in $\mathbb{R}$, then $\neg\phi$ is true in $\mathbb{R}$, so for the same reason $\neg\phi$ has a proof. Therefore, to decide whether ϕ is true, one can search through all possible proofs from the axioms of real closed fields until one proves either ϕ or $\neg\phi$.

Let $\mathcal{M}$ be a mathematical structure consisting of a set M and various other parts such as functions and binary operations. A subset X of M is called *definable*, with respect to some language $\mathcal{L}$ that describes $\mathcal{M}$, if there is an $\mathcal{L}$-formula ϕ with a free variable x such that $X = \{x \in M : \phi(x)\}$. Quantifier elimination gives us a good geometric understanding of the definable sets. If K is an ordered field, we say that $X \subseteq K^n$ is *semialgebraic* if it is a finite Boolean combination of sets of the form

$$\{x \in K^n : p(x) = 0\} \quad \text{and} \quad \{x \in K^n : q(x) > 0\},$$

where $p, q \in K[X_1, \dots, X_n]$. By quantifier elimination, the definable sets in a real closed field are easily shown to be exactly the semialgebraic sets.

A simple application of this fact is that if A is a semialgebraic subset of $\mathbb{R}^n$, then the closure of A is also semialgebraic. Indeed, the closure of A is, by definition, the set

$$\left\{x \in \mathbb{R}^n : \forall \epsilon > 0\ \exists y \in A\ \sum_{i=1}^{n} (x_i - y_i)^2 < \epsilon\right\}.$$

This is a definable set, and hence a semialgebraic set.

Semialgebraic subsets of the real line are particularly simple. For any real polynomial f in one variable, the set $\{x \in \mathbb{R} : f(x) > 0\}$ is a finite union of open intervals. Therefore, any semialgebraic subset of $\mathbb{R}$ is a finite union of points and intervals. This simple fact is the starting point of the modern model-theoretic approach to $\mathbb{R}$. Let $\mathcal{L}^*$ be a language extending $\mathcal{L}_{\text{or}}$ and let $\mathbb{R}^*$ denote the reals considered as an $\mathcal{L}^*$-structure. For example, below we will be interested in the case where $\mathcal{L}^* = \mathcal{L}_{\text{exp}}$ and $\mathbb{R}^* = \mathbb{R}_{\text{exp}}$. We say that $\mathbb{R}^*$ is *o-minimal* if every subset of $\mathbb{R}$ definable using $\mathcal{L}^*$-formulas is a finite union of points and intervals. The "o" in "o-minimal" stands for "ordered." $\mathbb{R}^*$ is o-minimal if every definable subset of $\mathbb{R}$ can be defined using only the ordering.

Pillay and Steinhorn introduced o-minimality, generalizing an earlier idea of van den Dries. It turned out to be a key definition, because although o-minimality is defined in terms of the one-dimensional set $\mathbb{R}$, it has remarkably strong consequences for definable subsets of $\mathbb{R}^n$ when $n > 1$.

To explain this, we inductively define a collection of basic sets called *cells* as follows.

- A subset X of $\mathbb{R}$ is a cell if and only if it is either a point or an interval.
- If X is a cell in $\mathbb{R}^n$ and f is a continuous definable function from X to $\mathbb{R}$, then the graph of f (which is a subset of $\mathbb{R}^{n+1}$) is a cell.
- If X is a cell in $\mathbb{R}^n$ and f and g are continuous definable functions from X to $\mathbb{R}$ such that $f(x) > g(x)$ for every $x \in X$, then $\{(x, y) : x \in X \text{ and } f(x) > y > g(x)\}$ is a cell, as are $\{(x, y) : x \in X \text{ and } f(x) > y\}$ and $\{(x, y) : x \in X \text{ and } y > f(x)\}$.

Cells are topologically simple definable sets that play the role of open intervals in $\mathbb{R}$. It is not hard to see that any cell is homeomorphic to $(0, 1)^n$ for some n. Remarkably, all definable sets can be decomposed into cells. The following theorem is a precise version of this statement.

Theorem.

(i) *If $\mathbb{R}^*$ is an o-minimal structure, then every definable set X can be partitioned into finitely many disjoint cells.*

(ii) *If $f : X \to \mathbb{R}$ is a definable function, then there is a partition of X into finitely many cells such that f is continuous on each cell.*

This is just the beginning. In any o-minimal structure, definable sets have many of the good topological and geometric properties of the semialgebraic sets. For example:

- Any definable set has finitely many connected components.
- Definable bounded sets can be definably triangulated.
- Suppose that X is a definable subset of $\mathbb{R}^{n+m}$. For each $a \in \mathbb{R}^m$, let X_a be the "cross-section" $\{x \in \mathbb{R}^n : (x, a) \in X\}$. Then there are only finitely many different homeomorphism types for the sets X_a.

As these results were known for semialgebraic sets, the real interest is in finding new o-minimal structures. The most interesting example is $\mathbb{R}_{\text{exp}}$. It is known that $\mathbb{R}_{\text{exp}}$ does not have quantifier elimination in the language $\mathcal{L}_{\text{exp}}$. Wilkie showed that the next best thing is true. We say that $\mathbb{R}^n$ is an *exponential variety* if it is the zero set of a finite system of exponential terms. For example, the set $\{(x, y, z) : x = \exp(y)^2 - z^3 \wedge \exp(\exp(z)) = y - x\}$ is an exponential variety.

Theorem. *Every $\mathcal{L}_{\text{exp}}$-definable subset of $\mathbb{R}^n$ is of the form*

$$\{x \in \mathbb{R}^n : \exists y \in \mathbb{R}^m\ (x, y) \in V\}$$

for some exponential variety $V \subseteq \mathbb{R}^{n+m}$.

In other words, the definable sets, though not exponential varieties themselves, are projections of exponential varieties, which makes them tractable. Indeed, a theorem from real analytic geometry, due to Khovanskii, states that every exponential variety has a finite number of connected components. Since this property is preserved by projections, it follows that every definable set has a finite number of connected components, and also that every definable subset of the real line is a finite union of points and intervals. Thus $\mathbb{R}_{\text{exp}}$ is o-minimal and all of the results above about definable sets in o-minimal structures apply.

Tarski asked if the theory of $\mathbb{R}_{\text{exp}}$ is decidable. This question remains open, but the answer is known to follow from the following conjecture of Schanuel in transcendental number theory.

Conjecture. *Suppose that* $\lambda_1, \dots, \lambda_n$ *are complex numbers that are linearly independent over* $\mathbb{Q}$*. Then the field* $\mathbb{Q}(\lambda_1, \dots, \lambda_n, \mathrm{e}^{\lambda_1}, \dots, \mathrm{e}^{\lambda_n})$ *has transcendence degree at least* n*.*

Macintyre and Wilkie have shown that if Schanuel's conjecture is true, then the theory of $\mathbb{R}_{\text{exp}}$ is decidable.

6 The Random Graph

Model-theoretic methods give interesting information about random GRAPHS [III.34]. Suppose we construct a graph as follows. The vertex set is the set $\mathbb{N}$ of all natural numbers $\mathbb{N}$. To decide whether we will have an edge between x and y (with $x \neq y$) we flip a coin, putting an edge there if and only if we get heads. Although these constructions are random, we will show below that, with probability 1, any two such graphs are isomorphic.

The proof depends on the following extension property. Let A and B be disjoint finite subsets of $\mathbb{N}$, and suppose that they have sizes n and m, respectively. We would like to find a vertex $x \in \mathbb{N}$ that is joined to every element of A and to no element of B. Now for any particular x, the probability that it does *not* have the desired property is $p = 1 - 2^{-(n+m)}$. Therefore, if we look at N different vertices, the probability that none of them has the desired property is p^N. Since this converges to zero with N, the probability that at least one $x \in \mathbb{N}$ has the property is 1. Moreover, since there are only countably many disjoint pairs (A, B) of finite sets, with probability 1 it is the case that for *every* such pair (A, B) one can find a vertex x that is joined to every vertex in A and to no vertex in B.

We can formalize this observation in a model-theoretic way. Let $\mathcal{L}_{\text{g}} = \mathcal{L}(\sim)$, where "$\sim$" is a binary relation symbol (which we read as "is joined to"). We let T be the $\mathcal{L}_{\text{g}}$-theory:

(i) $\forall x \forall y \; x \sim y \to y \sim x$;
(ii) $\forall x \; \neg(x \sim x)$;
(iii) $\Phi_{n,m}$ for $n, m \geqslant 0$.

Here $\Phi_{n,m}$ is the sentence

$$\forall x_1 \cdots \forall x_n \forall y_1 \cdots \forall y_m$$
$$\bigwedge_{i=1}^{n} \bigwedge_{j=1}^{m} x_i \neq y_j \to \exists z \left(\left(\bigwedge_{i=1}^{n} x_i \sim z \right) \wedge \left(\bigwedge_{i=1}^{m} \neg(y_i \sim z) \right) \right).$$

The first two sentences tell us that the relation "$\sim$" defines a graph, and for each pair (n, m) the sentence $\Phi_{n,m}$ tells us that the extension property holds for all pairs of disjoint sets A and B with A of size n and B of size m. Thus, a model of T is a graph for which the extension property holds for any pair of disjoint finite sets of vertices.

The argument above shows that with probability 1 the random graphs we constructed are models of T. Now let us see why they are isomorphic (again with probability 1). This will be an immediate consequence of the following theorem.

Theorem. *If* G_1 *and* G_2 *are any two countable models of* T*, then* G_1 *is isomorphic to* G_2*.*

Recall that an *isomorphism* between G_1 and G_2 means a bijection f from the vertex set of G_1 to the vertex set of G_2 such that x is joined to y in G_1 if and only if $f(x)$ is joined to $f(y)$ in G_2. The proof, which we shall now sketch, is a "back-and-forth" argument that gradually builds up an isomorphism between G_1 and G_2. First, let $a_0, a_1, \dots$ be an enumeration of the vertices of G_1 and let $b_0, b_1, \dots$ be an enumeration of the vertices of G_2. Let us set $f(a_0)$ to be b_0. Next, we choose an image for a_1: if a_1 is joined to a_0 then we need to find some vertex that is joined to b_0 and if a_1 is not joined to a_0 then we need to find a vertex that is not joined to b_0. Either way, we can do it because G is a model of T, so it satisfies the extension property. (The particular cases we use here are $\Phi_{1,0}$ and $\Phi_{0,1}$.)

It is tempting to continue finding images for a_2, a_3, and so on, in each case using the extension property to make sure that the images are joined to each other if and only if the original vertices are. The trouble with this is that we may not end up with a bijection, since for any particular b_j there is no guarantee that we will ever choose it as the image of some a_j. However, we can remedy this by alternately choosing an image for the first a_i that does not yet have an image, and a preimage for the first b_j that does not yet have a preimage. In this way we build the desired isomorphism.

It was not essential to use model theory to prove the above result. However, it has the following very nice model-theoretic consequence.

Corollary. *For any* $\mathcal{L}_{\text{g}}$*-sentence* ϕ *either* ϕ *is true in every model of* T *or* $\neg\phi$ *is true in every model of* T*. Moreover, there is an algorithm that will tell us which of* ϕ *or* $\neg\phi$ *is true in every model of* T*.*

To prove this, one first applies a slight strengthening of the compactness theorem, which allows one to conclude that if the result is false then there are *countable* models G_1 and G_2 of T such that ϕ is true in G_1 and $\neg\phi$ is true in G_2. But this shows that G_1 and G_2 are not isomorphic, and therefore directly contradicts the previous theorem.

To decide which of ϕ or $\neg\phi$ is true in every model of T, one searches through all possible proofs from the sentences of T. By the completeness theorem, one or other of the statements has a proof, so we will eventually find either a proof of ϕ or a proof of $\neg\phi$. At that point we will know which of ϕ and $\neg\phi$ is true in every model of T.

The theory T also gives us information about random finite graphs. Let $\mathcal{G}_N$ be the set of all graphs with vertices $\{1, 2, \dots, N\}$. We consider the probability measure on $\mathcal{G}_N$ in which we make all graphs equally likely. This is the same as constructing a random graph on N vertices, where for each i and j we toss an unbiased coin in order to decide whether i is joined to j. For any $\mathcal{L}_\mathrm{g}$-sentence ϕ, let us write $p_N(\phi)$ for the probability that a random graph on N vertices satisfies ϕ.

An easy variant of the argument for infinite graphs shows that for each extension axiom $\Phi_{n,m}$, the probability $p_N(\Phi_{n,m})$ tends to 1. Therefore, for any fixed M, if N is sufficiently large, then with very high probability a random graph on N vertices satisfies all the axioms $\Phi_{n,m}$ with $n, m \leqslant M$.

This observation allows us to use the theory T to get a good understanding of the asymptotic properties of random graphs. The following result is called a *zero-one law.*

Theorem. *Given any $\mathcal{L}_\mathrm{g}$-sentence ϕ, the probability $p_N(\phi)$ either tends to* 0 *or tends to* 1 *as $N \to \infty$. Moreover, T axiomatizes the set of statements ϕ such that the limit is* 1*, called the* almost sure theory of graphs*, which is a decidable theory.*

This follows from our previous results. We saw earlier that either ϕ is true in every model of T or $\neg\phi$ is true in every model of T. In the first case, by the completeness theorem there must be a proof of ϕ from T. Since proofs are finite, this proof can use only finitely many of the statements $\Phi_{n,m}$. Therefore, there exists some M such that if $G \models \Phi_{M,M}$, then $G \models \phi$. But if G is a random graph on N vertices, then the probability that $G \models \Phi_{M,M}$ tends to 1, and therefore the probability $p_N(\phi)$ that $G \models \phi$ tends to 1 as well. The same argument holds if $\neg\phi$ is true in every model of T and shows that $p_N(\neg\phi)$ tends to 1, which implies that $p_N(\phi)$ tends to 0.

Note the following interesting consequence of this result. It is not hard to prove that the probability that a random graph contains at least $\frac{1}{2}\binom{N}{2}$ edges converges to $\frac{1}{2}$ as N tends to infinity. Combining this simple observation with the theorem we can deduce that the property "contains at least as many edges as nonedges" cannot be expressed by a first-order formula in $\mathcal{L}_\mathrm{g}$. This is a purely syntactic result, but to prove it we made essential use of model theory.

Further Reading

Shoenfield (2001) is an excellent introduction to logic including the completeness and incompleteness theorems, basic computability theory, and elementary model theory.

The examples described here give only a small part of the flavor for modern model theory. Hodges (1993), Marker (2002), and Poizat (2000) are comprehensive introductions. Marker et al. (1995) contains several introductory articles on the model theory of fields.

In addition to providing tools for analyzing definability in particular structures, a major goal in model theory is proving structure theorems for wide classes of mathematical structures. A key feature is the development by Shelah of notions of dependence generalizing linear dependence in vector spaces and algebraic dependence in fields. Led by Hrushovski and Zilber, model theorists have studied the geometry of dependence and found that frequently it can be used to detect hidden algebraic structure.

In recent years, abstract model theory has found interesting applications in classical mathematics. Hrushovski used these ideas to give a model-theoretic proof of the Mordell–Lang conjecture for function fields in Diophantine geometry. Bouscaren (1998) is an excellent collection of survey articles leading up to Hrushovski's proof.

Bouscaren, E., ed. 1998. *Model Theory and Algebraic Geometry. An Introduction to E. Hrushovski's Proof of the Geometric Mordell–Lang Conjecture.* New York: Springer.

Hodges, W. 1993. *Model Theory.* Encyclopedia of Mathematics and Its Applications, volume 42. Cambridge: Cambridge University Press.

Marker, D. 2002. *Model Theory: An Introduction.* New York: Springer.

Marker, D., M. Messmer, and A. Pillay. 1995. *Model Theory of Fields.* New York: Springer.

Poizat, B. 2000. *A Course in Model Theory. An Introduction to Contemporary Mathematical Logic.* New York: Springer.

Shoenfield, J. 2001. *Mathematical Logic.* Natick, MA: A. K. Peters.