

3 Completeness

In this chapter, we're going to meet some important results about the powers and limitations of first-order logic. Our starting point is the completeness theorem, which shows that there is a mechanical way to check for first-order entailment. We'll see that this positive result is tightly connected to some negative results: the compactness theorem and the Löwenheim-Skolem theorems. These results concern the *size* of models, by which we mean the size of their domain. To fully appreciate their implications, I'll begin with some background about the sizes of sets, which will be needed in later chapters anyway.

3.1 Cardinalities

$\{\text{Athens}\}$ is a set with a one member: the city Athens. We say that $\{\text{Athens}\}$ has size 1. $\{\text{Athens}, \text{Berlin}\}$ has size 2. $\{\text{Athens}, \text{Berlin}, \text{Cairo}\}$ has size 3. And so on. It seems straightforward. But now consider the set \mathbb{N} of natural numbers 0, 1, 2, 3, What is its size?

The official set-theoretic term for the size of a set is *cardinality*. So $\{\text{Athens}, \text{Berlin}, \text{Cairo}\}$ has cardinality 3. Finite sets have a finite cardinalities, which are natural numbers. But infinite sets also have a cardinality. These aren't natural numbers, but numbers of a more general sort, called *cardinals*. The infinite cardinals are also known as the *alephs*, because they are conventionally written using the Hebrew letter 'ℵ' ("aleph"). For example, the cardinality of \mathbb{N} is called 'ℵ₀'. This is the smallest infinite cardinal. The next larger one is ℵ₁, followed by ℵ₂, and so on.

How do we determine the cardinality of an infinite set? The crucial idea goes back to Galileo and Hume, but was only fully developed by Georg Cantor in the 19th century. Following Galileo and Hume, Cantor stipulates that two sets have the same cardinality iff there is a one-to-one correspondence between their members.

To make this precise, we define the notion of a one-to-one correspondence, or bijection.

Definition 3.1

A function f from a set A to a set B is a *bijection* if it satisfies the following two conditions.

- (i) For every $x, y \in A$, if $f(x) = f(y)$ then $x = y$. (Injectivity)
- (ii) For every $b \in B$ there is some $a \in A$ such that $f(a) = b$. (Surjectivity)

Intuitively, a bijection pairs up each element of A with exactly one element of B , and vice versa, so that every element of either set gets a unique partner. As a shorthand, we say that sets A and B are *equinumerous* if there is a bijection from A to B . (In this case, there is always also a bijection from B to A .)

The Galileo-Hume-Cantor principle now says that *two sets have the same cardinality iff they are equinumerous*. Obviously, no finite set of numbers is equinumerous with \mathbb{N} . So \mathbb{N} has an infinite cardinality. We define ' \aleph_0 ' to name this cardinality. Using the Galileo-Hume-Cantor principle, we can determine, for any other set, whether it also has cardinality \aleph_0 .

Consider, for example, the set of odd numbers $1, 3, 5, 7, \dots$. The following function f is a bijection from \mathbb{N} to the set of odd numbers:

$$f(n) = 2n + 1.$$

The function maps 0 to 1, 1 to 3, 2 to 5, and so on. Every natural number is mapped to a unique odd number, and no odd number is left unmapped. So the set of odd numbers also has cardinality \aleph_0 .

Galileo found this paradoxical: how can there be as many odd numbers as natural numbers, given that the odd numbers are a proper subset of the natural numbers? Never mind, said Cantor: the resulting theory of cardinalities is consistent and mathematically fruitful, even if it may initially seem strange.

Sets that are equinumerous with \mathbb{N} are also called *countably infinite* or *denumerable*. A set is *countable* if it is either finite or countably infinite. The word 'countable' alludes to the fact that a bijection between \mathbb{N} and another set effectively assigns a "counter" to each member of the set. For example, the above bijection between \mathbb{N} and the odd numbers assigns the counter 0 to 1, 1 to 3, 2 to 5, and so on.

Exercise 3.1 Show that (a) the set of even natural numbers and (b) the set of integers $\dots, -2, -1, 0, 1, 2, \dots$ are both countably infinite.

Are there any *uncountable* sets? We might suspect that the set of *pairs* of natural numbers is uncountable. But not so. Cantor's *zig-zag method* shows that there is a bijection between \mathbb{N} and the set of pairs of natural numbers. We begin by arranging all pairs of natural numbers in a two-dimensional grid.

	0	1	2	3	4	...
0	(0,0)	(0,1)	(0,2)	(0,3)	(0,4)	...
1	(1,0)	(1,1)	(1,2)	(1,3)	(1,4)	...
2	(2,0)	(2,1)	(2,2)	(2,3)	(2,4)	...
3	(3,0)	(3,1)	(3,2)	(3,3)	(3,4)	...
4	(4,0)	(4,1)	(4,2)	(4,3)	(4,4)	...
...

We then define a path through this grid that visits each cell exactly once. The orange arrows indicate that path. It effectively enumerates all pairs $\langle x, y \rangle$, starting with $\langle 0, 0 \rangle$, followed by $\langle 0, 1 \rangle$, $\langle 1, 0 \rangle$, $\langle 2, 0 \rangle$, $\langle 1, 1 \rangle$, $\langle 0, 2 \rangle$, and so on. The enumeration amounts to a list of all the pairs. We get a bijection to \mathbb{N} by assigning to each pair its position in the list (starting with position 0). Thus $\langle 0, 0 \rangle$ is mapped to 0, $\langle 0, 1 \rangle$ to 1, $\langle 1, 0 \rangle$ to 2, and so forth.

(We can find a formula for this bijection. As we follow the arrow, the pairs *before* any given pair (x, y) comprise all the pairs on the left-to-top diagonals before (x, y) , which have 1, 2, 3, etc. pairs, plus the y pairs on the diagonal containing (x, y) itself. The total number of pairs preceding (x, y) is therefore $(1 + 2 + \dots + (x + y)) + y$. In exercise 1.4, you showed that $1 + 2 + \dots + (x + y) = (x + y)(x + y + 1)/2$. So the position of any pair (x, y) in the enumeration is $(x + y)(x + y + 1)/2 + y$.)

Exercise 3.2 Show that the set of ordered triples of natural numbers is countably infinite.

But it's true that not all sets are countable. Cantor established this with another powerful technique: *diagonalization*.

Theorem 3.1: Cantor's Theorem

The set of all sets of natural numbers is not countable.

Proof. Assume for reductio that the set of all sets of natural numbers is countable. This means that we can list them as S_0, S_1, S_2, \dots . Now consider the set $D = \{n \in \mathbb{N} : n \notin S_n\}$. This is a set of natural numbers, so it must be somewhere in the list. That is, there is some S_k such that $D = S_k$. Now the number k is either in D or not. If k is in D then by definition of D , k is not in S_k . This is impossible, as $D = S_k$. If k is not in D , then by definition of D , k is in S_k . Again, this is impossible. \square

We can again picture this method with a two-dimensional grid.

	0	1	2	3	4	...
S_0	✓	✗	✗	✓	✓	...
S_1	✗	✗	✗	✓	✓	...
S_2	✗	✗	✓	✗	✗	...
S_3	✓	✗	✗	✓	✗	...
S_4	✓	✗	✗	✗	✗	...
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots

Each row represents a set of natural numbers. Each column stands for a number. A checkmark indicates that the number is in the set, a cross that it is not. (So 0 is in S_0 , 1 is not in S_0 , 2 is not in S_0 , and so on.) Cantor's method now looks at the diagonal, in bold. It defines the new set D by *inverting* the diagonal, swapping crosses and checkmarks. In the picture, the inverted diagonal would begin with $\text{✗} \checkmark \text{✗} \text{✗} \checkmark \dots$, meaning that D does not contain 0, does contain 1, does not contain 2 and 3, does contain 4, and so on. The so-defined set D is called the *antidiagonal* of the grid. It can't be anywhere in the list. And so there can be no list of all sets of natural numbers.

Theorem 3.2 can be strengthened:

Theorem 3.2: Cantor’s Theorem (general version)

For any set A , the set of subsets of A has strictly greater cardinality than A .

Proof. The proof proceeds essentially as before. Suppose for reductio that there is a bijection $f : A \rightarrow \mathcal{P}(A)$, where $\mathcal{P}(A)$ (called the *power set* of A) is the set of all subsets of A . Define the antidiagonal set D as the set of all members X of A such that X is not in $f(X)$ – for short: $D = \{X \in A : X \notin f(X)\}$. Since D is a subset of A , it is in $\mathcal{P}(A)$. But it can’t be an output of f . For suppose it is. That is, suppose $D = f(X)$ for some $X \in A$. Either X is in D or not. If X is in D , then by definition of D , X is not in $f(X)$. But $f(X) = D$, so this is impossible. If X is not in D , then by definition of D , X is in $f(X)$. Again, this is impossible. \square

Cantor’s theorem reveals a hierarchy of infinities. Starting with \mathbb{N} , we can produce larger and larger infinities by taking power sets: $\mathcal{P}(\mathbb{N})$ is larger than \mathbb{N} , $\mathcal{P}(\mathcal{P}(\mathbb{N}))$ is even larger, $\mathcal{P}(\mathcal{P}(\mathcal{P}(\mathbb{N})))$ is larger than that, and so on, without end. There are infinitely many infinite cardinals.

Exercise 3.3 Show that if A and B are countably infinite sets, then their union $A \cup B$ is countably infinite.

Exercise 3.4 Show that if a first-order language has countably many non-logical symbols, then the set of all sentences of the language is countably infinite.

3.2 Planning the completeness proof

In section 1.4, we showed that all valid sentences in propositional logic are derivable from the axiom schemata A1–A3 by Modus Ponens. This particular result wasn’t easy to foresee, but it wasn’t a surprise that there is *some* mechanical way of checking if a sentence of propositional logic is valid. Models of propositional logic are essentially finitary. A model is a truth-value assignment to the sentence letters. Since each sentence of propositional logic is composed of finitely many sentence letters, and there are only finitely many ways of assigning truth-values to these sentence letters, it is to be expected that there is a finitary algorithm for deciding whether all of these assignments make the sentence true.

In first-order logic, the situation is very different. A first-order model consists of an

arbitrary set D together with an interpretation of the non-logical vocabulary, where, for example, every subset of D is a possible interpretation of any monadic predicate. Even if D itself is countable, this means that there are usually uncountably many models with that domain. And D doesn't have to be countable. It can have any cardinality whatsoever. The space of first-order models is enormous. Their number is vastly greater than \aleph_0 . There is, in fact, no aleph big enough to measure the number of first-order models. It is therefore astonishing that there is a finitary, mechanical method – a proof system – by which one can find every sentence that is true across the realm of all first-order models. Astonishing, but true, as Kurt Gödel showed in 1929.

Gödel's completeness theorem (not to be confused with his *incompleteness* theorems that we'll discuss in later chapters) is noteworthy for other reasons as well. For example, it supports the conjecture that all intuitively valid mathematical arguments can be formalized as deductions in the first-order calculus. Let's take for granted, for the sake of the argument, that every mathematical statement can be expressed in a first-order language. Now suppose some mathematical argument *can't* be replicated in the first-order calculus. By the completeness theorem, it follows that there is a model in which the premises are true but the conclusion false. This strongly suggests that the argument wasn't valid. We'll meet further applications of the completeness theorem later. First we need to prove it.

Our proof will follow Henkin's 1949 method, which we used in chapter 1 to prove the completeness of propositional logic. Let's recall the key steps.

We want to show that whenever a sentence A is entailed by a set of sentences Γ (meaning that every model that makes all members of Γ true also makes A true), then there is a deduction of A from Γ . For short: if $\Gamma \models A$ then $\Gamma \vdash A$. The proof is by contraposition. We assume $\Gamma \not\models A$ and derive $\Gamma \not\vdash A$, as follows.

1. Assume $\Gamma \not\models A$.
2. Infer that $\Gamma \cup \{\neg A\}$ is consistent.
3. Show that $\Gamma \cup \{\neg A\}$ can be extended to a maximal consistent set Γ^+ .
4. Construct a model \mathfrak{M} based on Γ^+ in which all members of Γ^+ are true.
5. Infer that $\Gamma \not\models A$.

As in chapter 1, we call a set of sentences *consistent* if no contradiction can be derived from it. That is, Γ is consistent iff there is no sentence A such that $\Gamma \vdash A$ and $\Gamma \vdash \neg A$. Equivalently (as you showed in exercise 1.8; the proof carries over from propositional logic), Γ is consistent iff $\Gamma \not\vdash \perp$.

Steps 2 and 5 are easy. The following lemma establishes step 2, and is proved just as

its propositional counterpart, lemma 1.2.

Lemma 3.1

$\Gamma \not\vdash A$ iff $\Gamma \cup \{\neg A\}$ is consistent.

Proof. Suppose $\Gamma \cup \{\neg A\}$ is inconsistent. Then $\Gamma \vdash \neg\neg A$ by RAA and so $\Gamma \vdash A$ by DNE. Contraposing, this means that if $\Gamma \not\vdash A$ then $\Gamma \cup \{\neg A\}$ is consistent. Conversely, suppose $\Gamma \vdash A$. Then $\Gamma, \neg A \vdash A$ by Mon and $\Gamma, \neg A \vdash \neg A$ by Mon and Id. So $\Gamma \cup \{\neg A\}$ is inconsistent. \square

It remains to fill in steps 3 and 4: we need to show that *every consistent set of sentences is satisfiable*.

Here is the plan. Let Γ be a consistent set of sentences in some first-order language \mathcal{L} (with identity and function symbols). We'll show how one can construct a model \mathfrak{M} in which all members of Γ are true. As in chapter 1, we first extend Γ to a maximal consistent set Γ^+ that will guide the construction of the model. For example, if Γ contains $Fa \vee Gb$, it's not obvious if our model should satisfy Fa or Gb or both. Γ^+ will directly contain Fa or Gb or both, so it answer this question.

Suppose we have Fa in Γ^+ . Our model \mathfrak{M} must then contain an object $\llbracket a \rrbracket^{\mathfrak{M}}$ denoted by a that falls in the extension $\llbracket F \rrbracket^{\mathfrak{M}}$ of F . The nature of the object $\llbracket a \rrbracket^{\mathfrak{M}}$ is irrelevant. It proves useful to choose *the individual constant* a as the object denoted by a . Then we can say that the extension of F comprises all individual constants c for which $Fc \in \Gamma^+$. (It might seem odd to have a model whose domain consists of individual constants, and in which each constant denotes itself. But nothing in the definition of a first-order model prevents us from doing this.)

Unfortunately, there are some complications. Suppose Γ contains $\exists xFx$. (Which is short for $\neg\forall x\neg Fx$.) We want this sentence to be true in our model \mathfrak{M} . So there must be some object in $\llbracket F \rrbracket^{\mathfrak{M}}$. But if $\llbracket F \rrbracket^{\mathfrak{M}}$ is defined as the set of individual constants c for which $Fc \in \Gamma^+$, there might be no such object, as there might be no constant c for which $Fc \in \Gamma^+$. We need to ensure that this doesn't happen. That is, we need to ensure that for each existential sentence $\exists xFx$ in Γ^+ , there is a “witnessing” sentence Fc in Γ^+ . The problem is that Γ may already contain $\neg Fc$ for every individual constant c . Then we can't add the required witness without making Γ inconsistent.

To get around this problem, we'll construct Γ^+ in an extend language \mathcal{L}^+ that adds new individual constants to \mathcal{L} . We can use these constants to ensure that every existential sentence in Γ has a witness.

Another complication arises from the presence of the identity symbol. If we let each

individual constant denote itself, any identity statement involving different individual constants will be false. After all, no constant is identical to any other constant. But $a = b$ is consistent, and might be in Γ and therefore in Γ^+ . So we'll actually let each constant denote a *set* of terms: the constant a will denote the set of closed terms t for which $a = t$ is in Γ^+ .

Let's fill in the details.

3.3 The completeness proof

We'll show that every consistent set of first-order sentences is satisfiable. As we've seen above (and as we'll spell out again below), from this it is only a small step to the completeness theorem.

So let Γ be a consistent set of sentences in a first-order language \mathcal{L} . Let \mathcal{L}^+ be an extension of \mathcal{L} with (countably) infinitely many new individual constants. We'll extend Γ to a maximal consistent set in \mathcal{L}^+ . First, though, we need to confirm that Γ itself is still consistent in the extended language \mathcal{L}^+ . Consistency is language-relative because the language determines which instances of the axioms are available. As we switch from \mathcal{L} to \mathcal{L}^+ , new axioms become available. We need to confirm that these new axioms don't allow deriving a contradiction from Γ .

Lemma 3.2

If Γ is a set of \mathcal{L} -sentences that is consistent within \mathcal{L} , and \mathcal{L}^+ extends \mathcal{L} by a set of new individual constants, then Γ is consistent within \mathcal{L}^+ .

Proof by contraposition. Assume that Γ is inconsistent within \mathcal{L}^+ . Then there is a deduction A_1, \dots, A_n of \perp from Γ , where each A_i is an \mathcal{L}^+ -sentence. Being finite, the deduction only uses finitely many of the new constants: call them c_1, \dots, c_k . The deduction also can't use more than finitely many of the old constants in \mathcal{L} . Since first-order languages have infinitely many constants, we can choose k distinct constants d_1, \dots, d_k from \mathcal{L} that don't occur in the deduction. Now consider the sequence of sentences A'_1, \dots, A'_n that results from A_1, \dots, A_n by replacing each c_i by d_i . It is easy to see that

- (i) if A_i is an axiom then so is A'_i ;
- (ii) if $A_i \in \Gamma$ then $A'_i \in \Gamma$ (sentences in Γ don't contain any new constants);
- (iii) if A_i follows from A_1, \dots, A_{i-1} by MP or Gen, then A'_i follows from A'_1, \dots, A'_{i-1} by MP or Gen.

So A'_1, \dots, A'_n is a deduction of \perp from Γ , showing that Γ is inconsistent within \mathcal{L} . \square

Next, we show that Γ can be extended to a maximal consistent set Γ^+ in which every existential sentence $\exists xA$ has a witness $A(x/c)$. Such sets are called *Henkin sets*.

Definition 3.2

A set of sentences Γ in a first-order language \mathcal{L} is a *Henkin set in \mathcal{L}* if it satisfies the following conditions.

- (i) Γ is consistent (within \mathcal{L}).
- (ii) For every \mathcal{L} -sentence A , either $A \in \Gamma$ or $\neg A \in \Gamma$. (Maximality)
- (iii) Whenever Γ contains a sentence of the form $\neg\forall xA$ then it also contains $\neg A(x/c)$ for some individual constant c . (Witnessing)

In the presence of the other conditions, (iii) is equivalent to the requirement that if Γ contains $\neg\forall x\neg A$ then it contains $A(x/c)$ for some c .

Exercise 3.5 Show that if Γ is a Henkin set and $\Gamma \vdash A$, then $A \in \Gamma$.

Exercise 3.6 Show that if Γ is a Henkin set (in a first-order language with identity) then for every closed term t of the language there is an individual constant c such that $t = c$ is in Γ .

Exercise 3.7 Show that if Γ is a Henkin set then Γ contains $\forall xA$ iff Γ contains $A(x/c)$ for every individual constant c .

Lemma 3.3

Every consistent set of \mathcal{L} -sentences Γ can be extended to a Henkin set Γ^+ in any language \mathcal{L}^+ that adds infinitely many individual constants to \mathcal{L} .

Proof. Let Γ be a consistent set of \mathcal{L} -sentences. Let A_1, A_2, \dots be a list of all \mathcal{L}^+ -formulas with exactly one free variable. We define a sequence of sets $\Gamma_0, \Gamma_1, \dots$ as

follows:

$$\begin{aligned}\Gamma_0 &:= \Gamma \\ \Gamma_{n+1} &:= \Gamma_n \cup \{\neg\forall x A_n \rightarrow \neg A_n(x/c_n)\},\end{aligned}$$

where x is the free variable in A_n and c_n is a new \mathcal{L}^+ -constant that does not occur in Γ_n . (There must be some such constant because \mathcal{L}^+ contains infinitely many constants that don't occur in Γ .) Let Γ' be the union $\bigcup_n \Gamma_n$ of all sets in this sequence. (That is, a sentence A is in Γ' iff it is in some Γ_n .)

We show that Γ' is consistent. Suppose not. Then there is a derivation of \perp from Γ and the “Henkin sentences” $\neg\forall x A_n \rightarrow \neg A_n(x/c_n)$. This derivation can use only finitely many of the Henkin sentences. So one of the Γ_n must be inconsistent. But we can show by induction on n that each Γ_n is consistent.

The base case, for $n = 0$, holds by assumption: Γ is consistent.

For the inductive step, assume Γ_n is consistent and suppose for reductio that Γ_{n+1} is inconsistent. By RAA, we then have

$$\Gamma_n \vdash \neg(\neg\forall x A_n \rightarrow \neg A_n(x/c_n)). \quad (1)$$

It's easy to show that if $\Gamma_n \vdash \neg(A \rightarrow B)$ then $\Gamma_n \vdash A$ and $\Gamma_n \vdash \neg B$. (Both $\neg(A \rightarrow B) \rightarrow A$ and $\neg(A \rightarrow B) \rightarrow \neg B$ are truth-functional tautologies.) So (1) implies

$$\Gamma_n \vdash \neg\forall x A_n, \text{ and} \quad (2)$$

$$\Gamma_n \vdash \neg\neg A_n(x/c_n). \quad (3)$$

From (3), we get $\Gamma_n \vdash A_n(x/c_n)$ by DNE. As c_n does not occur in Γ_n , Gen yields

$$\Gamma_n \vdash \forall x A_n. \quad (4)$$

(2) and (4) show that Γ_n is inconsistent, which contradicts the induction hypothesis.

Next, we extend Γ' to a maximal consistent set Γ^+ . The construction follows the proof of Lindenbaum's Lemma (lemma 1.4). Let S_1, S_2, \dots be a list of all \mathcal{L}^+ -sentences. Start-

ing with Γ' , we define another sequence of sets $\Gamma'_0, \Gamma'_1, \dots$:

$$\begin{aligned}\Gamma'_0 &:= \Gamma' \\ \Gamma'_{n+1} &:= \begin{cases} \Gamma'_n \cup \{S_n\} & \text{if } \Gamma'_n \cup \{S_n\} \text{ is consistent,} \\ \Gamma'_n \cup \{\neg S_n\} & \text{otherwise.} \end{cases}\end{aligned}$$

We show by induction that each Γ'_n is consistent. The *base case*, for $n = 0$, holds because Γ'_0 is Γ' , which we've just shown to be consistent. For the *inductive step*, assume that a set Γ'_n in the list is consistent. By lemma 1.3 (which only depends on RAA and therefore also holds for the first-order calculus), either $\Gamma'_n \cup \{S_n\}$ or $\Gamma'_n \cup \{\neg S_n\}$ is consistent. If $\Gamma'_n \cup \{S_n\}$ is consistent, then Γ'_{n+1} is $\Gamma'_n \cup \{S_n\}$ (by construction), so Γ'_{n+1} is consistent. If $\Gamma'_n \cup \{S_n\}$ is not consistent, then Γ'_{n+1} is $\Gamma'_n \cup \{\neg S_n\}$, so again Γ'_{n+1} is consistent.

Let Γ^+ be the union $\bigcup_n \Gamma'_n$ of $\Gamma'_0, \Gamma'_1, \dots$. Evidently, Γ is a subset of Γ^+ . We show that Γ^+ is a Henkin set.

Maximality holds because for each sentence S_n , one of S_n or $\neg S_n$ is in Γ'_{n+1} and therefore in Γ^+ .

To show consistency, suppose that Γ^+ is inconsistent. Then there are sentences A_1, \dots, A_n in Γ^+ from which \perp is deducible. All of these sentences have to occur somewhere on the list S_1, S_2, \dots . Let S_j be the first sentence from S_1, S_2, \dots that occurs after all the A_1, \dots, A_n . Then all A_1, \dots, A_n are in Γ'_j . So Γ'_j is inconsistent. But we've seen that all of the Γ'_n are consistent.

It remains to show that Γ^+ has the witnessing property: whenever $\neg \forall x A$ is in Γ^+ then Γ^+ contains a corresponding sentence $\neg A(x/c)$. Let A be any formula in which x is the only free variable. By construction, Γ' contains $\neg \forall x A \rightarrow \neg A(x/c)$, for some constant c . Since Γ^+ extends Γ' , it also contains this sentence. So if Γ^+ contains $\neg \forall x A$ then it contains $\neg A(x/c)$, by MP and exercise 3.5. \square

Next, we show how to read off a model from a Henkin set. The model's domain will consist of sets of closed terms, so that we can stipulate that each constant c denotes the set of all terms t for which $c = t$ is in the Henkin set. Let's have a closer look at these sets.

Definition 3.3

A binary relation R on some domain D is an *equivalence relation* if it is

- (i) *reflexive*: for every $x \in D$, xRx ;
- (ii) *symmetric*: for every $x, y \in D$, if xRy then yRx ; and
- (iii) *transitive*: for every $x, y, z \in D$, if xRy and yRz then xRz .

Lemma 3.4

If Γ is a Henkin set in a language \mathcal{L} then the relation R that holds between \mathcal{L} -terms t, s iff $t = s \in \Gamma$ is an equivalence relation.


Proof. By A6, $\vdash t = t$. So $t = t \in \Gamma$ by exercise 3.5. Symmetry and transitivity follow similarly from exercise 2.13. \square

An equivalence relation *partitions* the domain over which it is defined into distinct cells so that within each cell, all objects stand in the relation to one another. These cells are called *equivalence classes*. If R is an equivalence relation and x an object in the domain, we write ' $[x]_R$ ' for the equivalence class of R that contains x . That is, $[x]_R = \{y \in D : xRy\}$. If the relation R is clear from the context, we may simply write ' $[x]$ '.

When we're talking about a Henkin set Γ , the relevant equivalence relation is the one defined in lemma 3.4. In what follows, ' $[t]$ ' therefore denotes the set of all terms s for which $t = s \in \Gamma$.

The model \mathfrak{M} that we'll construct from a Henkin set Γ will have as its domain the set of all equivalence classes $[t]$, where t is a closed term in the language of Γ . By exercise 3.6, $[t]$ always contains an individual constant. So we can also say that D is the set of all $[c]$ where c is an individual constant in the language of Γ .

We'll stipulate that the interpretation function I of \mathfrak{M} assigns $[c]$ to each constant c . So we'll have $\llbracket c \rrbracket^{\mathfrak{M}} = [c]$. We'll extend this to function terms, so that $\llbracket f(c) \rrbracket^{\mathfrak{M}} = [f(c)]$. But we can't *directly* stipulate this, because interpretation functions don't assign denotations to complex terms. We have to interpret f as denoting a function on D . By definition 2.9, The denotation of $f(c)$ is the denotation of f applied to the denotation of c . Since the denotation of c is $[c]$, we want the denotation of f to be a function that returns $[f(c)]$



for input $[c]$. So we'll stipulate that for any one-place function symbol f and constant c , $I(f)$ is the function that maps $[c]$ to $[f(c)]$. Similarly for many-place function symbols.

This kind of stipulation can go wrong. Suppose $[c]$ contains two constants c and d , and $[f(c)] \neq [f(d)]$. Our stipulation would entail that $I(f)$ returns $[f(c)]$ for input $[c]$ and $[f(d)]$ for input $[d]$. But if c and d are both in $[c]$ then $[c] = [d]$. And a function can't return two different values for the same input. We must show that this problem can never arise.

A similar issue arises for the interpretation of predicates. To ensure that Ft is in Γ iff $\mathfrak{M} \models Ft$, we'll stipulate that $I(F)$ is the set of all $[t]$ for which $Ft \in \Gamma$. This set isn't well-defined if there are cases where $[t]$ contains two terms t and s for which $Ft \in \Gamma$ but $Fs \notin \Gamma$.

The following lemma shows that neither problem can arise.

Lemma 3.5

If f is an n -ary function symbol, P an n -ary predicate symbol, and $s_1, \dots, s_n, t_1, \dots, t_n$ are terms such that $[s_1] = [t_1], \dots, [s_n] = [t_n]$, then

- (i) $[f(s_1, \dots, s_n)] = [f(t_1, \dots, t_n)]$;
- (ii) $P(s_1, \dots, s_n) \in \Gamma$ iff $P(t_1, \dots, t_n) \in \Gamma$.

Proof. Assume $[s_i] = [t_i]$, for $i = 1, 2, \dots, n$. So Γ contains $s_i = t_i$, for each $i = 1, 2, \dots, n$.

(i). By A6 (and exercise 3.5), $f(s_1, \dots, s_n) = f(s_1, \dots, s_n)$ is in Γ . By n instances of A7 and MP, it follows that $f(s_1, \dots, s_n) = f(t_1, \dots, t_n)$ is in Γ , which entails that $[f(s_1, \dots, s_n)] = [f(t_1, \dots, t_n)]$.

(ii). Assume $P(s_1, \dots, s_n)$ is in Γ . By n instances of A7 and MP, $P(t_1, \dots, t_n)$ is in Γ as well. The converse holds by the same reasoning. \square

Lemma 3.5 ensures that the following definition is legitimate.

Definition 3.4

If Γ be a Henkin set in a language \mathcal{L} then the *Henkin model* \mathfrak{M}_Γ of Γ is defined as follows.

The domain D of \mathfrak{M}_Γ is the set of all $[c]$ where c is an individual constant of \mathcal{L} .

The interpretation function I of \mathfrak{M}_Γ maps

- (i) each individual constant c to $[c]$,
- (ii) each n -ary function symbol f to the function on D that maps any n objects $[t_1], \dots, [t_n]$ to $[f(t_1, \dots, t_n)]$,
- (iii) each (non-logical) n -ary predicate symbol P to the set of all n -tuples $\langle [t_1], \dots, [t_n] \rangle$ such that $P(t_1, \dots, t_n) \in \Gamma$.

Now remember what we're trying to achieve. We want to show that every consistent set of sentences is satisfiable. We've shown in lemma 3.3 that every such set can be extended to a Henkin set. Definition 3.4 tells us how to construct a model from this Henkin set. It remains to show that all members of the Henkin set (and therefore all members of the original set) are true in this model.

We'll need the following two facts.

Lemma 3.6

If \mathfrak{M} is a Henkin model and t a closed term then $\llbracket t \rrbracket^{\mathfrak{M}} = [t]$.

Proof. The proof is by induction on complexity of t . The base case is covered by clause (i) in definition 3.4. So let t be $f(t_1, \dots, t_n)$. Then

$$\begin{aligned} \llbracket f(t_1, \dots, t_n) \rrbracket^{\mathfrak{M}} &= \llbracket f \rrbracket^{\mathfrak{M}}(\llbracket t_1 \rrbracket^{\mathfrak{M}}, \dots, \llbracket t_n \rrbracket^{\mathfrak{M}}) && \text{by def. 2.10} \\ &= \llbracket f \rrbracket^{\mathfrak{M}}([t_1], \dots, [t_n]) && \text{by ind. hyp.} \\ &= [f(t_1, \dots, t_n)] && \text{by def. 3.4. } \square \end{aligned}$$

Lemma 3.7

If \mathfrak{M} is a Henkin model and A a formula then $\mathfrak{M} \models \forall x A$ iff $\mathfrak{M} \models A(x/c)$ for all individual constants c .

Proof. We first show that if $\mathfrak{M}^{d:[c]}$ is a model just like \mathfrak{M} except that it assigns $[c]$ to d , and d does not occur in A , then

$$\mathfrak{M}^{d:[c]} \models A(x/d) \text{ iff } \mathfrak{M} \models A(x/c). \quad (1)$$

There are two cases to consider. If d is c then $\mathfrak{M}^{d:[c]}$ is \mathfrak{M} and (1) holds trivially. If d is a constant other than c , then d does not occur in $A(x/c)$. So $\mathfrak{M}^{d:[c]}$ and \mathfrak{M}

agree on the interpretation of all symbols in $A(x/c)$ and we have $\mathfrak{M}^{d:[c]} \models A(x/c)$ iff $\mathfrak{M} \models A(x/c)$ by the coincidence lemma (lemma 2.1). Also, by the extensionality lemma (lemma 2.2), $\mathfrak{M}^{d:[c]} \models A(x/d)$ iff $\mathfrak{M}^{d:[c]} \models A(x/c)$. So (1) holds.

Now, by definition 2.10, $\mathfrak{M} \models \forall xA$ iff $\mathfrak{M}' \models A(x/d)$ for every model \mathfrak{M}' that differs from \mathfrak{M} at most in the object assigned to d , where d is the alphabetically first constant that does not occur in A . Since each such \mathfrak{M}' has the same domain as \mathfrak{M} , the set of such \mathfrak{M}' is precisely the set of variants $\mathfrak{M}^{d:[c]}$ of \mathfrak{M} . So $\mathfrak{M} \models \forall xA$ iff $\mathfrak{M}^{d:[c]} \models A(x/d)$ for every constant c , which, by (1), is the case iff $\mathfrak{M} \models A(x/c)$ for every constant c . \square

Lemma 3.8: Truth Lemma

If \mathfrak{M} is the Henkin model of a Henkin set Γ in a language \mathcal{L} , then for every \mathcal{L} -sentence A , $\mathfrak{M} \models A$ iff $A \in \Gamma$.

Proof by induction on complexity of A .

Base case: A is atomic. There are two subcases.

Assume that A is an identity sentence $s = t$. Then $\mathfrak{M} \models s = t$ iff $\llbracket s \rrbracket^{\mathfrak{M}} = \llbracket t \rrbracket^{\mathfrak{M}}$ by definition 2.10, iff $\llbracket s \rrbracket = \llbracket t \rrbracket$ by lemma 3.6, iff $s = t \in \Gamma$, by definition of the equivalence classes.

Assume next that A is an atomic sentence $P(t_1, \dots, t_n)$, where P is non-logical. Then $\mathfrak{M} \models P(t_1, \dots, t_n)$ iff $(\llbracket t_1 \rrbracket^{\mathfrak{M}}, \dots, \llbracket t_n \rrbracket^{\mathfrak{M}}) \in \llbracket P \rrbracket^{\mathfrak{M}}$ by definition 2.10, iff $(\llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket) \in \llbracket P \rrbracket^{\mathfrak{M}}$ by lemma 3.6, iff $P(t_1, \dots, t_n) \in \Gamma$ by definition 3.4.

Inductive step: A is composed of other sentences. We have three subcases. The first two, where A is $\neg B$ or $B \rightarrow C$, are easy and go exactly as in lemma 1.6. I'll skip them.

For the final subcase, assume that A is $\forall xB$, for some formula B . We have: $\mathfrak{M} \models \forall xB$ iff $\mathfrak{M} \models B(x/c)$ for every constant c by lemma 3.7, iff $B(x/c) \in \Gamma$ for every constant c by induction hypothesis, iff $\forall xB \in \Gamma$ by exercise 3.7. \square

With that, we have all the ingredients for the completeness proof.

Theorem 3.3: Completeness of the first-order calculus (Gödel 1929)

If $\Gamma \models A$ then $\Gamma \vdash A$.

Proof. We argue by contraposition. Assume $\Gamma \not\models A$. Then $\Gamma \cup \{\neg A\}$ is consistent, by lemma 3.1. By lemma 3.3, $\Gamma \cup \{\neg A\}$ can be extended to a Henkin set Γ^+ in an extended language \mathcal{L}^+ . By the Truth Lemma, the Henkin model of Γ^+ satisfies every sentence in Γ^+ , and therefore every sentence in $\Gamma \cup \{\neg A\}$. So there is a model that satisfies all sentences in Γ but doesn't satisfy A . So $\Gamma \not\models A$. \square

Technically, the model that figures in this proof is not a model for the original language \mathcal{L} of Γ and A , because it also interprets the added individual constants. If this bothers you, we can define an \mathcal{L} -model that falsifies $\Gamma \models A$ by restricting the interpretation function of the Henkin model to \mathcal{L} -constants, without changing the domain. This is called the *reduct* of the Henkin model to \mathcal{L} .

Exercise 3.8 Assume that $\Gamma_0, \Gamma_1, \dots$ are consistent sets of sentences in a first-order language \mathcal{L} , and that each Γ_i is a subset of Γ_{i+1} . Show that their union $\bigcup_i \Gamma_i$ is consistent.

3.4 Unintended Models

We've now shown that the first-order calculus is both sound (theorem 2.4) and complete (theorem 3.3): $\Gamma \models A$ iff $\Gamma \vdash A$. As Gödel pointed out, this has a curious consequence:

Theorem 3.4: Compactness of first-order logic (Gödel 1929)

If every finite subset of a set of sentences is satisfiable then the set itself is satisfiable.

Proof. Let Γ be a set of first-order sentences. Assume that Γ is not satisfiable. So $\Gamma \models \perp$. By completeness, it follows that $\Gamma \vdash \perp$: there is a deduction of \perp from Γ . This deduction can only use finitely many sentences from Γ . So there is a finite subset Γ' of Γ for which $\Gamma' \vdash \perp$. By soundness, $\Gamma' \models \perp$. \square

Why is this curious? Well, for one, it is easy to come up with cases where a conclusion is entailed by infinitely many premises, but not by any finite subset of those premises. For example, consider the premises 'I like the number 0', 'I like the number 1', 'I like the number 2', and so on, for all natural numbers. Together, these premises entail 'I like all natural numbers'. But no finite subset of the premises does. This doesn't contradict the compactness theorem because the inference isn't *logically* valid: it depends on the interpretation of '0', '1', '2', ..., and 'natural number'. Still, it's curious that first-order

logic doesn't allow any such inference to be valid. More directly, compactness implies that there is no way to fully pin down the interpretation of '0', '1', '2', etc. in a first-order language. Let's think through this more carefully.

Many branches of mathematics can be seen as studying certain *mathematical structures*. A mathematical structure consists of a set of objects together with some operations and relations on these objects. For example, arithmetic studies operations and relations on the natural numbers. A formalized, first-order theory of arithmetic will have nonlogical symbols for, say, addition ('+'), multiplication ('×'), and the less-than relation ('<'). We might add an individual constant '*n*' for each number *n*, but for the sake of economy we can instead have a single constant '0' for 0 and another symbol '*s*' for the successor function that maps each number *n* to its successor *n* + 1. Instead of '1', we can then write '*s*(0)'; '2' is '*s*(*s*(0))', and so on.

In *logic*, we abstract away from the meaning of the nonlogical symbols. In *arithmetic*, we don't. In arithmetic, ' $1 + 2 = 3$ ' (that is, ' $s(0) + s(s(0)) = s(s(s(0)))$ ') is a definite claim about the natural numbers. We say that it has an *intended interpretation*, or an *intended model*.

The intended model of arithmetic, also known as the *standard model* of arithmetic, or \mathfrak{A} , has as its domain the set of natural numbers \mathbb{N} ; it interprets '0' as denoting 0, '+' as denoting the addition function +, '×' as denoting the multiplication function ×, '*s*' as denoting the successor function *s*, and '<' as denoting the less-than relation <. We can represent this as a list: $\langle \mathbb{N}, 0, +, \times, s, < \rangle$. The list represents the “structure” of the natural numbers. It identifies a set \mathbb{N} and some operations and relations on that set that are picked out by the nonlogical symbols of the language. (0 counts as a zero-ary operation.)

Of course, the language of arithmetic also has unintended models. For example, we can let the domain be { Athens, Berlin } and use an interpretation function that maps '0' to Athens, '+' and '×' to the function that maps any pair of cities to Athens, '*s*' to the function that maps each city to itself, and '<' to the empty relation. In this model, ' $s(0) + s(s(0)) = s(s(s(0)))$ ' is true, but so is ' $s(0) = 0$ '.

Exercise 3.9 Is Lagrange's Theorem (every natural number is the sum of four squares) true in this model?

We might hope that such unintended models can always be ruled out by laying down sufficiently many postulates in the formal language of arithmetic. For example, if our theory of arithmetic contains ' $s(0) \neq 0$ ' then the above model (with Athens and Berlin)

is no longer a model of the theory. So we can rule out *some* unintended models. The compactness theorem dashes any hope of ruling out *all* unintended models. Let $\text{Th}(\mathfrak{A})$ be the set of all sentences true in the standard model \mathfrak{A} . This is called the *theory of \mathfrak{A}* . It contains all postulates we could possibly lay down, assuming that they should all be true in the standard model. By compactness, there is a model of $\text{Th}(\mathfrak{A})$ whose domain includes junk elements that clearly aren't natural numbers.

Theorem 3.5: Non-standard Models of Arithmetic

$\text{Th}(\mathfrak{A})$ has models in which some object is not reachable from 0 by finitely many applications of the successor function. (Such models are called *non-standard models* of arithmetic.)

Proof. Let \mathcal{L}_A^+ be the language of arithmetic with an added individual constant c . Let $\text{Th}(\mathfrak{A})^+$ be the set of sentences that extends $\text{Th}(\mathfrak{A})$ by

$$c \neq 0, \quad c \neq s(0), \quad c \neq s(s(0)), \quad c \neq s(s(s(0))), \quad \dots$$

So $\text{Th}(\mathfrak{A})^+$ says that c is different from 0, 1, 2, and so on, for all natural numbers. Every finite subset of $\text{Th}(\mathfrak{A})^+$ is obviously satisfiable: it is true in the standard model \mathfrak{A} , interpreting c as a sufficiently large number. By the compactness theorem, it follows that $\text{Th}(\mathfrak{A})^+$ is satisfiable. Any model of $\text{Th}(\mathfrak{A})^+$ must have an element (denoted by c) that is not identical to any natural number: it can't be reached from 0 by finitely many applications of the successor function. The reduct of any such model to the original language of arithmetic (discarding the interpretation of c) is a model of $\text{Th}(\mathfrak{A})$. \square

Compactness thus reveals a deep expressive limitation of first-order logic: the structure of the natural numbers can't be captured in a first-order language.

Here is a simpler example. For each $n > 0$, it is easy to find a first-order sentence S_n that is true in all and only the models with exactly n elements. For example, $\exists x \exists y (x \neq y \wedge \forall z (z = x \vee z = y))$ is true in all and only models with exactly two elements. (Compare exercise 2.16.) In that sense, we can capture the intended size of a domain, as long as that size is a fixed finite number. But there is no first-order sentence that is true in all and only the models with infinitely many elements.

To see why, let S_∞ be such a sentence. Its negation $\neg S_\infty$ would be true in all and only the finite models. Now consider the set consisting of $\neg S_\infty$ together with $\neg S_1, \neg S_2, \neg S_3, \dots$. All finite subsets of this set are satisfiable. By compactness, the whole set is satisfiable. But there is no model whose domain not infinite, and yet also has no finite size.

The following theorems show that the situation is even worse. In section 3.1, we saw that there are many levels of infinity: many infinite cardinalities. None of them can be captured in first-order logic, insofar as there is no sentence (nor even a set of sentences) that would force a model to have a particular infinite cardinality, or even to fall in any non-trivial range of infinite cardinalities.

Theorem 3.6: The (Downward) Löwenheim-Skolem Theorem

If a set of sentences in a countable first-order language has a model, then it has a countable model.

Proof. Let Γ be such a set. By lemma 3.3, Γ can be extended to a Henkin set Γ^+ . By lemma 3.8, the Henkin model of Γ^+ is a model of Γ . Its domain consists of equivalence classes of closed terms in the language of Γ^+ . If the language of Γ is countable, then so is the language of Γ^+ . So it has only countably many closed terms. \square

Theorem 3.7: The Upward Löwenheim-Skolem Theorem (Tarski 1935)

If a set of sentences in a countable first-order language has an infinite model, then it has models of every infinite cardinality.

Proof sketch. The proof requires relaxing our stipulation that a first-order language must only have countably many individual constants. The completeness theorem can be proved without this assumption (although the proof becomes more complicated, as we can no longer appeal to enumerations of the sentences in the language). Compactness still follows from completeness, and we get a slightly generalized downward Löwenheim-Skolem theorem according to which every satisfiable set of sentences has a model whose cardinality is at most equal to the cardinality of the language.

Now let Γ be a set of first-order sentences in a countable language with an infinite model. From the downward theorem, we know that Γ has a countably infinite model \mathfrak{M} . Let κ be any cardinal greater than \aleph_0 . Expand the language of Γ by κ new individual constants. For each pair of these constants c_i and c_j , add the sentence $c_i \neq c_j$ to Γ , thereby creating the set Γ^+ . All finite subsets of Γ^+ are satisfied in \mathfrak{M} , with the (finitely many) new constants in the set interpreted as distinct objects in \mathfrak{M} . By the compactness theorem (for uncountable languages), Γ^+ has a model \mathfrak{M}^+ . All the κ new constants denote distinct objects in this model, so \mathfrak{M}^+ has at least κ objects. By the generalized

downward theorem, it follows that Γ^+ has a model whose cardinality is exactly κ . The reduct of that model to the language of Γ is a model of Γ with cardinality κ . \square

The downward theorem was proved by Thoralf Skolem in 1920, building on an earlier proof sketch by Leopold Löwenheim. The (ill-named) upward theorem, due to Alfred Tarski, requires compactness and could only be proved after 1929. It entails, among other things, that $\text{Th}(\mathfrak{A})$ has non-standard models of every infinite cardinality.

Exercise 3.10 (a) Can you find a first-order sentence that is only true in infinite models? (Hint: let f be a function that is injective, but not surjective.) (b) Is the negation of this sentence only true in finite models?

Exercise 3.11 Explain why every satisfiable set of first-order sentences has a model whose domain is a set of natural numbers.

Exercise 3.12 Consider a first-order language with a binary predicate symbol ' P ' for the "parent" relation, so that Pxy means (on its intended interpretation) that x is a parent of y . We can then define the "grandparent" relation $P^2(x, y)$ as $\exists z_1 (Pxz_1 \wedge Pz_1y)$, the "great-grandparent" relation $P^3(x, y)$ as $\exists z_1 \exists z_2 (Pxz_1 \wedge Pz_1z_2 \wedge Pz_2y)$, and so on. Show that there is no way to define the "ancestor" relation (no matter what other nonlogical symbols we add to the language): there can be no formula $A(x, y)$ that is equivalent to the infinite disjunction $P(x, y) \vee P^2(x, y) \vee P^3(x, y) \vee \dots$

Exercise 3.13 Let \mathfrak{A}^c be exactly like the standard model of arithmetic \mathfrak{A} , except that the number 2 is replaced by Julius Caesar: the domain of \mathfrak{A}^c is $\{0, 1, \text{Caesar}, 3, 4, \dots\}$; ' s ' denotes a function that maps 0 to 1, 1 to Caesar, Caesar to 3, ...; ' $+$ ' denotes a function that maps 1 and 1 to Caesar, 1 and Caesar to 3, and so on. Is \mathfrak{A}^c a model of $\text{Th}(\mathfrak{A})$? In what way is it less interesting than the non-standard models of theorem 3.5?