When we ask how *many* different models a sentence or set of sentences may have of a given size, the answer is disappointing: there are always an unlimited number (a nonenumerable infinity) of models if there are any at all. To give a completely trivial example, consider the empty language, with identity but no nonlogical predicates, for which an interpretation is just a nonempty set to serve as domain. And consider the sentence $\exists x \forall y(y = x)$, which says there is just one thing in the domain. For any object a you wish, the interpretation whose domain is $\{a\}$, the set whose only element is a, is a model of this sentence. So for each real number, or each point on the line, we get a model.

Of course, these models all 'look alike': each consists of just one thing, sitting there doing nothing, so to speak. The notion of *isomorphism*, which we are about to define, is a technically precise way of saying what is meant by 'looking alike' in the case of *non*trivial languages. Two interpretations $\mathcal{P}$ and $\mathcal{Q}$ of the same language L are *isomorphic* if and only if there is a *correspondence* j between individuals p in the domain $|\mathcal{P}|$ and individuals q in the domain $|\mathcal{Q}|$ subject to certain conditions. (The definition of correspondence, or total, one-to-one, onto function, has been given in the problems at the end of Chapter 1.) The further conditions are that for every n-place predicate R and all $p_1, \ldots, p_n$ in $|\mathcal{P}|$ we have

$$\text{(I1)} \qquad R^{\mathcal{P}}(p_1, \ldots, p_n) \quad \text{if and only if} \quad R^{\mathcal{Q}}(j(p_1), \ldots, j(p_n))$$

and for every constant c we have

$$\text{(I2)} \qquad j(c^{\mathcal{P}}) = c^{\mathcal{Q}}.$$

If function symbols are present, it is further required that for every n-place function symbol f and all $p_1, \ldots, p_n$ in $|\mathcal{P}|$ we have

$$\text{(I3)} \qquad j(f^{\mathcal{P}}(p_1, \ldots, p_n)) = f^{\mathcal{Q}}(j(p_1), \ldots, j(p_n)).$$

12.3 Example (Inverse order and mirror arithmetic). Consider the language with a single two-place predicate $<$, the interpretation with domain the natural numbers $\{0, 1, 2, 3, \ldots\}$ and with $<$ denoting the usual strict less-than order relation, and by contrast the interpretation with domain the nonpositive integers $\{0, -1, -2, -3, \ldots\}$ and with $<$ denoting the usual strict greater-than relation. The correspondence associating n with $-n$ is an isomorphism, since m is less than n if and only if $-m$ is greater than $-n$, as required by (I1).

If we also let **0** denote zero, let $'$ denote the predecessor function, which takes x to $x - 1$, let $+$ denote the addition function, and let $\cdot$ denote the function taking x and y to the negative of their product, $-xy$, then we obtain an interpretation isomorphic to the standard interpretation of the language of arithmetic. For the following equations show (I3) to be fulfilled:

$$\begin{aligned} -x - 1 &= -(x + 1) \\ (-x) + (-y) &= -(x + y) \\ -(-x)(-y) &= -xy. \end{aligned}$$

Generalizing our completely trivial example, in the case of the empty language, where an interpretation is just a domain, two interpretations are isomorphic if and only if there is a correspondence between their domains (that is, if and only if they are equinumerous, as defined in the problems at the end of Chapter 1). The analogous property for nonempty languages is stated in the next result.

12.4 Proposition. Let X and Y be sets, and suppose there is a correspondence j from X to Y. Then if $\mathcal{Y}$ is any interpretation with domain Y, there is an interpretation $\mathcal{X}$ with domain X such that $\mathcal{X}$ is isomorphic to $\mathcal{Y}$. In particular, for any interpretation with a finite domain having n elements, there is an isomorphic interpretation with domain the set $\{0, 1, 2, \ldots, n - 1\}$, while for any interpretation with a denumerable domain there is an isomorphic interpretation with domain the set $\{0, 1, 2, \ldots\}$ of natural numbers.

Proof: For each relation symbol R, let $R^{\mathcal{X}}$ be the relation that holds for $p_1, \ldots, p_n$ in X if and only if $R^{\mathcal{Y}}$ holds for $j(p_1), \ldots, j(p_n)$. This makes (I1) hold automatically. For each constant c, let $c^{\mathcal{X}}$ be the unique p in X such that $j(p) = c^{\mathcal{Y}}$. (There will be such a p because j is onto, and it will be unique because j is one-to-one.) This makes (I2) hold automatically. If function symbols are present, for each function symbol f, let $f^{\mathcal{X}}$ be the function on X whose value for $p_1, \ldots, p_n$ in X is the unique p such that $j(p) = f^{\mathcal{Y}}(j(p_1), \ldots, j(p_n))$. This makes (I3) hold automatically.

The next result is a little more work. Together with the preceding, it implies what we hinted earlier, that a sentence or set of sentences has an unlimited number of models if it has any models at all: given one model, by the preceding proposition there will be an unlimited number of interpretations isomorphic to it, one for each set equinumerous with its domain. By the following result, these isomorphic interpretations will all be models of the given sentence or set of sentences.

12.5 Proposition (Isomorphism lemma). If there is an isomorphism between two interpretations $\mathcal{P}$ and $\mathcal{Q}$ of the same language L, then for every sentence A of L we have

(1) $$\mathcal{P} \models A \quad \text{if and only if} \quad \mathcal{Q} \models A.$$

Proof: We first consider the case where identity and function symbols are absent, and proceed by induction on complexity. First, for an atomic sentence involving a nonlogical predicate R and constants $t_1, \ldots, t_n$, the atomic clause in the definition of truth gives

$$\mathcal{P} \models R(t_1, \ldots, t_n) \quad \text{if and only if} \quad R^{\mathcal{P}}\big(t_1^{\mathcal{P}}, \ldots, t_n^{\mathcal{P}}\big)$$
$$\mathcal{Q} \models R(t_1, \ldots, t_n) \quad \text{if and only if} \quad R^{\mathcal{Q}}\big(t_1^{\mathcal{Q}}, \ldots, t_n^{\mathcal{Q}}\big)$$

while the clause (I1) in the definition of isomorphism gives

$$R^{\mathcal{P}}\big(t_1^{\mathcal{P}}, \ldots, t_n^{\mathcal{P}}\big) \quad \text{if and only if} \quad R^{\mathcal{Q}}\big(j\big(t_1^{\mathcal{P}}\big), \ldots, j\big(t_n^{\mathcal{P}}\big)\big)$$

and the clause (I2) in the definition of isomorphism gives

$$R^{\mathcal{Q}}\big(j\big(t_1^{\mathcal{P}}\big), \ldots, j\big(t_n^{\mathcal{P}}\big)\big) \quad \text{if and only if} \quad R^{\mathcal{Q}}\big(t_1^{\mathcal{Q}}, \ldots, t_n^{\mathcal{Q}}\big).$$

Together the four displayed equivalences give (1) for $R(t_1, \ldots, t_n)$.

Second, suppose (1) holds for less complex sentences than $\sim F$, including the sentence F. Then (1) for $\sim F$ is immediate from this assumption together with the negation clause in the definition of truth, by which we have

$$\mathcal{P} \models \sim F \quad \text{if and only if} \quad \text{not } \mathcal{P} \models F$$
$$\mathcal{Q} \models \sim F \quad \text{if and only if} \quad \text{not } \mathcal{Q} \models F.$$

The case of junctions is similar.

Third, suppose (1) holds for less complex sentences than $\forall x\, F(x)$, including sentences of the form $F(c)$. For any element p of $|\mathcal{P}|$, if we extend the language by adding a new constant c and extend the interpretation $\mathcal{P}$ so that c denotes p, then there is one and only one way to extend the interpretation $\mathcal{Q}$ so that j remains an isomorphism of the extended interpretations; namely, we extend the interpretation $\mathcal{Q}$ so that c denotes $j(p)$, and therefore clause (I2) in the definition of isomorphism still holds for the extended language. By our assumption that (1) holds for $F(c)$ it follows on the one hand that

(2) $$\mathcal{P} \models F[p] \quad \text{if and only if} \quad \mathcal{Q} \models F[j(p)].$$

By the universal quantifier clause in the definition of truth

$$\mathcal{P} \models \forall x\, F(x) \quad \text{if and only if} \quad \mathcal{P} \models F[p] \text{ for all } p \text{ in } |\mathcal{P}|.$$

Hence

$$\mathcal{P} \models \forall x\, F(x) \quad \text{if and only if} \quad \mathcal{Q} \models F[j(p)] \text{ for all } p \text{ in } |\mathcal{P}|.$$

On the other hand, again by the universal quantifier clause in the definition of truth we have

$$\mathcal{Q} \models \forall x\, F(x) \quad \text{if and only if} \quad \mathcal{Q} \models F[q] \text{ for all } q \text{ in } |\mathcal{Q}|.$$

But since j is a correspondence, and therefore is onto, *every* q in $|\mathcal{Q}|$ is of the form $j(p)$, and (1) follows for $\forall x\, F(x)$. The existential-quantifier case is similar.

If identity is present, we have to prove (1) also for atomic sentences involving $=$. That is, we have to prove

$$p_1 = p_2 \quad \text{if and only if} \quad j(p_1) = j(p_2).$$

But this is simply the condition that j is one-to-one, which is part of the definition of being a correspondence, which in turn is part of the definition of being an isomorphism.

If function symbols are present, we must first prove as a preliminary that for any closed term t we have

$$j(t^{\mathcal{P}}) = t^{\mathcal{Q}}. \tag{3}$$

This is proved by induction on complexity of terms. For constants we have (3) by clause (I2) in the definition of isomorphism. And supposing (3) holds for $t_1, \ldots, t_n$, then it holds for $f(t_1, \ldots, t_n)$ since by clause (I3) in the definition of isomorphism we have

$$\begin{aligned} j((f(t_1, \ldots, t_n))^{\mathcal{P}}) &= j\big(f^{\mathcal{P}}\big(t_1^{\mathcal{P}}, \ldots, t_n^{\mathcal{P}}\big)\big) \\ &= f^{\mathcal{Q}}\big(j\big(t_1^{\mathcal{P}}\big), \ldots, j\big(t_n^{\mathcal{P}}\big)\big) = f^{\mathcal{Q}}\big(t_1^{\mathcal{Q}}, \ldots, t_n^{\mathcal{Q}}\big) = (f(t_1, \ldots, t_n))^{\mathcal{Q}}. \end{aligned}$$

The proof given above for the atomic case of (1) now goes through even when the t_i are complex closed terms rather than constants, and no further changes are required in the proof.

12.6 Corollary (Canonical-domains lemma).

(a) Any set of sentences that has a finite model has a model whose domain is the set $\{0, 1, 2, \ldots, n\}$ for some natural number n.

(b) Any set of sentences having a denumerable model has a model whose domain is the set $\{0, 1, 2, \ldots\}$ of natural numbers.

Proof: Immediate from Propositions 12.4 and 12.5.

Two models that are isomorphic are said to be of the same *isomorphism type*. The intelligent way to count the models of a given size that a sentence has is to count not literally the number of models (which is always a nonenumerable infinity if it is nonzero), but the number of isomorphism types of models. The import of the rather abstract results of this section should become clearer as they are illustrated concretely in the next section.

12.3 The Löwenheim–Skolem and Compactness Theorems

We have seen that there are sentences that have only infinite models. One might wonder whether there are sentences that have only *nonenumerable* models. We have also seen that there are enumerable sets of sentences that have only infinite models, though every finite subset has a finite model. One might wonder whether there are sets of sentences that have *no models at all*, though every finite subset has a model. The answer to both these questions is negative, according to the following pair of theorems. They are basic results in the theory of models, with many implications about the existence, size, and number of models.

12.14 Theorem (Löwenheim–Skolem theorem). If a set of sentences has a model, then it has an enumerable model.

12.15 Theorem (Compactness theorem). If every finite subset of a set of sentences has a model, then the whole set has a model.

We explore a few of the implications of these theorems in the problems at the end this chapter. We stop here just to note three immediate implications.

12.16 Corollary (Overspill principle). If a set of sentences has arbitrarily large finite models, then it has a denumerable model.

Proof: Let Γ be a set of sentences having arbitrarily large finite models, and for each m let I_m be the sentence with identity but no nonlogical symbols considered in Example 12.1, which is true in a model if and only if the model has size $\geq m$. Let

$$\Gamma^* = \Gamma \cup \{I_1, I_2, I_3, \ldots\}$$

be the result of adding all the I_m to Γ. Any finite subset of Γ^* is a subset of $\Gamma \cup \{I_1, I_2, \ldots, I_m\}$ for some m, and since Γ has a model of size $\geq m$, such a set has a model. By the compactness theorem, therefore, Γ^* has a model. Such a model is of course a model of Γ, and being also a model of each I_m, it has size $\geq m$ for all finite m, and so is infinite. By the Löwenheim–Skolem theorem, we could take it to be enumerable.

A set Γ of sentences is (*implicationally*) *complete* if for every sentence A in its language, either A or $\sim A$ is a consequence of Γ, and *denumerably categorical* if any two denumerable models of Γ are isomorphic.

12.17 Corollary (Vaught's test). If Γ is a denumerably categorical set of sentences having no finite models, then Γ is complete.

Proof: Suppose Γ is not complete, and let A be some sentence in its language such that neither A nor $\sim A$ is a consequence of Γ. Then both $\Gamma \cup \{\sim A\}$ and $\Gamma \cup \{A\}$ are satisfiable, and by the Löwenheim–Skolem theorem they have enumerable models $\mathcal{P}^-$ and $\mathcal{P}^+$. Since Γ has no finite models, $\mathcal{P}^-$ and $\mathcal{P}^+$ must be denumerable. Since Γ is denumerably categorical, they must be isomorphic. But by the isomorphism lemma, since A is untrue in one and true in the other, they can*not* be isomorphic. So the assumption that Γ is not complete leads to a contradiction, and Γ must be complete after all.

Thus if Γ is any of the examples of the preceding section in which we found there was only one isomorphism type of denumerable model, then adding the sentences $I_1, I_2, I_3, \ldots$ to Γ (in order to eliminate the possibility of finite models) produces an example that is complete.

The Löwenheim–Skolem theorem also permits a sharpening of the statement of the canonical-domains lemma (Lemma 12.6).

12.18 Corollary (Canonical-domains theorem).

(a) Any set of sentences that has a model, has a model whose domain is either the set of natural numbers $< n$ for some positive n, or else the set of all natural numbers.

(b) Any set of sentences not involving function symbols or identity that has a model, has a model whose domain is the set of all natural numbers.

Proof: (a) is immediate from the Löwenheim–Skolem theorem and Corollary 12.6. For (b), given a set of sentences Γ not involving function symbols or identity, if Γ has a model, apply part (a) to get, at worst, a model $\mathcal{Y}$ with domain the finite set $\{0, 1, \ldots, n-1\}$ for some n. Let f be the function from the set of all natural numbers to this finite set given by $f(m) = \min(m, n-1)$. Define an interpretation $\mathcal{X}$ with domain the set of all natural numbers by assigning to each k-place relation symbol R as denotation the relation $R^{\mathcal{X}}$ that holds for $p_1, \ldots, p_k$ if and only if $R^{\mathcal{Y}}$ holds for $f(p_1), \ldots, f(p_k)$. Then f has all the properties of an isomorphism except for not being one-to-one. Examining the proof of the isomorphism lemma (Proposition 12.5), which tells us the same sentences are true in isomorphic interpretations, we see that the property of being one-to-one was used *only* in connection with sentences involving identity. Since the sentences in Γ do not involve identity, they will be true in $\mathcal{X}$ because they are true in $\mathcal{Y}$.

The remainder of this section is devoted to an advance description of what will be done in the following two chapters, which contain proofs of the Löwenheim–Skolem and compactness theorems and a related result. Our preview is intended to enable the readers who are familiar with the contents of an introductory textbook to decide how much of this material they need or want to read. The next chapter, Chapter 13, is devoted to a proof of the compactness theorem. Actually, the proof shows that if every finite subset of a set Γ has a model, then Γ has an *enumerable* model. This version of the compactness theorem implies the Löwenheim–Skolem theorem, since if a set has a model, so does every subset, and in particular every finite subset. An optional final section 13.5 considers what happens if we admit *nonenumerable* languages. (It turns out that the compactness theorem still holds, but the 'downward' Löwenheim–Skolem theorem fails, and one gets instead an 'upward' theorem to the effect that any set of sentences having an infinite model has a nonenumerable model.)

Problems

12.1 By the *spectrum* of a sentence C (or set of sentences Γ) is meant the set of all positive integers n such that C (or Γ) has a finite model with a domain having exactly n elements. Consider a language with just two nonlogical symbols, a one-place predicate P and a one-place function symbol f. Let A be the following sentence:

$$\forall x_1 \forall x_2 (f(x_1) = f(x_2) \rightarrow x_1 = x_2)\ \&$$
$$\forall y \exists x (f(x) = y)\ \&$$
$$\forall x \forall y (f(x) = y \rightarrow (Px \leftrightarrow \sim Py)).$$

Show that the spectrum of A is the set of all even positive integers.

12.2 Give an example of a sentence whose spectrum is the set of all odd positive integers.

12.3 Give an example of a sentence whose spectrum is the set of all positive integers that are perfect squares.

12.4 Give an example of a sentence whose spectrum is the set of all positive integers divisible by three.

12.5 Consider a language with just one nonlogical symbol, a two-place predicate Q. Let $\mathcal{U}$ be the interpretation in which the domain consists of the four sides of a square, and the denotation of Q is the relation between sides of being parallel. Let $\mathcal{V}$ be the interpretation in which the domain consists of the four vertices of a square, and the denotation of Q is the relation between vertices of being diagonally opposite. Show that $\mathcal{U}$ and $\mathcal{V}$ are isomorphic.

12.6 Consider a language with just one nonlogical symbol, a two-place predicate $<$. Let $\mathcal{Q}$ be the interpretation in which the domain is the set of real numbers *strictly greater than zero and strictly less than one* and the denotation of $<$ is the usual order relation. Let $\mathcal{R}$ be the interpretation in which the domain is the set of *all* real numbers and the denotation of $<$ is the usual order relation. Show that $\mathcal{Q}$ and $\mathcal{R}$ are isomorphic.

12.7 Let L be a language whose only nonlogical symbols are a two-place function symbol § and a two-place predicate $<$. Let $\mathcal{P}$ be the interpretation of this language in which the domain is the set of *positive* real numbers, the denotation of § is the usual *multiplication* operation, and the denotation of $<$ is the usual order relation. Let $\mathcal{Q}$ be the interpretation of this language in which the domain is the set of *all* real numbers, the denotation of § is the usual *addition* operation, and the denotation of $<$ is the usual order relation. Show that $\mathcal{P}$ and $\mathcal{Q}$ are isomorphic.

12.8 Write $\mathcal{A} \cong \mathcal{B}$ to indicate that $\mathcal{A}$ is isomorphic to $\mathcal{B}$. Show that for all interpretations $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$ of the same language the following hold:

(a) $\mathcal{A} \cong \mathcal{A}$;

(b) if $\mathcal{A} \cong \mathcal{B}$, then $\mathcal{B} \cong \mathcal{A}$;

(c) if $\mathcal{A} \cong \mathcal{B}$ and $\mathcal{B} \cong \mathcal{C}$, then $\mathcal{A} \cong \mathcal{C}$.

12.9 By *true arithmetic* we mean the set Γ of all sentences of the language of arithmetic that are true in the standard interpretation. By a *nonstandard model of arithmetic* we mean a model of this Γ that (unlike the model in Example 12.3) is not isomorphic to the standard interpretation. Let Δ be the set of sentences obtained by adding a constant c to the language and adding the sentences $c \neq \mathbf{0}$, $c \neq \mathbf{1}$, $c \neq \mathbf{2}$, and so on, to Γ. Show that any model of Δ would give us a nonstandard model of arithmetic.

12.10 Consider the language with just the one nonlogical symbol $\equiv$ and the sentence Eq whose models are precisely the sets with equivalence relations, as in the examples in section 12.2.

(a) For each n, indicate how to write down a sentence B_n such that the models of Eq & B_n will be sets with equivalence relations *having at least n equivalence classes*.

(b) For each n, indicate how to write down a formula $F_n(x)$ such that in a model of Eq, an element a of the domain will satisfy $F_n(x)$ if and only if there are at least n elements in the equivalence class of a.

(c) For each n, indicate how to write down a sentence C_n that is true in a model of Eq if and only if there are exactly n equivalence classes.

(d) For each n, indicate how to write down a formula $G_n(x)$ that is satisfied by an element of the domain if and only if its equivalence class has exactly n elements.

12.11 For each m and n indicate how to write down a sentence D_{mn} that is true in a model of Eq if and only if there are at least m equivalence classes with exactly n elements.

12.12 Show that if two models of *Eq* are isomorphic, then the equivalence relations of the models have the same signature.

12.13 Suppose E_1 and E_2 are equivalence relations on denumerable sets X_1 and X_2 both having the signature $\sigma(n) = 0$ for $n \geq 1$ and $\sigma(0) = \infty$, that is, both having infinitely many equivalence classes, all infinite. Show that the models involved are isomorphic.

12.14 Show that two denumerable models of *Eq* are isomorphic if and only if they have the same signature.

In the remaining problems you may, when relevant, use the Löwenheim–Skolem and compactness theorems, even though the proofs have been deferred to the next chapter.

12.15 Show that:

(**a**) Γ is unsatisfiable if and only if $\sim C_1 \vee \cdots \vee \sim C_m$ is valid for some $C_1, \ldots, C_m$ in Γ.

(**b**) D is a consequence of Γ if and only if D is a consequence of some finite subset of Γ.

(**c**) D is a consequence of Γ if and only if $\sim C_1 \vee \cdots \vee \sim C_m \vee D$ is valid for some $C_1, \ldots, C_m$ in Γ.

12.16 For any prime $p = 2, 3, 5, \ldots$, let $D_p(x)$ be the formula $\exists y\ \mathbf{p} \cdot y = x$ of the language of arithmetic, so that for any natural number n, $D_p(\mathbf{n})$ is true if and only if p divides n without remainder. Let S be any set of primes. Say that a nonstandard model $\mathcal{M}$ of arithmetic *encrypts* S if there is an individual m in the domain $|\mathcal{M}|$ such that $\mathcal{M} \models D_p[m]$ for all p belonging to S, and $\mathcal{M} \models \sim D_p[m]$ for all p not belonging to S. Show that for any set S of primes there is a denumerable nonstandard model of arithmetic that encrypts S.

12.17 Show that there are nonenumerably many isomorphism types of denumerable nonstandard models of arithmetic.

12.18 Show that if two sentences have the same enumerable models, then they are logically equivalent.

12.19 Work with a language whose only nonlogical symbol is a single two-place predicate $<$. Consider the set of sentences of this language that are true in the interpretation where the domain is the set of real numbers and the denotation of the predicate is the usual order on real numbers. According to the Löwenheim–Skolem theorem, there must be an *enumerable* model of this set of sentences. Can you guess what one is?

The next several problems provide a significant example of a denumerably categorical set of sentences.

12.20 Work with a language whose only nonlogical symbol is a single two-place predicate $<$. The models of the following sentence *LO* of the language are called *linear orders*:

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

Such a model $\mathcal{A}$ will consist of a nonempty set $|\mathcal{A}|$ or A and a two-place relation $<^{\mathcal{A}}$ or $<_A$ on it. Show that the above sentence implies

$$\forall x \forall y \sim (x < y \;\&\; y < x).$$

12.21 Continuing the preceding problem, a *finite partial isomorphism* between linear orders $(A, <_A)$ and $(B, <_B)$ is a function j from a finite subset of A onto a finite subset of B such that for all a_1 and a_2 in the domain of j, $a_1 <_A a_2$ if and only if $j(a_1) <_A j(a_1)$. Show that if j is a finite partial isomorphism from a linear order $(A, <_A)$ to the rational numbers with their usual order, and a is any element of A not in the domain of j, then j can be extended to a finite partial isomorphism whose domain is the domain of j together with a. (Here *extended* means that the new isomorphism assigns the same rational numbers as the old to elements of A there were already in the domain of the old.)

12.22 Continuing the preceding problem, if $j_0, j_1, j_2, \ldots$ are finite partial isomorphisms from an enumerable linear order to the rational numbers with their usual order, and if each j_{i+1} is an extension of the preceding j_i, and if every element of A is in the domain of one of the j_i (and hence of all j_k for $k \geq i$), then $(A, <_A)$ is isomorphic to some *suborder* of the rational numbers with their usual order. (Here *suborder* means a linear order $(B, <_B)$ where B is some subset of the rational numbers, and $<_B$ the usual order on rational numbers as it applies to elements of this subset.)

12.23 Continuing the preceding problem, show that every enumerable linear order $(A, <_A)$ is isomorphic to a suborder of the rational numbers with their usual order.

12.24 Continuing the preceding problem, a linear order is said to be *dense* if it is a model of

$$\forall x \forall y (x < y \rightarrow \exists z (x < z \;\&\; z < y)).$$

It is said to have *no endpoints* if it is a model of

$$\sim\exists x \forall y (x < y \vee x = y) \;\&\; \sim\exists x \forall y (x = y \vee y < x).$$

Which of the following is dense: the natural numbers, the integers, the rational numbers, the real numbers, in each case with their usual order? Which have no endpoints?

12.25 Continuing the preceding problem, show that the set of sentences whose models are the dense linear orders without endpoints is denumerably categorical.

12.26 A linear order is said to have endpoints if it is a model of

$$\exists x \forall y (x < y \vee x = y) \;\&\; \exists x \forall y (x = y \vee y < x).$$

Show that the set of sentences whose models are the dense linear orders with endpoints is denumerably categorical.

12.27 How many isomorphism types of denumerable dense linear orders are there?

13

The Existence of Models

This chapter is entirely devoted to the proof of the compactness theorem. Section 13.1 outlines the proof, which reduces to establishing two main lemmas. These are then taken up in sections 13.2 through 13.4 to complete the proof, from which the Löwenheim–Skolem theorem also emerges as a corollary. The optional section 13.5 discusses what happens if nonenumerable languages are admitted: compactness still holds, but the Löwenheim–Skolem theorem in its usual 'downward' form fails, while an alternative 'upward' theorem holds.

13.1 Outline of the Proof

Our goal is to prove the compactness theorem, which has already been stated in the preceding chapter (in section 12.3). For convenience, we work with a version of first-order logic in which the only logical operators are $\sim$, $\vee$, and $\exists$, that is, in which & and $\forall$ are treated as unofficial abbreviations. The hypothesis of the theorem, it will be recalled, is that every finite subset of a given set of sentences is satisfiable, and the conclusion we want to prove is that the set itself is satisfiable, or, as we more elaborately put it, belongs to the set S of all satisfiable sets of sentences. As a first step towards the proof, we set down some properties enjoyed by this target set S. The reason for not including & and $\forall$ officially in the language is simply that in this and subsequent lemmas we would need four more clauses, two for & and two for $\forall$. These would not be difficult to prove, but they would be tedious.

13.1 Lemma (Satisfaction properties lemma). Let S be the set of all sets Γ of sentences of a given language such that Γ is satisfiable. Then S has the following properties:

(S0) If Γ is in S and Γ_0 is a subset of Γ, then Γ_0 is in S.
(S1) If Γ is in S, then for no sentence A are both A and $\sim A$ in Γ.
(S2) If Γ is in S and $\sim\sim B$ is in Γ, then $\Gamma \cup \{B\}$ is in S.
(S3) If Γ is in S and $(B \vee C)$ is in Γ, then either $\Gamma \cup \{B\}$ is in S or $\Gamma \cup \{C\}$ is in S.
(S4) If Γ is in S and $\sim(B \vee C)$ is in Γ, then $\Gamma \cup \{\sim B\}$ is in S and $\Gamma \cup \{\sim C\}$ is in S.
(S5) If Γ is in S and $\{\exists x\, B(x)\}$ is in Γ, and the constant c does not occur in Γ or $\exists x B(x)$, then $\Gamma \cup \{B(c)\}$ is in S.
(S6) If Γ is in S and $\sim\exists x\, B(x)$ is in Γ, then for every closed term t, $\Gamma \cup \{\sim B(t)\}$ is in S.

(S7) If Γ is in S, then $\Gamma \cup \{t = t\}$ is in S for any closed term t of the language of Γ.
(S8) If Γ is in S and $B(s)$ and $s = t$ are in Γ, then $\Gamma \cup \{B(t)\}$ is in S.

Proof: These have been established in Chapter 10. (S0) and (S1) were mentioned just before Example 10.4. (S2) appeared as Example 10.4(g), where it was derived from Example10.3(a). (S4), (S6), and (S8) can be derived in exactly the same way from Example 10.3(c), 10.3(e), and 10.3(f), as remarked after the proof of Example 10.4. (S3), (S5), and (S7) were established in Example 10.5.

We call (S0)–(S8) the *satisfaction properties*. Of course, at the outset we do not know that the set we are interested in belongs to S. Rather, what we are given is that it belongs to the set S^* of all sets of sentences whose every finite subset belongs to S. (Of course, once we succeed in proving the compactness theorem, S and S^* will turn out to be the *same* set.) It will be useful to note that S^* shares the above properties of S.

13.2 Lemma (Finite character lemma). If S is a set of sets of sentences having the satisfaction properties, then the set S^* of all sets of formulas whose every finite subset is in S also has properties (S0)–(S8).

Proof: To prove (S0) for S^*, note that if every finite subset of Γ is in S, and Γ_0 is subset of Γ, then every finite subset of Γ_0 is in S, since any finite subset of Γ_0 is a finite subset of Γ. To prove (S1) for S^*, note that if every finite subset of Γ is in S, then Γ cannot contain both A and $\sim A$, else $\{A, \sim A\}$ would be a finite subset of Γ, though $\{A, \sim A\}$ is not in S by property (S1) of S. To prove (S2) for S^*, note that if every finite subset of $\Gamma \cup \{\sim\sim B\}$ is in S, then any finite subset of $\Gamma \cup \{B\}$ is either a finite subset of Γ and hence of $\Gamma \cup \{\sim\sim B\}$ and therefore is in S, or else is of form $\Gamma_0 \cup \{B\}$ where Γ_0 is a finite subset of Γ. In the latter case, $\Gamma_0 \cup \{\sim\sim B\}$ is a finite subset of $\Gamma \cup \{\sim\sim B\}$ and therefore in S, so $\Gamma \cup \{B\}$ is in S by property (S2) of S. Thus the finite subset $\Gamma_0 \cup \{B\}$ is in S^*. (S4)–(S8) for S^* follow from (S4)–(S8) for S *exactly* as in the case of (S2). It remains only to prove (S3) for S^*.

So suppose every finite subset of $\Gamma \cup \{(B \vee C)\}$ is in S, but that it is not the case that every finite subset of $\Gamma \cup \{B\}$ is in S, or in other words that there is some finite subset of $\Gamma \cup \{B\}$ that is not in S. This cannot just be a subset of Γ, since then it would be a finite subset of $\Gamma \cup \{(B \vee C)\}$ and would be in S. So it must be of the form $\Gamma_0 \cup \{B\}$ for some finite subset Γ_0 of Γ. We now claim that every finite subset of $\Gamma \cup \{C\}$ is in S. For any such set is either a finite subset of Γ and therefore in S, or is of form $\Gamma_1 \cup \{C\}$ for some finite subset Γ_1 of Γ. In the latter case, $\Gamma_0 \cup \Gamma_1 \cup \{(B \vee C)\}$ is a finite subset of $\Gamma \cup \{(B \vee C)\}$ and so is in S. It follows that either $\Gamma_0 \cup \Gamma_1 \cup \{B\}$ or $\Gamma_0 \cup \Gamma_1 \cup \{C\}$ is in S by property (S3) of S. But if $\Gamma_0 \cup \Gamma_1 \cup \{B\}$ were in S, then by property (S1) of S, $\Gamma_0 \cup \{B\}$ would be in S, which it is not. So it must be that $\Gamma_0 \cup \Gamma_1 \cup \{C\}$ is in S and hence $\Gamma_1 \cup \{C\}$ is in S by property (S0) of S.

By these preliminary manoeuvres, we have reduced proving the compactness theorem to proving the following lemma, which is a kind of converse to Lemma 13.1. In stating it we suppose we have available an infinite set of constants not occurring in the set of sentences we are interested in.

13.3 Lemma (Model existence lemma). Let L be a language, and L^+ a language obtained by adding infinitely many new constants to L. If S^* is a set of sets of sentences of L^+ having the satisfaction properties, then every set of sentences of L in S^* has a model in which each element of the domain is the denotation of some closed term of L^+.

Note that the condition that every element of the domain is the denotation of some closed term guarantees that, since we are working in an enumerable language, the domain will be enumerable, which means that we get not only the compactness but also the Löwenheim–Skolem theorem, as remarked in the preceding chapter (following the statement of the two theorems in section 12.3).

So it 'only' remains to prove Lemma 13.3. The conclusion of Lemma 13.3 asserts the existence of an interpretation in which every element of the domain is the denotation of some closed term of the relevant language, and we begin by listing some properties that the set of all sentences true in such an interpretation would have to have.

13.4 Proposition (Closure properties lemma). Let L^+ be a language and $\mathcal{M}$ an interpretation thereof in which every element of the domain is the denotation of some closed term. Then the set Γ^* of sentences true in $\mathcal{M}$ has the following properties:

(C1) For no sentence A are both A and $\sim A$ in Γ^*.
(C2) If $\sim\sim B$ is in Γ^*, then B is in Γ^*.
(C3) If $B \vee C$ is in Γ^*, then either B is in Γ^* or C is in Γ^*.
(C4) If $\sim(B \vee C)$ is in Γ^*, then both $\sim B$ and $\sim C$ are in Γ^*.
(C5) If $\exists x B(x)$ is in Γ^*, then for some closed term t of L^+, $B(t)$ is in Γ^*.
(C6) If $\sim\exists x B(x)$ is in Γ^*, then for every closed term t of L^+, $\sim B(t)$ is in Γ^*.
(C7) For every closed term t of L^+, $t = t$ is in Γ^*.
(C8) If $B(s)$ and $s = t$ are in Γ^*, then $B(t)$ is in Γ^*.

Proof: For (C1), for no A are both A and $\sim A$ true in the same interpretation. For (C2), anything implied by anything true in a given interpretation is itself true in that interpretation, and B is implied by $\sim\sim B$. Similarly for (C4) and (C6)–(C8).

For (C3), any interpretation that makes a disjunct true must make at least one of its disjuncts true.

For (C5), if $\exists x B(x)$ is true in a given interpretation, then $B(x)$ is satisfied by some element m of the domain, and if that element m is the denotation of some closed term t, then $B(t)$ is true.

We call the properties (C1)–(C8) the *closure properties*. Actually, it is not Proposition 13.4 itself that will be useful to us here, but the following converse.

13.5 Lemma (Term models lemma). Let Γ^* be a set of sentences with the closure properties. Then there is an interpretation $\mathcal{M}$ in which every element of the domain is the denotation of some closed term, such that every sentence in Γ^* is true in $\mathcal{M}$.

To prove Lemma 13.3, it would suffice to prove the foregoing lemma plus the following one.

13.6 Lemma (Closure lemma). Let L be a language, and L^+ a language obtained by adding infinitely many new constants to L. If S^* is a set of sets of sentences of L^+ having the satisfaction properties, then every set Γ of sentences of L in S^* can be extended to a set Γ^* of sentences of L^+ having the closure properties.

Sections 13.2 and 13.3 will be devoted to the proof of the term models lemma, Lemma 13.5. As in so many other proofs, we consider first, in section 13.2, the case where identity and function symbols are absent, so that (C7) and (C8) may be ignored, and the only closed terms are constants, and then, in section 13.3, consider the additional complications that arise when identity is present, as well as those created by the presence of function symbols. The proof of the closure lemma, Lemma 13.6, will be given in section 13.4, with an alternative proof, avoiding any dependence on the assumption that the language is enumerable, to be outlined in the optional section 13.5.

13.2 The First Stage of the Proof

In this section we are going to prove the term models lemma, Lemma 13.5, in the case where identity and function symbols are absent. So let there be given a set Γ^* with the closure properties (C1)–(C6), as in the hypothesis of the lemma to be proved. We want to show that, as in the conclusion of that lemma, there is an interpretation in which every element of the domain is the denotation of some constant of the language of Γ^*, in which every sentence in Γ^* will be true.

To specify an interpretation $\mathcal{M}$ in this case, we need to do a number of things. To begin with, we must specify the domain $|\mathcal{M}|$. Also, we must specify for each constant c of the language which element $c^{\mathcal{M}}$ of the domain is to serve as its denotation. Moreover, we must do all this in such a way that *every* element of the domain is the denotation of *some* constant. This much is easily accomplished: simply pick for each constant c some object $c^{\mathcal{M}}$, picking a distinct object for each distinct constant, and let the domain consist of these objects.

To complete the specification of the interpretation, we must specify for each predicate R of the language what relation $R^{\mathcal{M}}$ on elements of the domain is to serve as its denotation. Moreover, we must do so in such a way that it will turn out that for every sentence B in the language we have

(1) $$\text{if} \quad B \text{ is in } \Gamma^* \quad \text{then} \quad \mathcal{M} \models B.$$

What we do is to specify $R^{\mathcal{M}}$ in such a way that (1) *automatically* becomes true for *atomic* B. We define $R^{\mathcal{M}}$ by the following condition:

$$R^{\mathcal{M}}\big(c_1^{\mathcal{M}}, \ldots, c_n^{\mathcal{M}}\big) \quad \text{if and only if} \quad R(c_1, \ldots, c_n) \text{ is in } \Gamma^*.$$

Now the definition of truth for atomic sentences reads as follows:

$$\mathcal{M} \models R(c_1, \ldots, c_n) \quad \text{if and only if} \quad R^{\mathcal{M}}\big(c_1^{\mathcal{M}}, \ldots, c_n^{\mathcal{M}}\big).$$

We therefore have the following:

(2) $$\mathcal{M} \models R(c_1, \ldots, c_n) \quad \text{if and only if} \quad R(c_1, \ldots, c_n) \text{ is in } \Gamma^*$$

and this implies (1) for atomic B.

We also have (1) for negated atomic sentences. For if $\sim R(c_1, \ldots, c_n)$ is in Γ^*, then by property (C1) of Γ^*, $R(c_1, \ldots, c_n)$ is *not* in Γ^*, and therefore by (2), $R(c_1, \ldots, c_n)$ is *not* true in $\mathcal{M}$, and so $\sim R(c_1, \ldots, c_n)$ *is* true in $\mathcal{M}$, as required.

To prove (1) for other formulas, we proceed by induction on complexity. There are three cases, according as A is a negation, a disjunction, or an existential quantification. However, we divide the negation case into subcases. Apart from the subcase of the negation of an atomic sentence, which we have already handled, there are three of these: the negation of a negation, the negation of a disjunction, and the negation of an existential quantification. So there are five cases in all:

to prove (1) for $\sim\sim B$	assuming (1) for B
to prove (1) for $B_1 \vee B_2$	assuming (1) for each B_i
to prove (1) for $\sim(B_1 \vee B_2)$	assuming (1) for each $\sim B_i$
to prove (1) for $\exists x B(x)$	assuming (1) for each $B(c)$
to prove (1) for $\sim\exists x B(x)$	assuming (1) for each $\sim B(c)$.

The five cases correspond to the five properties (C2)–(C6), which are just what is needed to prove them.

If $\sim\sim B$ is in Γ^*, then B is in Γ^* by property (C2). Assuming that (1) holds for B, it follows that B is true in $\mathcal{M}$. But then $\sim B$ is untrue, and $\sim\sim B$ is true as required. If $B_1 \vee B_2$ is in Γ^*, then B_i is in Γ^* for at least one of $i = 1$ or 2 by property (C3) of Γ^*. Assuming (1) holds for this B_i, it follows that B_i is true in $\mathcal{M}$. But then $B_1 \vee B_2$ is true as required. If $\sim(B_1 \vee B_2)$ is in Γ^*, then each $\sim B_i$ is in Γ^* for $i = 1$ or 2 by property (C4) of Γ^*. Assuming (1) holds for the $\sim B_i$, it follows that each $\sim B_i$ is true in $\mathcal{M}$. But then each B_i is untrue, so $B_1 \vee B_2$ is untrue, so $\sim(B_1 \vee B_2)$ is true as required.

In connection with existential quantification, note that since every individual in the domain is the denotation of some constant, $\exists x B(x)$ will be true if and only if $B(c)$ is true for some constant c. If $\exists x B(x)$ is in Γ^*, then $B(c)$ is in Γ^* for some constant c by property (C5) of Γ^*. Assuming (1) holds for this $B(c)$, it follows that $B(c)$ is true in $\mathcal{M}$. But then $\exists x B(x)$ is true as required. If $\sim\exists x B(x)$ is in Γ^*, then $\sim B(c)$ is in Γ^* for every constant c by property (C6) of Γ^*. Assuming (1) holds for each $\sim B(c)$, it follows that $\sim B(c)$ is true in $\mathcal{M}$. But then $B(c)$ is untrue for each c, and so $\exists x B(x)$ is untrue, and $\sim\exists x B(x)$ is true as required. We are done with the case without identity or function symbols.

13.3 The Second Stage of the Proof

In this section we want to extend the result of the preceding section to the case where identity is present, and then to the case where function symbols are also present. Before describing the modifications of the construction of the preceding section needed to accomplish this, we pause for a lemma.

13.7 Lemma. Let Γ^* be a set of sentences with properties (C1)–(C8). For closed terms t and s write $t \equiv s$ to mean that the sentence $t = s$ is in Γ^*. Then the following hold:

(E1) $t \equiv t$.
(E2) If $s \equiv t$, then $t \equiv s$.

(E3) If $t \equiv s$ and $s \equiv r$, then $t \equiv r$.

(E4) If $t_1 \equiv s_1, \ldots, t_n \equiv s_n$, then for any predicate R, $R(t_1, \ldots, t_n)$ is in Γ^* if and only if $R(s_1, \ldots, s_n)$ is in Γ^*.

(E5) If $t_1 \equiv s_1, \ldots, t_n \equiv s_n$, then for any function symbol f, $f(t_1, \ldots, t_n) = f(s_1, \ldots, s_n)$ is in Γ^*.

Proof: (E1) is simply a restatement of (C7). For (E2), let $B(x)$ be the formula $x = s$. We now know that the sentence $B(s)$, which is to say the sentence $s = s$, is in Γ^*, so if $s = t$ is in Γ^*, it follows by (C8) that the sentence $B(t)$, which is to say the sentence $t = s$, is in Γ^*. For (E3), let $B(x)$ be the formula $x = r$. If $t = s$ is in Γ^*, then we now know $s = t$ is in Γ^*, and if $B(s)$, which is $s = r$, is in Γ^*, it follows by (C8) that $B(t)$, which is $t = r$, is in Γ^*. For (E4), if all $t_i = s_i$ are in Γ^* and $R(s_1, \ldots, s_n)$ is in Γ^*, then repeated application of (C8) tells us that $R(t_1, s_2, s_3, \ldots, s_n)$ is in Γ^*, that $R(t_1, t_2, s_3, \ldots, s_n)$ is in Γ^*, and so on, and finally that $R(t_1, \ldots, t_n)$ is in Γ^*. This gives the 'only if' direction of (E4). For the 'if' direction, if all $t_i = s_i$ are in Γ^*, then so are all $s_i = t_i$, so if $R(t_1, \ldots, t_n)$ is in Γ^*, then by the direction we have already proved, $R(s_1, \ldots, s_n)$ is in Γ^*. For (E5), the proof just given for (E4) applies not only to atomic formulas $R(x_1, \ldots, x_n)$ but to arbitrary formulas $F(x_1, \ldots, x_n)$. Applying this fact where F is the formula $f(t_1, \ldots, t_n) = f(x_1, \ldots, x_n)$ gives (E5).

Note that (E1)–(E3) say that $\equiv$ is an equivalence relation. If we write $[t]$ for the equivalence class of t, then (E4) and (E5) may be rewritten as follows:

(E4′) If $[t_1] = [s_1], \ldots, [t_n] = [s_n]$, then for any predicate R, $R(t_1, \ldots, t_n)$ is in Γ^* if and only if $R(s_1, \ldots, s_n)$ is in Γ^*

(E5′) If $[t_1] = [s_1], \ldots, [t_n] = [s_n]$, then for any function symbol f, $[f(t_1, \ldots, t_n)] = [f(s_1, \ldots, s_n)]$.

We now return to the proof of the term models lemma, taking up the case where identity is present but function symbols are absent, so the only closed terms are constants. To specify the domain for our interpretation, instead of picking a distinct object for each distinct constant, we pick a distinct object C^* for each distinct *equivalence class* C of constants. We let the domain of the interpretation consist of these objects, and for the denotations of constants we specify the following:

$$c^{\mathcal{M}} = [c]^*. \tag{3}$$

Since $[c] = [d]$ if and only if $c = d$ is in Γ^*, we then have:

$$c^{\mathcal{M}} = d^{\mathcal{M}} \quad \text{if and only if} \quad c = d \text{ is in } \Gamma^*.$$

This is (the analogue of) (2) of the preceding section for atomic sentences involving the logical predicate =, and gives us (1) of the preceding section for such sentences and their negations.

What remains to be done is to define the denotation $R^{\mathcal{M}}$ for a nonlogical predicate R, in such a way that (2) of the preceding section will hold for atomic sentences involving nonlogical predicates. *From that point, the rest of the proof will be exactly the same* as where identity was not present. Towards framing the definition of $R^{\mathcal{M}}$, note that (E4′) allows us to give the following definition:

$$R^{\mathcal{M}}(C_1{*}, \ldots, C_n{*}) \quad \text{if and only if} \quad \begin{array}{l} R(c_1, \ldots, c_n) \text{ is in } \Gamma{*} \\ \text{for } \textit{some} \text{ or equivalently } \textit{any} \\ c_i \text{ with } C_i = [c_i]. \end{array}$$

Thus

$$R^{\mathcal{M}}([c_1], \ldots, [c_n]) \quad \text{if and only if} \quad R(c_1, \ldots, c_n) \text{ is in } \Gamma{*}.$$

Together with (3), this gives (2) of the preceding section. Since as already indicated the proof is the same from this point on, we are done with the case with identity but without function symbols.

For the case with function symbols, we pick a distinct object $T{*}$ for each equivalence class of *closed terms*, and let the domain of the interpretation consist of these objects. Note that (3) above still holds for constants. We must now specify for each function symbol f what function $f^{\mathcal{M}}$ on this domain is to serve as its denotation, and in such a way that (3) will hold for all closed terms. From that point, the rest of the proof will be exactly the same as in the preceding case where function symbols were not present.

(E5′) allows us to give the following definition:

$$f^{\mathcal{M}}(T_1{*}, \ldots, T_n{*}) = T{*} \qquad \text{where} \quad \begin{array}{l} T = [f(t_1, \ldots, t_n)] \\ \text{for } \textit{some} \text{ or equivalently } \textit{any} \\ t_i \text{ with } T_i = [t_i]. \end{array}$$

Thus

(4) $$f^{\mathcal{M}}([t_1]{*}, \ldots, [t_n]{*}) = [f(t_1, \ldots, t_n)]{*}.$$

We can now prove by induction on complexity that (3) above, which holds by definition for constants, in fact holds for any closed term t. For suppose (3) holds for $t_1, \ldots, t_n$, and consider $f(t_1, \ldots, t_n)$. By the general definition of the denotation of a term we have

$$(f(t_1, \ldots, t_n))^{\mathcal{M}} = f^{\mathcal{M}}\big(t_1^{\mathcal{M}}, \ldots, t_n^{\mathcal{M}}\big).$$

By our induction hypothesis about the t_i we have

$$t_i^{\mathcal{M}} = [t_i]{*}.$$

Putting these together, we get

$$(f(t_1, \ldots, t_n))^{\mathcal{M}} = f^{\mathcal{M}}([t_1]{*}, \ldots, [t_n]{*}).$$

And this together with the definition (4) above gives

$$(f(t_1, \ldots, t_n))^{\mathcal{M}} = [f(t_1, \ldots, t_n)]{*}.$$

which is precisely (3) above for the closed term $f(t_1, \ldots, t_n)$. Since, as already indicated, the proof is the same from this point on, we are done.

13.4 The Third Stage of the Proof

What remains to be proved is the closure lemma, Lemma 13.6. So let there be given a language L, a language L^+ obtained by adding infinitely many new constants to L, a set S^* of sets of sentences of L^+ having the satisfaction properties (S0)–(S8), and a set Γ of sentences of L in S^*, as in the hypotheses of the lemma to be proved. We want to show that, as in the conclusion of that lemma, Γ can be extended to a set Γ^* of sentences of L^+ with closure properties (C1)–(C8).

The idea of the proof will be to obtain Γ^* as the union of a sequence of sets $\Gamma_0, \Gamma_1, \Gamma_2, \ldots$, where each Γ_n belongs to S^* and each contains all earlier sets Γ_m for $m < n$, and where Γ_0 is just Γ. (C1) will easily follow, because if A and $\sim A$ were both in Γ^*, A would be in some Γ_m and $\sim A$ would be in some Γ_n, and then both would be in Γ_k, where k is whichever of m and n is the larger. But since Γ_k is in S^*, this is impossible, since (S0) says precisely that no element of S^* contains both A and $\sim A$ for any A.

What need to be worried about are (C2)–(C8). We have said that each Γ_{k+1} will be a set in S^* containing Γ_k. In fact, each Γ_{k+1} be obtained by adding to Γ_k a *single* sentence B_k, so that $\Gamma_{k+1} = \Gamma_k \cup \{B_k\}$. (It follows that each Γ_n will be obtained by adding only finitely many sentences to Γ, and therefore will involve only finitely many of the constants of L^+ that are not in the language L of Γ, leaving at each stage infinitely many as yet unused constants.) At each stage, having Γ_k in S^*, we are free to choose as B_k any sentence such that $\Gamma_k \cup \{B_k\}$ is still in S^*. But we must make the choices in such a way that in the end (C2)–(C8) hold.

Now how can we arrange that Γ^* fulfills condition (C2), for example? Well, if $\sim\sim B$ is in Γ^*, it is in some Γ_m. *If* we can so arrange matters that whenever m and B are such that $\sim\sim B$ is in Γ_m, then B is in Γ_{k+1} for some $k \geq m$, *then* it will follow that B is in Γ^*, as required by (C2). To achieve this, it will be more than enough if we can so arrange matters that the following holds:

$$\text{If} \quad \sim\sim B \text{ is in } \Gamma_m, \quad \text{then} \quad \text{for some } k \geq m, \quad \Gamma_{k+1} = \Gamma_k \cup \{B\}.$$

But *can* we so arrange matters that this holds? Well, what does (S2) tell us? If $\sim\sim B$ is in Γ_m, then $\sim\sim B$ will still be in Γ_k for any $k \geq m$, since the sets get larger. Since each Γ_k is to be in S^*, (S2) promises that $\Gamma_k \cup \{B\}$ will be in S^*. That is:

$$\text{If} \quad \sim\sim B \text{ is in } \Gamma_m, \quad \text{then} \quad \text{for any } k \geq m, \quad \Gamma_k \cup \{B\} \text{ is in } S^*.$$

So we *could* take $\Gamma_{k+1} = \Gamma_k \cup \{B\}$ if we chose to do so.

To understand better what is going on here, let us introduce some suggestive terminology. If $\sim\sim B$ is in Γ_m, let us say that *the demand for admission of B is raised* at stage m; and if $\Gamma_{k+1} = \Gamma_k \cup \{B\}$, let us say that *the demand is granted* at stage k. What is required by (C2) is that *any* demand that is raised at *any* stage m should be granted at some later stage k. And what is promised by (S2) is that at any stage k, any *one* demand raised at any *one* earlier stage m could be granted. There is a gap here between what is demanded and what is promised, since it may well be that there are infinitely many demands raised at stage m, which is to say, infinitely many sentences of form $\sim\sim B$ in Γ_m, and in any case, there are infinitely many stages m at which new demands may arise—and all this only considering demands of the type associated with condition (C2), whereas there are several other conditions, also raising demands, that we also wish to fulfill.

Let us look at these. The relationship between (C3)–(C8) and (S3)–(S8) is exactly the same as between (C2) and (S2). Each of (C2)–(C8) corresponds to a demand of a certain type:

(C2) If $\sim\sim B$ is in Γ_m, then for some $k \geq m$, $\Gamma_{k+1} = \Gamma_k \cup \{B\}$.
(C3) If $B \vee C$ is in Γ_m, then for some $k \geq m$, $\Gamma_{k+1} = \Gamma_k \cup \{B\}$ or $\Gamma_k \cup \{C\}$.
(C4) If $\sim(B \vee C)$ or $\sim(C \vee B)$ is in Γ_m, then for some $k \geq m$, $\Gamma_{k+1} = \Gamma_k \cup \{\sim B\}$.
(C5) If $\exists x B(x)$ is in Γ_m, then for some $k \geq m$, for some constant c, $\Gamma_{k+1} = \Gamma_k \cup \{B(c)\}$.
(C6) If $\sim\exists x B(x)$ is in Γ_m and t is a closed term in the language of Γ_m, then for some $k \geq m$, $\Gamma_{k+1} = \Gamma_k \cup \{\sim B(t)\}$.
(C7) If t is a closed term in the language of Γ_m, then for some $k \geq m$, $\Gamma_{k+1} = \Gamma_k \cup \{t = t\}$.
(C8) If $B(s)$ and $s = t$ are in Γ_m, where s and t are closed terms $B(x)$ a formula, then for some $k \geq m$, $\Gamma_{k+1} = \Gamma_k \cup \{B(t)\}$.

Each of (S2)–(S8) promises that any one demand of the relevant type can be granted:

(S2) If $\sim\sim B$ is in Γ_m, then for any $k \geq m$, $\Gamma_k \cup \{B\}$ is in S^*.
(S3) If $B \vee C$ is in Γ_m, then for any $k \geq m$, $\Gamma_k \cup \{B\}$ or $\Gamma_k \cup \{C\}$ is in S^*.
(S4) If $\sim(B \vee C)$ or $\sim(C \vee B)$ is in Γ_m, then for any $k \geq m$, $\Gamma_k \cup \{\sim B\}$ is in S^*.
(S5) If $\exists x B(x)$ is in Γ_m, then for any $k \geq m$, for any as yet unused constant c, $\Gamma_k \cup \{B(c)\}$ is in S^*.
(S6) If $\sim\exists x B(x)$ is in Γ_m and t is a closed term in the language of Γ_m, then for any $k \geq m$, $\Gamma_k \cup \{\sim B(t)\}$ is in S^*.
(S7) If t is a closed term in the language of Γ_m, then for any $k \geq m$, $\Gamma_k \cup \{t = t\}$ is in S^*.
(S8) If $B(s)$ and $s = t$ are in Γ_m, where s and t are closed terms $B(x)$ a formula, then for any $k \geq m$, $\Gamma_k \cup \{B(t)\}$ is in S^*.

At any stage k of the construction, we can grant *any one demand we choose* from among those that have been raised at earlier stages, but for the construction to succeed, we must make our successive choices so that in the end *any demand that is ever raised at any stage* is granted at some later stage. Our difficulty is that at each stage many different demands may be raised. Our situation is like that of Herakles fighting the hydra: every time we chop off one head (grant one demand), multiple new heads appear (multiple new demands are raised). At least in one respect, however, we have made progress: we have succeeded in redescribing our problem in abstract terms, eliminating all details about which particular formulas are of concern.

And indeed, with this redescription of the problem we are now not far from a solution. We need only recall two facts. First, our languages are *enumerable*, so that at each stage, though an infinity of demands may be raised, it is still only an *enumerable* infinity. Each demand may be worded 'admit such-and-such a sentence' (or 'admit one or the other of two such-and-such sentences'), and an enumeration of the sentences of our language therefore gives rise to an enumeration of all the demands raised at any given stage. Thus each demand that is ever raised may be described as the ith demand raised at stage m, for some numbers i and m, and so may be described by a pair of numbers (i, m). Second, we have seen in Chapter 1 that there is a way—in fact, there are many ways—of coding any pair of numbers by a single number $j(i, m)$, and if one looks closely at this coding, one easily sees that $j(i, m)$ is greater than m (and greater than i). We can solve our problem, then, by proceeding as follows. At stage k, see what pair (i, m) is coded by k, and grant the ith demand that was raised at stage $m < k$. In this way, though we grant only one demand at a time, all the demands that are ever raised will eventually be granted.

The proof of the compactness theorem is now complete.

In the remaining problems, for simplicity assume that function symbols are absent, though the results indicated extend to the case where they are present.

13.8 An *embedding* of one interpretation $\mathcal{P}$ in another interpretation $\mathcal{Q}$ is a function j fulfilling all the conditions in the definition of isomorphism in section 13.1, except that j need not be onto. Given an interpretation $\mathcal{P}$, let $L^{\mathcal{P}}$ be the result of adding to the language a constant c_p for each element p of the domain $|\mathcal{P}|$, and and let $\mathcal{P}*$ be the extension of $\mathcal{P}$ to an interpretation of $L^{\mathcal{P}}$ in which each c_p denotes the corresponding p. The set $\Delta(\mathcal{P})$ of all atomic and negated atomic sentences of $L^{\mathcal{P}}$, whether involving a nonlogical predicate R or the logical predicate $=$, that are true in $\mathcal{P}*$, is called the *diagram* of $\mathcal{P}$. Show that if $\mathcal{Q}$ is any interpretation of the language of $\mathcal{P}$ that can be extended to a model $\mathcal{Q}*$ of $\Delta(\mathcal{P})$, then there is an embedding of $\mathcal{P}$ into $\mathcal{Q}$.

13.9 A sentence is called *existential* if and only if it is of the form $\exists x_1 \ldots \exists x_n F$ where F contains no further quantifiers (universal or existential). A sentence is said to be *preserved upwards* if and only if, whenever it is true in an interpretation $\mathcal{P}$, and there is an embedding of $\mathcal{P}$ in another interpretation $\mathcal{Q}$, then it is true in $\mathcal{Q}$. Show that every existential sentence is preserved upwards.

13.10 Let A be a sentence that is preserved upwards, $\mathcal{P}$ a model of A, and $\Delta(\mathcal{P})$ the diagram of $\mathcal{P}$. Show that $\Delta \cup \{\sim A\}$ is unsatisfiable, and that some finite subset of $\Delta \cup \{\sim A\}$ is unsatisfiable.

13.11 Let A be a sentence of a language L that is preserved upwards. Show that:

(a) $\mathcal{P}$ is a model of A if and only if there is a quantifier-free sentence B of the language $L^{\mathcal{P}}$ such that B implies A and $\mathcal{P}^*$ is a model of B.

(b) $\mathcal{P}$ is a model of A if and only if there is an existential sentence B of the language L such that B implies A and $\mathcal{P}$ is a model of B.

13.12 Let A be a sentence that is preserved upwards, and Γ the set of existential sentences of the language of A that imply A. Writing $\sim\Gamma$ for the set of negations of elements of Γ, show that:

(a) $\{A\} \cup \sim\Gamma$ is unsatisfiable.

(b) $\{A\} \cup \sim\Gamma_0$ is unsatisfiable for some finite subset Γ_0 of Γ.

(c) $\{A\} \cup \{\sim B\}$ is unsatisfiable for some single element of Γ.

13.13 Let A be a sentence that is preserved upwards. Show that A is logically equivalent to an existential sentence (in the same language).

13.14 A sentence is called *universal* if and only if it is of the form $\forall x_1 \ldots \forall x_n F$ where F contains no further quantifiers (universal or existential). A sentence is said to be *preserved downwards* if and only if, whenever it is true in an interpretation $\mathcal{Q}$, and there is an embedding of $\mathcal{P}$ in another interpretation $\mathcal{Q}$, then it is true in $\mathcal{P}$. Prove that a sentence is preserved downwards if and only if it is logically equivalent to a universal sentence (in the same language).

13.15 The proof in the preceding several problems involves (at the step of Problem 13.10) applying the compactness theorem to a language that may be nonenumerable. How could this feature be avoided?