(i). For each n-placed relation symbol P in $\mathscr{L}$, we define an n-placed relation R' on the set C by:
for all $c_1, \ldots, c_n \in C$,
(2) $R'(c_1 \ldots c_n)$ iff $P(c_1 \ldots c_n) \in T$.
By our axioms of identity, we have

$$\vdash P(c_1 \ldots c_n) \wedge c_1 \equiv d_1 \wedge \ldots \wedge c_n \equiv d_n \rightarrow P(d_1 \ldots d_n).$$

So $\sim$ is what is called a *congruence relation* for the relation R' on C. It follows that we may define a relation R on A by
(3) $R(\tilde{c}_1 \ldots \tilde{c}_n)$ iff $P(c_1 \ldots c_n) \in T$.
By (2), the definition (3) is independent of the representatives of the equivalence classes $\tilde{c}_1, \ldots, \tilde{c}_n$. This relation R is the interpretation of the symbol P in $\mathfrak{A}$.

(ii). Now consider a constant symbol d of $\mathscr{L}$. From predicate logic, we have

$$\vdash (\exists v_0)(d \equiv v_0).$$

So $(\exists v_0)(d \equiv v_0) \in T$, and, because T has witnesses, there is a constant $c \in C$ such that

$$(d \equiv c) \in T.$$

The constant c may not be unique, but its equivalence class is unique because, using our axioms of identity,

$$\vdash (d \equiv c \wedge d \equiv c' \rightarrow c \equiv c').$$

The constant d is interpreted in the model $\mathfrak{A}$ by the (uniquely determined) element $\tilde{c}$ of A. In particular, if $d \in C$, then d is interpreted by its own equivalence class $\tilde{d}$ in $\mathfrak{A}$, because $(d \equiv d) \in T$.

(iii). We handle the function symbols in a similar way. Let F be any m-placed function symbol of $\mathscr{L}$, and let $c_1, \ldots, c_m \in C$. As before, we have

$$(\exists v_0)(F(c_1 \ldots c_m) \equiv v_0) \in T,$$

and because T has witnesses, there is a constant $c \in C$ such that

$$(F(c_1 \ldots c_m) \equiv c) \in T.$$

Once more, we have a slight difficulty because c may not be unique, and we use our axioms of identity to obtain:

$$\vdash (F(c_1 \ldots c_m) \equiv c \wedge c_1 \equiv d_1 \wedge \ldots \wedge c_m \equiv d_m \wedge c \equiv d) \rightarrow F(d_1 \ldots d_m) \equiv d.$$

This shows that a function G can be defined on the set A of equivalence classes by the rule

(4) $G(\tilde{c}_1 \ldots \tilde{c}_m) = \tilde{c}$ iff $(F(c_1 \ldots c_m) \equiv c) \in T$.
We leave the detailed steps of (4) to the reader. We interpret the function symbol F by the function G in the model $\mathfrak{A}$.

We have now specified the universe set and the interpretation of each symbol of $\mathscr{L}$ in $\mathfrak{A}$, so we have completed the definition of the model $\mathfrak{A}$. We have pointed out that the interpretation of each constant $c \in C$ in $\mathfrak{A}$ is its equivalence class $\tilde{c}$, and it follows that every element $\tilde{c} \in A$ is the interpretation of some constant $c \in C$.

We proceed to prove that $\mathfrak{A}$ is a model of T. First of all, using (4) as the first step of an induction, we easily show that

(5) for every term t of $\mathscr{L}$ with no free variables and for every constant $c \in C$,

$$\mathfrak{A} \models t \equiv c \quad \text{if and only if} \quad (t \equiv c) \in T.$$

Using the fact that C is a set of witnesses for T, we obtain from (5):
(6) for any two terms t_1, t_2 of $\mathscr{L}$ with no free variables,

$$\mathfrak{A} \models t_1 \equiv t_2 \quad \text{if and only if} \quad (t_1 \equiv t_2) \in T,$$

(7) for any atomic formula $P(t_1 \ldots t_n)$ of $\mathscr{L}$ containing no free variables,

$$\mathfrak{A} \models P(t_1 \ldots t_n) \quad \text{if and only if} \quad P(t_1 \ldots t_n) \in T.$$

Combining (6) and (7) will form a basis for proving:
(8) for any sentence φ of $\mathscr{L}$,

$$\mathfrak{A} \models \varphi \quad \text{if and only if} \quad \varphi \in T.$$

(8) has an unusual proof in that it is proved by induction on the length of the sentences of $\mathscr{L}$. The reader will see that the reason why this could be done is because T is maximal consistent and has witnesses in $\mathscr{L}$. Without a great deal of trouble, we have for sentences φ, ψ of $\mathscr{L}$

$$\mathfrak{A} \models \neg\varphi \quad \text{if and only if} \quad (\neg\varphi) \in T,$$

and

$$\mathfrak{A} \models \varphi \wedge \psi \quad \text{if and only if} \quad (\varphi \wedge \psi) \in T.$$

Suppose $\varphi = (\exists x)\psi$. If $\mathfrak{A} \models \varphi$, then for some $\tilde{c} \in A$, $\mathfrak{A} \models \psi[\tilde{c}]$. This means that $\mathfrak{A} \models \psi(c)$, where $\psi(c)$ is obtained from ψ by replacing all free occurrences of x by c. So $\psi(c) \in T$ and because

$$\vdash \psi(c) \rightarrow (\exists x)\psi,$$

we have $\varphi \in T$. On the other hand, if $\varphi \in T$, then because T has witnesses, there exists a constant $c \in C$ such that

$$T \vdash (\exists x)\psi \rightarrow \psi(c).$$

As T is maximal, $\psi(c) \in T$, so $\mathfrak{A} \models \psi(c)$. This gives $\mathfrak{A} \models \psi[\tilde{c}]$ and $\mathfrak{A} \models \varphi$. This shows that $\mathfrak{A}$ is a model of T. ⊣

Note that a converse of Lemma 2.1.2 is very easily proved, and, in fact:

LEMMA 2.1.3. *Let C be a set of constant symbols of $\mathscr{L}$, and let T be a set of sentences of $\mathscr{L}$. If T has a model $\mathfrak{A}$ such that every element of $\mathfrak{A}$ is an interpretation of some constant $c \in C$, then T can be extended to a consistent $\bar{T}$ in $\mathscr{L}$ for which C is a set of witnesses.*

For the proof of Lemma 2.1.3, simply let $\bar{T}$ be the set of all sentences of $\mathscr{L}$ true in $\mathfrak{A}$.

The model $\mathfrak{A}$ constructed from the constants $c \in C$ of $\mathscr{L}$ by taking suitable equivalence classes is said to be *built up from the set C of constants of $\mathscr{L}$*. Since every $a \in A$ is the interpretation of some $c \in C$, we see immediately that $|A| \leqslant |C|$. We now supply the proofs of three theorems from Chapter 1.

THEOREM 1.3.21 (Extended Completeness Theorem). *Let Σ be a set of sentences of $\mathscr{L}$. Then Σ is consistent if and only if Σ has a model.*

PROOF. The consistency of Σ if Σ has a model is a straightforward argument. So assume Σ is consistent. By Lemma 2.1.1, we consider extensions $\bar{\Sigma}$ of Σ and $\bar{\mathscr{L}}$ of $\mathscr{L}$ ($\|\bar{\mathscr{L}}\| = \|\mathscr{L}\|$), so that $\bar{\Sigma}$ has witnesses in $\bar{\mathscr{L}}$. By Lemma 2.1.2, let $\mathfrak{A}$ be a model of $\bar{\Sigma}$. $\mathfrak{A}$ is a model for the expanded language $\bar{\mathscr{L}}$, so let $\mathfrak{B}$ be the model for $\mathscr{L}$ which is the reduct of $\mathfrak{A}$ to $\mathscr{L}$. Because sentences in Σ do not involve the constants of $\bar{\mathscr{L}}$ not in $\mathscr{L}$, we see that $\mathfrak{B}$ is a model of Σ. ⊣

COROLLARY 2.1.4 (Downward Löwenheim–Skolem–Tarski Theorem). *Every consistent theory T in $\mathscr{L}$ has a model of power at most $\|\mathscr{L}\|$.*

PROOF. In the proof above we may choose $\mathfrak{A}$ so that every element is a constant, and we have $|B| = |A| \leqslant \|\bar{\mathscr{L}}\| = \|\mathscr{L}\|$. ⊣

Corollary 2.1.4 gives the original theorem of Löwenheim (1915): If a sentence has a model, then it has a countable (finite or infinite) model.

THEOREM 1.3.20 (Gödel's Completeness Theorem). *A sentence of $\mathscr{L}$ is a theorem of $\mathscr{L}$ if and only if it is valid.*

PROOF. We need only concern ourselves with one direction of the theorem. If a sentence σ is not a theorem of $\mathscr{L}$, then $\{\neg\sigma\}$ is consistent in $\mathscr{L}$. By Theorem 1.3.21, $\{\neg\sigma\}$ will have a model in which σ cannot hold. Hence σ is not valid. ⊣

THEOREM 1.3.22 (Compactness Theorem). *A set of sentences Σ has a model if and only if every finite subset of Σ has a model.*

PROOF. If every finite subset of Σ has a model, then every finite subset of Σ is consistent. So Σ is consistent and Theorem 1.3.21 shows that Σ has a model. ⊣

We conclude this section with a representative list of applications or consequences of the completeness and compactness theorems. Some additional exercises can be found at the end.

COROLLARY 2.1.5. *If a theory T has arbitrarily large finite models, then it has an infinite model.*

PROOF. Let T be a theory in $\mathscr{L}$ with arbitrarily large finite models. Consider the expansion $\mathscr{L}' = \mathscr{L} \cup \{c_n : n \in \omega\}$, where c_n is a list of distinct constant symbols not in $\mathscr{L}$. Consider the set Σ of $\mathscr{L}'$ defined by

$$\Sigma = T \cup \{\neg(c_n \equiv c_m) : n < m < \omega\}.$$

Any finite subset Σ' of Σ will involve at most the constants $c_0, \ldots, c_m$, say, for some m. Let $\mathfrak{A}$ be a model of T with at least $m+1$ elements, and let $a_0, \ldots, a_m$ be a list of $m+1$ distinct elements of $\mathfrak{A}$. We can verify easily that the model $(\mathfrak{A}, a_0, \ldots, a_m)$ for the finite expansion $\mathscr{L}'' = \mathscr{L} \cup \{c_0, \ldots, c_m\}$ of $\mathscr{L}$ is a model of Σ'. So, by Theorem 1.3.22, Σ has a model. The reduction of this model to $\mathscr{L}$ gives a model of T which is clearly infinite. ⊣

COROLLARY 2.1.6 (Upward Löwenheim–Skolem–Tarski Theorem). *If T has infinite models, then it has infinite models of any given power $\alpha \geqslant \|\mathscr{L}\|$.*

PROOF. The proof is similar to that of Corollary 2.1.5. Let c_ξ, $\xi < \alpha$, be a list of distinct constant symbols not in $\mathscr{L}$, and consider the set of sentences

$\Sigma = T \cup \{\neg(c_\xi \equiv c_\eta) : \xi < \eta < \alpha\}$. Every finite subset Σ' of Σ will involve at most a finite number of the constants c_ξ. Hence any infinite model of T can be expanded to a model of Σ'. By Theorem 1.3.22, Σ has a model $\mathfrak{A}$ and by Corollary 2.1.4, this model is of power at most

$$\|\mathscr{L} \cup \{c_\xi : \xi < \alpha\}\| = \alpha.$$

On the other hand, the interpretations of the constants c_ξ in $\mathfrak{A}$ must give distinct elements of A. So $\alpha \leqslant |A| \leqslant \alpha$ and $|A| = \alpha$. ⊣

A result first published by Skolem (1934) is the following:

Corollary 2.1.7. *There exist nonstandard models of complete number theory.*

Proof. Recall from 1.4.11 that complete number theory is the set of all sentences holding in the standard model $\langle \omega, +, \cdot, S, 0\rangle$ of number theory. Since this theory has an infinite model, it has models of all infinite powers. A noncountable model of complete number theory clearly cannot be standard. ⊣

A simple but powerful device in model theory is the *method of diagrams*. Let $\mathfrak{A}$ be a model for $\mathscr{L}$. We expand the language $\mathscr{L}$ to a new language

$$\mathscr{L}_A = \mathscr{L} \cup \{c_a : a \in A\}$$

by adding a new constant symbol c_a for each element $a \in A$. It is understood that if $a \neq b$, then c_a and c_b are different symbols. We may then expand $\mathfrak{A}$ to the model

$$\mathfrak{A}_A = (\mathfrak{A}, a)_{a \in A}$$

for $\mathscr{L}_A$ by interpreting each new constant c_a by the element a. The *diagram of* $\mathfrak{A}$, denoted by $\Delta_{\mathfrak{A}}$, is the set of all atomic sentences and negations of atomic sentences of $\mathscr{L}_A$ which hold in the model $\mathfrak{A}_A$.

If X is a subset of A, then we let $\mathscr{L}_X$ be the language $\mathscr{L} \cup \{c_a : a \in X\}$ and $\mathfrak{A}_X = (\mathfrak{A}, a)_{a \in X}$ be the obvious expansion of $\mathfrak{A}$ to $\mathscr{L}_X$. If f is a mapping from X into the set of elements B of a model $\mathfrak{B}$ for $\mathscr{L}$, then $(\mathfrak{B}, fa)_{a \in X}$ is the expansion of $\mathfrak{B}$ to a model for $\mathscr{L}_X$ formed by interpreting each c_a by fa.

The method of adding new constant symbols for elements of a model is used again and again in model theory. The following proposition illustrates the usefulness of diagrams.

Proposition 2.1.8. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be models for $\mathscr{L}$ and let $f : A \to B$. Then the following are equivalent*:

(a). *f is an isomorphic embedding of $\mathfrak{A}$ into $\mathfrak{B}$.*
(b). *There is an extension $\mathfrak{C} \supset \mathfrak{A}$ and an isomorphism $g : \mathfrak{C} \cong \mathfrak{B}$ such that $g \supset f$.*
(c). *$(\mathfrak{B}, fa)_{a \in A}$ is a model of the diagram of $\mathfrak{A}$.*

Proof. The implication from (b) to (a) is trivial. If (a) holds, one can extend the set A to a set C and extend the function f to a one to one function g from C onto B. Then define the relations of $\mathfrak{C}$ by the rule

$$\mathfrak{C} \models R[c_1 \ldots c_n] \text{ iff } \mathfrak{B} \models R[gc_1 \ldots gc_n],$$

and similarly for functions. This will make (b) hold for $\mathfrak{C}$ and g.

To prove the equivalence of (a) and (c), use the fact that by Proposition 1.3.18, for each formula $\varphi(x_1 \ldots x_n)$ and all $a_1, \ldots, a_n$ in A,

$$\mathfrak{A} \models \varphi[a_1 \ldots a_n] \text{ if and only if } \mathfrak{A}_A \models \varphi(a_1 \ldots a_n)$$

and

$$\mathfrak{B} \models \varphi[fa_1 \ldots fa_n] \text{ if and only if } (\mathfrak{B}, fa)_{a \in A} \models \varphi(a_1 \ldots a_n). \dashv$$

Proposition 2.1.8 shows that the following three conditions are equivalent:

(a′) *$\mathfrak{A}$ is isomorphically embeddable in $\mathfrak{B}$.*
(b′) *$\mathfrak{B}$ is isomorphic to an extension of $\mathfrak{A}$.*
(c′) *$\mathfrak{B}$ can be expanded to a model of the diagram of $\mathfrak{A}$.*

In the special case that $A \subset B$ and f is the identity mapping from A into B, Proposition 2.1.8 shows that $\mathfrak{A}$ is a submodel of $\mathfrak{B}$ if and only if $\mathfrak{B}_A$ is a model of the diagram of $\mathfrak{A}$.

Corollary 2.1.9. *Suppose that $\mathscr{L}$ has no function or constant symbols. Let T be a theory in $\mathscr{L}$ and $\mathfrak{A}$ be a model for $\mathscr{L}$. Then $\mathfrak{A}$ is isomorphically embedded in some model of T if and only if every finite submodel of $\mathfrak{A}$ is isomorphically embedded in some model of T.*

Proof. We skip the easy direction and suppose that every finite submodel of $\mathfrak{A}$ is isomorphically embedded in some model of T. We show that the set $\Sigma = T \cup \Delta_{\mathfrak{A}}$ is consistent. Every finite subset Σ' of Σ contains at most a finite number of the new constants, say $c_{a_1}, \ldots, c_{a_m}$. Because the language $\mathscr{L}$ has no function or constant symbols, the finite set $A' = \{a_1, \ldots, a_m\}$ generates a finite submodel $\mathfrak{A}'$ of $\mathfrak{A}$. Let $\mathfrak{B}'$ be a model of T in which $\mathfrak{A}'$ is isomorphically embedded. We see without difficulty that $\Sigma' \subset T \cup \Delta_{\mathfrak{A}'}$. So, by Proposition 2.1.8, $\mathfrak{B}'$ can be expanded to a model of Σ', and hence Σ' has a model. By compactness, Σ has a model $\mathfrak{B}$. By Proposition 2.1.8 again, the reduct of $\mathfrak{B}$ to $\mathscr{L}$ gives a model of T in which $\mathfrak{A}$ is isomorphically embedded. ⊣

We next consider two applications from the theory of fields (see 1.4.9).

COROLLARY 2.1.10. *Let T be a theory in the language $\mathscr{L} = \{+, \cdot, 0, 1\}$ which has as models fields with arbitrary high finite characteristics. Then T has a model which is a field of characteristic* 0.

PROOF. Let T' be the theory of fields and consider the set

$$\Sigma = T \cup T' \cup \{p1 \not\equiv 0 : \text{all primes } p\}.$$

Recall from Chapter 1 that $p1$ is our abbreviation for the term $1 + \dots + 1$, p times, of the language $\mathscr{L}$. A finite subset Σ' of Σ will involve a highest prime, say p. Let $\mathfrak{A}$ be a model of T which is a field, so $\mathfrak{A}$ is also a model of T', and such that the characteristic of $\mathfrak{A}$ is higher than p. Then $\mathfrak{A}$ is a model of Σ', whence by compactness, Σ has a model. This model is a model of T, is a field, and has characteristic 0. ⊣

COROLLARY 2.1.11. *There exist non-Archimedean ordered fields elementarily equivalent to the ordered field of real numbers.*

PROOF. An ordered field $\langle F, +, \cdot, 0, 1, \leqslant \rangle$ is Archimedean iff for any two positive elements a, b in F there is an n such that $na \geqslant b$. This is not expressible in first-order logic.

Let T be the set of all sentences of $\mathscr{L} = \{+, \cdot, 0, 1, \leqslant\}$ holding in the ordered field of reals. Let c be a constant symbol different from 0 and 1. Let

$$\Sigma = T \cup \{n1 \leqslant c : n \in \omega\}.$$

For every finite subset Σ' of Σ, there is an expansion of the reals to a model of Σ'. By compactness, Σ has a model in which c has an interpretation b. In this model, both 1 and b are positive; yet no finite multiple of 1 can exceed b. ⊣

Corollary 2.1.11 is the very beginning of a branch of model theory called *nonstandard analysis*. The model theory of nonstandard analysis will be developed in Section 4.4.

Consider $\Delta_{\mathfrak{A}}$, the diagram of $\mathfrak{A}$ introduced earlier. We see that Proposition 2.1.8 gives an intimate connection between models of $\Delta_{\mathfrak{A}}$ and models in which $\mathfrak{A}$ can be isomorphically embedded. By the *positive diagram of* $\mathfrak{A}$ we mean the subset of $\Delta_{\mathfrak{A}}$ which consists only of atomic sentences (no negations of atomic sentences). We shall see that positive diagrams are associated with the following notion of homomorphic embedding.

Given models $\mathfrak{A}$ and $\mathfrak{A}'$ for $\mathscr{L}$, $\mathfrak{A}$ is *homomorphic to* $\mathfrak{A}'$ iff there is a function f mapping A onto A' satisfying the following:

(i). For each n-placed relation R of $\mathfrak{A}$ and the corresponding relation R' of $\mathfrak{A}'$, and all elements $x_1, \dots, x_n$ of A,

$$\text{if} \quad R(x_1 \dots x_n), \quad \text{then} \quad R'(f(x_1) \dots f(x_n)).$$

(ii). For each m-placed function G of $\mathfrak{A}$ and the corresponding G' of $\mathfrak{A}'$, and for all $x_1, \dots, x_m$ of A,

$$f(G(x_1 \dots x_m)) = G'(f(x_1) \dots f(x_m)).$$

(iii). For each constant x of $\mathfrak{A}$, $f(x)$ is the corresponding constant of $\mathfrak{A}'$.

A function f satisfying the above is called a *homomorphism of* $\mathfrak{A}$ *onto* $\mathfrak{A}'$. We write $\mathfrak{A} \simeq_f \mathfrak{A}'$ to indicate that f is such a homomorphism; if it is not necessary to indicate f, we write $\mathfrak{A} \simeq \mathfrak{A}'$ for $\mathfrak{A}$ is homomorphic to $\mathfrak{A}'$. In this case we also say $\mathfrak{A}'$ is *a homomorphic image of* $\mathfrak{A}$. $\mathfrak{A}$ is *homomorphically embedded in* $\mathfrak{A}'$ iff $\mathfrak{A}$ is homomorphic to some submodel of $\mathfrak{A}'$. See Exercise 2.1.3 for some elementary properties of these notions. The next proposition corresponds to Proposition 2.1.8.

PROPOSITION 2.1.12. *Let $\mathfrak{A}$, $\mathfrak{B}$ be models for $\mathscr{L}$. Then $\mathfrak{A}$ is homomorphically embedded in $\mathfrak{B}$ if and only if some expansion of $\mathfrak{B}$ is a model of the positive diagram of $\mathfrak{A}$.*

COROLLARY 2.1.13. *Every partial order on a set X can be extended to a simple order on X.*

PROOF. Suppose that $\leqslant$ partially orders X. Let $\mathfrak{A} = \langle X, \leqslant \rangle$. Let $\{c_x : x \in X\}$ be distinct constants for $x \in X$ and let Δ be the positive diagram of $\mathfrak{A}$. Let

$$\Sigma = \Delta \cup \{c_x \not\equiv c_y : x \neq y \text{ in } X\} \cup \{\sigma\},$$

where σ is the sentence which expresses that $\leqslant$ is a simple order (see 1.4.1). Let Σ' be a finite subset of Σ involving, say, the elements $x_1, \dots, x_n$ and the corresponding constants. We need the following fact:

(1) Every partial order $\leqslant$ on $\{x_1, \dots, x_n\}$ can be extended to a simple order $\leqslant'$ on $\{x_1, \dots, x_n\}$ so that $\leqslant$ is preserved, i.e., if $x_i \leqslant x_j$, then $x_i \leqslant' x_j$.

The proof of (1) is not difficult and proceeds by induction on n. Assuming (1), we see that $\langle \{x_1, \dots, x_n\}, \leqslant' \rangle$ is a model of Σ'. By compactness, Σ has a simply ordered model $\langle Y, \leqslant' \rangle$, in which there is an element y_x corresponding to each constant c_x. Clearly the set $\{y_x : x \in X\}$ is simply ordered by $\leqslant'$. If $x \leqslant z$, then $y_x \leqslant' y_z$, and if $x \neq z$, then $y_x \neq y_z$. Using the inverse of the 1–1 function $y : x \to y_x$, we can induce a simple order on X which extends $\leqslant$. ⊣

EXERCISES

2.1.1. Show that there are also *countable* nonstandard models of complete number theory.

2.1.2. Prove the representation theorem for Boolean algebras (Proposition 1.4.4) by the method of diagrams.

[*Hint*: (a). Every atomic Boolean algebra is isomorphic to a field of sets. (b). Every finite subset of a Boolean algebra generates a finite, therefore atomic, Boolean algebra. (c). If $\mathfrak{A}$ is isomorphically embedded in a field of sets, then $\mathfrak{A}$ is isomorphic to a field of sets.]

2.1.3. Prove the following. The homomorphism relation $\simeq$ is reflexive and transitive. It is not symmetric nor antisymmetric. If $\mathfrak{A} \simeq \mathfrak{B}$, then $|A| \geqslant |B|$. A sentence σ is called *positive* iff it is built up from atomic formulas using only $\wedge, \vee, \exists, \forall$. If $\mathfrak{A} \simeq \mathfrak{B}$, σ is a positive sentence, and $\mathfrak{A} \models \sigma$, then $\mathfrak{B} \models \sigma$. Compare this with Exercise 1.3.5.

2.1.4. Prove the assertion (1) in Corollary 2.1.13.

2.1.5. Show that every ordered field is equivalent to a non-Archimedean ordered field.

2.1.6. Show that every group which has elements of arbitrarily large finite order is equivalent to a group which has an element of infinite order.

2.1.7. Show that every model of ZF is equivalent to a (countable) model $\langle A, E\rangle$ which has an infinite sequence

$$\dots Ex_2 Ex_1 Ex_0.$$

Therefore every model of ZF is equivalent to a countable model which is not isomorphic to a transitive model.

2.1.8. Let $\mathfrak{A} = \langle A, \leqslant, \dots\rangle$ be an infinite model such that $\leqslant$ well orders A. Show that there is a model $\mathfrak{A}' = \langle A', \leqslant', \dots\rangle$ equivalent to $\mathfrak{A}$ such that $\leqslant'$ is not a well ordering.

2.1.9. Show that every infinite model $\mathfrak{A}$ for a language $\mathscr{L}$ has an equivalent model $\mathfrak{B}$ of power $||\mathscr{L}||$ such that not every element of B is a constant of $\mathfrak{B}$.

2.1.10. Let $\mathscr{L}$ have no function or constant symbols. Let T be a theory in $\mathscr{L}$ and $\mathfrak{A}$ be a model for $\mathscr{L}$. Then $\mathfrak{A}$ is homomorphically embedded in some model of T if and only if every finite submodel of $\mathfrak{A}$ is homomorphically

embedded in some model of T. (This is a homomorphism version of Corollary 2.1.9.)

2.1.11. Let $\mathfrak{A}$ be an arbitrary infinite model and let $\alpha \geqslant ||\mathscr{L}||$. Then there is a model $\mathfrak{B}$ equivalent to $\mathfrak{A}$ such that for every formula $\varphi(x)$ with one free variable, if $\varphi(x)$ is satisfied by infinitely many different elements of $\mathfrak{B}$, then $\varphi(x)$ is satisfied by α different elements of $\mathfrak{B}$.

2.1.12. A model $\mathfrak{C}$ is said to be *finitely generated* iff there is a finite set $X \subset B$ which generates $\mathfrak{C}$ (see Exercise 1.3.9). Let T be a theory in $\mathscr{L}$ and let $\mathfrak{A}$ be a model for $\mathscr{L}$. Then $\mathfrak{A}$ is isomorphically embedded in some model of T if and only if every finitely generated submodel of $\mathfrak{A}$ is isomorphically embedded in some model of T. (Compare with Corollary 2.1.9.)

2.1.13

(i). If T_1 and T_2 are two theories such that $T_1 \cup T_2$ has no models, then there is a sentence φ such that $T_1 \models \varphi$ and $T_2 \models \neg\varphi$.

(ii). If T_1 and T_2 are two theories such that for all $\mathfrak{A}$, $\mathfrak{A}$ is a model of T_1 iff $\mathfrak{A}$ is not a model of T_2, then T_1 and T_2 are finitely axiomatizable.

2.1.14. Let $T_1 \subset T_2 \subset T_3 \subset \dots$ be a strictly increasing sequence of closed theories in $\mathscr{L}$. Show that their union $T = \bigcup_{n<\omega} T_n$ is a consistent closed theory in $\mathscr{L}$ and it is not finitely axiomatizable.

2.1.15. Let T_n, $n \in \omega$, be a strictly increasing sequence of closed theories in a finite language $\mathscr{L}$. Prove that $\bigcup_n T_n$ has an infinite model.

2.1.16. Let T be a finitely axiomatizable theory with only a countable number of complete extensions in a language $\mathscr{L}$. Prove that T has a finitely axiomatizable complete extension in $\mathscr{L}$.

2.1.17. Prove that every complete theory T in a countable language has a model $\mathfrak{A}$ of power $\leqslant 2^\omega$ such that for every $\mathfrak{B} \models T$ and every $S \subseteq B$ there is an $R \subseteq A$ such that $(\mathfrak{B}, S)$ is elementarily equivalent to $(\mathfrak{A}, R)$.

2.1.18*. Let Δ be the theory of dense linear order without endpoints. Prove the following lemma (a), and then use (a) and the Löwenheim–Skolem–Tarski Theorem to give a simpler proof of Theorem 1.5.3 on the elimination of quantifiers for the theory Δ.
(a). Let $\mathfrak{A}$ and $\mathfrak{B}$ be countable models of $\Delta, a_1, \dots, a_n \in A$, and $b_1, \dots, b_n \in B$. If $a_1, \dots, a_n$ and $b_1, \dots, b_n$ satisfy the same arrangement, then $(\mathfrak{A}, a_1 \dots a_n) \cong (\mathfrak{B}, b_1 \dots b_n)$.

2.1.19*. Let $\mathscr{L} = \emptyset$ be the language of pure identity theory. Prove the following lemmas (a) and (b), and then use (a), (b), and the Löwenheim–Skolem–Tarski Theorem to give a simpler proof of Theorem 1.5.7 on the elimination of quantifiers for pure identity theory.
(a). Let $\mathfrak{A}$ and $\mathfrak{B}$ be models for $\mathscr{L}$ of the same cardinality, $a_1, \ldots, a_n \in A$, and $b_1, \ldots, b_n \in B$. If $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ satisfy the same arrangement, then $(\mathfrak{A}, a_1 \ldots a_n) \cong (\mathfrak{B}, b_1 \ldots b_n)$.
(b). Let $\varphi(x_1 \ldots x_n)$ be a formula of $\mathscr{L}$ and $\theta(x_1 \ldots x_n)$ be an arrangement. Let S be the set of all cardinals α such that $\varphi \wedge \theta$ is satisfiable in a model of cardinality α. Then either S or its complement is a finite set of finite cardinals.

2.1.20. This and some of the following exercises are designed to show an alternative method of proving the extended completeness theorem for countable languages. A generalization of this method to noncountable languages is also given later.

Let $\mathscr{L}$ be a countable language, and let T be a consistent set of sentences closed under $\vdash$. We aim to prove that T has a model. The starting point of our discussion is the countable Lindenbaum algebra $\mathfrak{B}_{\mathscr{L}}$. We have already seen in Exercise 1.4.11 that T corresponds to a filter in $\mathfrak{B}_0$, the Lindenbaum algebra of sentences of $\mathscr{L}$. It is also easy to show that the set

$$\Phi = \{\varphi : T \vdash \varphi \quad \text{and} \quad \varphi \text{ is a formula of } \mathscr{L}\}$$

is a filter in $\mathfrak{B}$. For simplicity, we shall now operate in the quotient algebra $\mathfrak{B}/\Phi$. In other words, the equivalence classes of this new algebra are given by sets of formulas

$$(\psi) = \{\varphi : T \vdash \varphi \leftrightarrow \psi\},$$

with its unit element given by the set Φ, and its zero element given by

$$\{\varphi : T \vdash \neg\varphi\}.$$

We denote this quotient algebra by $\mathfrak{B}_T$ and call it the *Lindenbaum algebra of* T. $\mathfrak{B}_T$ is obviously a countable Boolean algebra.

2.1.21. Let $\mathfrak{A}$ be any Boolean algebra and let Y be a subset of A. The *sum* of Y, or the l.u.b. of Y, is defined to be the unique $y \in A$ such that

$x \leqslant y$ for all $x \in Y$ (i.e. y is an *upper bound* for Y), and

if $z \in A$ is any upper bound for Y, then $y \leqslant z$.

We denote the sum of Y if it exists by $\bigvee Y$, or if the elements of Y are indexed by I, $\bigvee_{i \in I} y_i$. In an entirely similar manner, we can define the

product of Y, or the g.l.b. of Y, and denote it by $\bigwedge Y$ or $\bigwedge_{i \in I} y_i$. Sums and products of arbitrary $Y \subset A$ do not necessarily exist. When they do exist, they satisfy the following identities (assume that $\bigvee_{i \in I} y_i$ exists):

$$(\bigvee_{i \in I} y_i) + x = \bigvee_{i \in I} (y_i + x),$$

$$(\bigvee_{i \in I} y_i) \cdot x = \bigvee_{i \in I} (y_i \cdot x),$$

$$\overline{\bigvee_{i \in I} y_i} = \bigwedge_{i \in I} \overline{y_i}.$$

These identities imply, of course, that the sums and products on the right-hand side also exist. We leave the duals involving $\bigwedge$ to the reader.

Let φ be any formula of $\mathscr{L}$. Let $\varphi(k/p)$ be the formula obtained from φ by first replacing all bound occurrences of v_p in φ by v_j, the first variable in the sequence $v_0, v_1, \ldots$, not occurring in φ, and then replacing all free occurrences of v_k by v_p. Show that in the Boolean algebra $\mathfrak{B}_T$,

$$\bigvee_{p \in \omega} (\varphi(k/p)) = ((\exists v_k)\varphi),$$

$$\bigwedge_{p \in \omega} (\varphi(k/p)) = ((\forall v_k)\varphi).$$

Thus sums and products of certain sets of substitution instances of a single formula φ always exist and correspond to existential and universal quantification of φ. Note that the number of such sums (and products) in $\mathfrak{B}_T$ is countable.

2.1.22. An ultrafilter D on $\mathfrak{A}$ is said to *preserve the sum* $\bigvee_{i \in I} y_i$ iff

$$\bigvee_{i \in I} y_i \in D \quad \text{if and only if some} \quad y_i \in D.$$

Similarly, *D preserves the product* $\bigwedge_{i \in I} y_i$ iff

$$\bigwedge_{i \in I} y_i \in D \quad \text{if and only if all} \quad y_i \in D.$$

Prove the following: Given a countable sequence of products $\bigwedge X_0$, $\bigwedge X_1$, ..., $\bigwedge X_n$, ... of $\mathfrak{A}$. Then there exists an ultrafilter D on $\mathfrak{A}$ which preserves each product.

[*Hint*: Pick a sequence $x_n \in X_n$ such that no finite product of elements of the form $\bigwedge X_n + \bar{x}_n$ is equal to zero. Now consider any ultrafilter D which has as elements all $\bigwedge X_n + \bar{x}_n$.]

There is also a corresponding result about countable sequences of sums.

2.1.23. Let D be any ultrafilter on $\mathfrak{B}_T$ which preserves all the products of Exercise 2.1.21. We shall now construct a model of T from the variables

$v_0, v_1, \dots$ of $\mathscr{L}$. Since the procedure is quite similar to that of Lemma 2.1.2, we ask the reader to fill in all the details.

First define equivalence by

$$v_i \sim v_j \quad \text{iff} \quad (v_i \equiv v_j) \in D.$$

The equivalence classes are denoted by $\tilde{v}_i$.

Let c be a constant symbol of $\mathscr{L}$. Since $\vdash (\exists v_0)(c \equiv v_0)$, and D preserves sums, we see that for some i, $(c \equiv v_i) \in D$. Let the interpretation of c be the class $\tilde{v}_i$.

Let t be any term of $\mathscr{L}$ (this includes the cases of function symbols and constant symbols) and v_p be a variable not occurring in t. Then $\vdash (\exists v_p)(t \equiv v_p)$. Since D preserves sums, for some j, $(t \equiv v_j) \in D$. Let the interpretation of the term t (defined on equivalence classes $\tilde{v}_i$) be $\tilde{v}_j$.

Finally, let P be a relation symbol of $\mathscr{L}$. We define the relation R by

$$R(\tilde{v}_{i_1} \dots \tilde{v}_{i_n}) \quad \text{iff} \quad (P(v_{i_1} \dots v_{i_n})) \in D.$$

In this way, we have defined in an unambiguous manner a model $\mathfrak{A}$ for $\mathscr{L}$ with universe the set of equivalence classes $\tilde{v}_i$. Now prove by induction on the formulas $\varphi(v_0 \dots v_n)$ of $\mathscr{L}$ that

$$\mathfrak{A} \models \varphi[\tilde{v}_0 \dots \tilde{v}_n] \quad \text{iff} \quad (\varphi(v_0 \dots v_n)) \in D.$$

To pass through the cases of $\forall$ or $\exists$, we again need the fact that D preserves sums. Since $\varphi \in T$ implies that $(\varphi) \in D$, this shows that $\mathfrak{A}$ is a model of T.

2.1.24. If the language $\mathscr{L}$ is uncountable, then the number of sums and products corresponding to $\exists$ and $\forall$ in $\mathfrak{B}_T$ is also uncountable. Even though, in general, Exercise 2.1.17 fails for uncountable sequences of products in an arbitrary Boolean algebra, there is, nevertheless, a version of it which holds for the algebra $\mathfrak{B}_T$. This is because every formula φ contains only a finite number of symbols. The generalization of Exercise 2.1.17 is as follows (the proof is straightforward):

Let $\mathfrak{A}$ be a Boolean algebra, α be an infinite cardinal, and $\bigwedge X_\beta$, $\beta < \alpha$, be a sequence of products of $\mathfrak{A}$. Suppose that for all $\beta < \alpha$ and all filters E on $\mathfrak{A}$ generated by fewer than α elements,

$$\text{whenever} \quad X_\beta \subset E, \quad \text{then} \quad \bigwedge X_\beta \in E.$$

Then there exists an ultrafilter D on $\mathfrak{A}$ which preserves each product $\bigwedge X_\beta$.

Using this result, a generalization of the proof in Exercise 2.1.18 can be given for noncountable languages $\mathscr{L}$. (A technical detail should be mentioned: Before proceeding with the proof we must first expand $\mathscr{L}$ to a language $\mathscr{L}'$ with $\|\mathscr{L}\|$ new constant symbols. This is apparently necessary, see the proof of Lemma 2.1.1.)

2.2. Refinements of the method. Omitting types and interpolation theorems

In this section, we shall give two refinements of the method used in Section 2.1 to construct countable models with additional properties. The first refinement will lead us to the omitting types theorem. At the moment, the possible ramifications of this technique to noncountable languages and models are not yet fully understood. We shall mention only a couple of results for noncountable languages.

The starting point of our discussion is the notion of a set Σ of formulas of $\mathscr{L}$ *in the (free) variables* $x_1, \dots, x_n$. Here we are using $x_1, x_2, \dots$ as names for arbitrary free variables of $\mathscr{L}$. We could just as well use $v_{m_1}, v_{m_2}, \dots$, but we abhor double subscripts. The following is a precise definition: Σ is a set of formulas of $\mathscr{L}$ in the (free) variables $x_1, \dots, x_n$ (symbolically, $\Sigma = \Sigma(x_1 \dots x_n)$) iff $x_1, \dots, x_n$ are distinct individual variables and every formula σ in Σ contains at most the variables $x_1, \dots, x_n$ free. We now introduce the convention $\sigma = \sigma(x_1 \dots x_n)$, as we did for $\varphi = \varphi(v_0 \dots v_n)$. If $\sigma = \sigma(x_1 \dots x_n)$, then the notation

$$\mathfrak{A} \models \sigma[a_1 \dots a_n]$$

means that the sequence $a_1, \dots, a_n$ of A satisfies σ in $\mathfrak{A}$ (see the section on satisfaction). It is useful also to introduce the notation

$$\mathfrak{A} \models \Sigma[a_1 \dots a_n]$$

to mean that for every $\sigma \in \Sigma$, $a_1, \dots, a_n$ satisfies σ in $\mathfrak{A}$; in this case we say that $a_1, \dots, a_n$ *satisfies*, or *realizes*, Σ *in* $\mathfrak{A}$. If $c_1, \dots, c_n$ is a sequence of constant symbols, then $\sigma(c_1 \dots c_n)$ denotes the sentence formed by simultaneously replacing each free occurrence of x_i, $1 \leqslant i \leqslant n$, in σ by the corresponding c_i. Sometimes we shall replace just *some* of the variables by constants. If $m \leqslant n$, the notation $\sigma(c_1 \dots c_m x_{m+1} \dots x_n)$ is self-explanatory.

For reasons explained in Section 2.1 (before Lemma 2.1.1), we must be careful to use the above notation only in a context where the list of variables $x_1, \dots, x_n$ is given. A completely unambiguous notation can be introduced, but at great cost in readability. For example, we could use the notation

$$\mathfrak{A} \models \sigma[a_1/x_1 \dots a_n/x_n] \quad \text{for} \quad \mathfrak{A} \models \sigma[a_1 \dots a_n],$$

$$\sigma(c_1/x_1 \dots c_m/x_m x_{m+1} \dots x_n) \quad \text{for} \quad \sigma(c_1 \dots c_m x_{m+1} \dots x_n).$$

Let Σ be a set of formulas in the variables $x_1, \dots, x_n$, and let $\mathfrak{A}$ be a model for $\mathscr{L}$. We say that $\mathfrak{A}$ *realizes* Σ iff some n-tuple of elements of A satisfies Σ in $\mathfrak{A}$. We say that $\mathfrak{A}$ *omits* Σ iff $\mathfrak{A}$ does not realize Σ. The phrase *Σ is satisfiable in* $\mathfrak{A}$ has exactly the same meaning as $\mathfrak{A}$ *realizes* Σ. Σ is *consistent* iff Σ is satisfiable in some model.

Example 2.2.1. Let T be Peano arithmetic and let $\Sigma(x)$ be the set

$$\{0 \not\equiv x, S0 \not\equiv x, SS0 \not\equiv x, \dots\}.$$

An element is said to be *nonstandard* iff it realizes $\Sigma(x)$. The standard model of T omits $\Sigma(x)$, while all the nonstandard models realize $\Sigma(x)$.

Example 2.2.2. Let T be the theory of ordered fields and let $\Sigma(x)$ be the set

$$\{1 \leqslant x, 1+1 \leqslant x, 1+1+1 \leqslant x, \dots\}.$$

An element is said to be *positive infinite* iff it realizes $\Sigma(x)$. An ordered field omits $\Sigma(x)$ if and only if it is Archimedean. The ordered fields of rationals and reals omit $\Sigma(x)$. Non-Archimedean ordered fields were constructed in the last section using the compactness theorem.

Example 2.2.3. Let T be the theory of Abelian groups and let $\Sigma(x)$ be the set

$$\{x \not\equiv 0, 2x \not\equiv 0, 3x \not\equiv 0, \dots\}.$$

Elements which realize $\Sigma(x)$ are said to be of *infinite order*. An Abelian group which omits $\Sigma(x)$ is said to be a *torsion group*. Thus in a torsion group, every element has a finite multiple which is zero.

Example 2.2.4. Here is an example of a set of formulas with infinitely many variables. Let T be the theory of partial order and let Σ be the set

$$\{x_1 < x_0, x_2 < x_1, x_3 < x_2, \dots\}.$$

A model $\mathfrak{A}$ of T omits Σ iff $\mathfrak{A}$ is a *well founded* partial ordering. A linear ordering $\mathfrak{A}$ omits Σ iff it is a well ordering.

Example 2.2.5. By a *type* $\Gamma(x_1 \dots x_n)$ in the variables $x_1, \dots, x_n$ we mean a maximal consistent set of formulas of $\mathscr{L}$ in these variables. Given any model $\mathfrak{A}$ and n-tuple $a_1, \dots, a_n \in A$, the set $\Gamma(x_1 \dots x_n)$ of all formulas $\gamma(x_1 \dots x_n)$ satisfied by $a_1, \dots, a_n$ is a type, and, in fact, is the unique type realized by $a_1, \dots, a_n$. It is called the *type of* $a_1, \dots, a_n$ in $\mathfrak{A}$.

Example 2.2.6. Let $\mathfrak{A}$ be the ordered field of real numbers. Then any two distinct elements $a, b \in A$ have different types. For if $a < b$, there is a rational number r with $a < r < b$; hence a satisfies $x < r$, while b does not. Thus $\mathfrak{A}$ realizes 2^ω different types in one variable.

The next proposition answers the question: When is a set of formulas realized by some model of a theory T? Its proof is a simple application of the compactness theorem.

Proposition 2.2.7. *Let T be a theory and let $\Sigma = \Sigma(x_1 \dots x_n)$. The following are equivalent*:

(i). *T has a model which realizes Σ.*

(ii). *Every finite subset of Σ is realized in some model of T.*

(iii). $T \cup \{(\exists x_1 \dots x_n)(\sigma_1 \wedge \dots \wedge \sigma_m) : m < \omega, \sigma_1, \dots, \sigma_m \in \Sigma\}$ *is consistent.*

We shall say that a formula $\sigma(x_1 \dots x_n)$ is *consistent with* a theory T iff there is a model $\mathfrak{A}$ of T which realizes $\{\sigma\}$, and we say that $\Sigma(x_1 \dots x_n)$ is *consistent with* T iff T has a model which realizes Σ. Thus (i)–(iii) above are all equivalent to the statement that Σ is consistent with T.

We now take up the question: When is a set Σ of formulas in $x_1, \dots, x_n$ *omitted* in some model of a theory T? This is a more difficult question, and we need more than the compactness theorem to answer it. The key theorem of this section, Theorem 2.2.9, gives a necessary and sufficient condition for T to have a model which omits Σ. The ω-completeness theorem 2.2.13 is one of a long list of consequences of it. We shall use Theorem 2.2.9 in the next section, and again later on. If Σ is a finite set of formulas, then there is no problem in determining whether Σ can be omitted, because the sentence

$$\varphi = (\exists x_1 \dots x_n)(\sigma_1 \wedge \dots \wedge \sigma_m),$$

where $\Sigma = \{\sigma_1, \dots, \sigma_m\}$, and its negation $\neg\varphi$ express, respectively, that Σ is realized or omitted. Thus the interesting case is where Σ is infinite.

Let us first take another look at Lemma 2.1.2. So far, we have only used the property that every element of $\mathfrak{A}$ is the interpretation of a constant $c \in C$ in a simple way, to show that $|A| \leqslant |C|$. In this section, we shall make much more use of that property of $\mathfrak{A}$.

The central idea in dealing with our problem is the notion of a theory locally realizing a set of formulas.

Let $\Sigma = \Sigma(x_1 \dots x_n)$ be a set of formulas of $\mathscr{L}$. A theory T in $\mathscr{L}$ is said to *locally realize* Σ iff there is a formula $\varphi(x_1 \dots x_n)$ in $\mathscr{L}$ such that:

(i). φ is consistent with T.
(ii). For all $\sigma \in \Sigma$, $T \models \varphi \to \sigma$.

That is, every n-tuple in a model of T which satisfies φ realizes Σ.

We say that T *locally omits* Σ iff T does not locally realize Σ. Thus T locally omits Σ if and only if for every formula $\varphi(x_1 \dots x_n)$ which is consistent with T, there exists $\sigma \in \Sigma$ such that $\varphi \wedge \neg \sigma$ is consistent with T.

For complete theories we have a simple proposition:

Proposition 2.2.8. *Let T be a complete theory in $\mathscr{L}$, and let $\Sigma = \Sigma(x_1 \dots x_n)$ be a set of formulas of $\mathscr{L}$. If T has a model which omits Σ, then T locally omits Σ.*

Proof. The proposition may be restated as follows: If T locally realizes Σ, then every model of T realizes Σ. Suppose T locally realizes Σ and let $\varphi(x_1 \dots x_n)$ be a formula consistent with T such that $T \models \varphi \to \sigma$, $\sigma \in \Sigma$.

Let $\mathfrak{A}$ be a model of T. Since T is complete, $T \models (\exists x_1 \dots x_n)\varphi$. So some n-tuple $a_1, \dots, a_n$ satisfies φ in $\mathfrak{A}$. Then $a_1, \dots, a_n$ satisfies each $\sigma \in \Sigma$, and hence realizes Σ in $\mathfrak{A}$. ⊣

The omitting types theorem is a converse of the above proposition. It holds, in fact, for arbitrary consistent theories in a countable language.

Theorem 2.2.9 (Omitting Types Theorem). *Let T be a consistent theory in a countable language $\mathscr{L}$, and let $\Sigma(x_1 \dots x_n)$ be a set of formulas. If T locally omits Σ, then T has a countable model which omits Σ.*

Proof. To simplify notation, let $\Sigma(x)$ be a set of formulas in one variable x. Suppose T locally omits $\Sigma(x)$. Let $C = \{c_0, c_1, \dots\}$ be a countable set of new constant symbols not already in $\mathscr{L}$ and let $\mathscr{L}' = \mathscr{L} \cup C$. Then $\mathscr{L}'$ is countable. Arrange all the sentences of $\mathscr{L}'$ in a list $\varphi_0, \varphi_1, \varphi_2, \dots$. We shall construct an increasing sequence of consistent theories

$$T = T_0 \subset T_1 \subset \dots \subset T_m \subset \dots$$

such that:

(1). Each T_m is a consistent theory of $\mathscr{L}'$ which is a finite extension of T.
(2). Either $\varphi_m \in T_{m+1}$ or $(\neg \varphi_m) \in T_{m+1}$.
(3). If $\varphi_m = (\exists x)\psi(x)$ and $\varphi_m \in T_{m+1}$, then $\psi(c_p) \in T_{m+1}$, where c_p is the first constant not occurring in T_m or φ_m.
(4). There is a formula $\sigma(x) \in \Sigma(x)$ such that $(\neg \sigma(c_m)) \in T_{m+1}$.

Assuming we already have the theory T_m, we construct T_{m+1} as follows: Let $T_m = T \cup \{\theta_1, \dots, \theta_r\}$, $r > 0$, and let $\theta = \theta_1 \wedge \dots \wedge \theta_r$. Let $c_0, \dots, c_n$ contain all the constants from C occurring in θ. Form the formula $\theta(x_m)$ of $\mathscr{L}$ by replacing each constant c_i by x_i (renaming bound variables if necessary), and prefixing by $\exists x_i$, $i \neq m$. Then $\theta(x_m)$ is consistent with T. Therefore, for some $\sigma(x) \in \Sigma(x)$, $\theta(x_m) \wedge \neg \sigma(x_m)$ is consistent with T. Put the sentence $\neg \sigma(c_m)$ into T_{m+1}. This makes (4) hold.

If φ_m is consistent with $T_m \cup \{\neg \sigma(c_m)\}$, put φ_m into T_{m+1}. Otherwise put $(\neg \varphi_m)$ into T_{m+1}. This takes care of (2). If $\varphi_m = (\exists x)\psi(x)$ is consistent with $T_m \cup \{\neg \sigma(c_m)\}$, put $\psi(c_p)$ into T_{m+1}. This takes care of (3). The theory T_{m+1} is a consistent finite extension of T_m. Thus (1)–(4) hold for T_{m+1}.

Let $T_\omega = \bigcup_{n<\omega} T_n$. From (1) and (2) we see that T_ω is a maximal consistent theory in $\mathscr{L}'$. Let $\mathfrak{B}' = (\mathfrak{B}, b_0, b_1, \dots)$ be a countable model of T_ω, and let $\mathfrak{A}' = (\mathfrak{A}, b_0, b_1, \dots)$ be the submodel of $\mathfrak{B}'$ generated by the constants $b_0, b_1, \dots$. We then see from (3) that

$$A = \{b_0, b_1, \dots\}.$$

Moreover, using (3) and the completeness of T_ω, we can show by induction on the complexity of a sentence φ in $\mathscr{L}'$ that

$$\mathfrak{A}' \models \varphi, \qquad \mathfrak{B}' \models \varphi, \qquad T_\omega \models \varphi$$

are all equivalent. Thus $\mathfrak{A}'$ is a model of T_ω and hence $\mathfrak{A}$ is a model of T. Finally, condition (4) ensures that $\mathfrak{A}$ omits Σ. ⊣

When T is a complete theory, we see that locally omitting $\Sigma(x_1 \dots x_n)$ is a necessary and sufficient condition for T to have a model omitting Σ. Here is a necessary and sufficient condition which works in general.

Corollary 2.2.10. *Let $\mathscr{L}$ be countable. A theory T has a (countable) model omitting $\Sigma(x_1 \dots x_n)$ if and only if some complete extension of T locally omits $\Sigma(x_1 \dots x_n)$.*

Example 2.2.11. Consider the language $\mathscr{L} = \{+, \cdot, S, 0\}$. We abbreviate $1 = S0$, $2 = SS0$, $3 = SSS0$, ... By an *ω-model* we mean a model $\mathfrak{A}$ in which

$$A = \{0, 1, 2, 3, \dots\} \quad ,$$

that is, $\mathfrak{A}$ omits the set $\{x \neq 0, x \neq 1, x \neq 2, \dots\}$. A theory T in $\mathscr{L}$ is said to be *ω-consistent* iff there is no formula $\varphi(x)$ of $\mathscr{L}$ such that

$$T \vDash \varphi(0), \quad T \vDash \varphi(1), \quad T \vDash \varphi(2), \ldots$$

and

$$T \vDash (\exists x) \neg \varphi(x).$$

T is said to be *ω-complete* iff for every formula $\varphi(x)$ of $\mathscr{L}$ we have

$$T \vDash \varphi(0),\ T \vDash \varphi(1),\ T \vDash \varphi(2), \ldots \quad \text{implies} \quad T \vDash (\forall x)\varphi(x).$$

It follows from the omitting types theorem that:

PROPOSITION 2.2.12. *Let T be a consistent theory in $\mathscr{L}$.*

(i). *If T is ω-complete, then T has an ω-model.*

(ii). *If T has an ω-model, then T is ω-consistent.*

PROOF. (i).. We show that T locally omits the set $\Sigma(x) = \{x \not\equiv 0, x \not\equiv 1, \ldots\}$. Suppose $\theta(x)$ is consistent with T. Then $T \vDash (\forall x) \neg \theta(x)$ fails. By ω-completeness, there is an n such that not $T \vDash \neg \theta(n)$. Hence $\theta(n)$ is consistent with T, so $\theta(x) \wedge \neg x \not\equiv n$ is consistent with T. Thus T locally omits $\Sigma(x)$.

(ii). Trivial. ⊣

The *ω-rule* is the following infinite rule of proof: From $\varphi(0), \varphi(1), \varphi(2), \ldots$, infer $(\forall x)\varphi(x)$, where $\varphi(x)$ is any formula of $\mathscr{L}$. *ω-logic* is formed by adding the ω-rule to the axioms and rules of inference of the first-order logic $\mathscr{L}$ and allowing infinitely long proofs. We have the following completeness theorem for ω-logic.

PROPOSITION 2.2.13 (ω-Completeness Theorem). *A theory T in $\mathscr{L}$ is consistent in ω-logic if and only if T has an ω-model.*

PROOF. Let T' be the set of all sentences of $\mathscr{L}$ provable from T in ω-logic. Then T is consistent in ω-logic if and only if T' is consistent in $\mathscr{L}$. Moreover, T' is ω-complete. Therefore T' has an ω-model if and only if T' is consistent. ⊣

The formulation of ω-logic above is aimed at studying the standard model of arithmetic. A useful generalization, which we shall call generalized ω-logic, is aimed at studying ordinary models for first order logic enriched by a symbol for the set of natural numbers.

EXAMPLE 2.2.11′. Let $\mathscr{L}'$ be a countable language which has among its symbols a special unary relation symbol N and special constant symbols

$0, 1, 2, \ldots$. By an *ω-model* for $\mathscr{L}'$ we mean a model $\mathfrak{A}$ for $\mathscr{L}'$ in which N is interpreted by the set ω of natural numbers, and $0, 1, 2, \ldots$ are interpreted by themselves. In an ω-model, ω is a subset of the universe A, but we allow A to contain elements outside of ω or even to be uncountable.

Let T_N be the special set of sentences

$$T_N = \{N(m) : m < \omega\} \cup \{\neg m \equiv n : m < n < \omega\}$$

which state that the natural numbers are distinct and belong to N. T_N holds in every ω-model for $\mathscr{L}'$. A theory T in $\mathscr{L}'$ is said to be *ω-consistent* iff there is no formula $\varphi(x)$ of $\mathscr{L}'$ such that

$$T_N \cup T \vDash \varphi(0),\ T_N \cup T \vDash \varphi(1),\ T_N \cup T \vDash \varphi(2), \ldots$$

and

$$T_N \cup T \vDash (\exists x)(N(x) \wedge \neg\varphi(x)).$$

T is said to be *ω-complete* iff for every formula $\varphi(x)$ of $\mathscr{L}'$ we have

$$T_N \cup T \vDash \varphi(0),\ T_N \cup T \vDash \varphi(1),\ T_N \cup T \vDash \varphi(2), \ldots$$

implies

$$T_N \cup T \vDash (\forall x)(N(x) \rightarrow \varphi(x)).$$

The *ω-rule* for $\mathscr{L}'$ is the infinite rule: From $\varphi(0), \varphi(1), \varphi(2), \ldots$, infer $(\forall x)(N(x) \rightarrow \varphi(x))$. By *generalized ω-logic* we mean first order logic for the language $\mathscr{L}'$ with T_N added as an additional set of logical axioms and the ω-rule added as an additional rule of proof.

Propositions 2.2.12 and 2.2.13 take the following form for generalized ω-logic.

PROPOSITION 2.2.12′. *Let T be a theory in $\mathscr{L}'$ such that $T_N \cup T$ is consistent.*

(i). *If T is ω-complete, then T has an ω-model.*

(ii). *If T has an ω-model, then T is ω-consistent.*

PROPOSITION 2.2.13′. *A theory T in $\mathscr{L}'$ is consistent in generalized ω-logic if and only if T has an ω-model.*

The following example shows that the omitting types theorem fails for sets of formulas with infinitely many free variables.

EXAMPLE 2.2.14. Let T be the theory of dense linear order without endpoints. Thus T is complete. Let $\Sigma(x_0 x_1 x_2 \ldots)$ be the set

$$\{x_1 < x_0, x_2 < x_1, x_3 < x_2, \ldots\}.$$

As we observed before, a model $\mathfrak{A}$ omits Σ if and only if $\mathfrak{A}$ is a well ordering. But T has no well ordered models, so no model of T omits Σ. However, T does locally omit Σ, because if $\varphi(x_0 x_1 \ldots x_n)$ is consistent with T, then $\varphi \wedge \neg x_{n+2} < x_{n+1}$ is consistent with T.

The omitting types theorem can be generalized to the case of countably many sets of formulas.

THEOREM 2.2.15 (Extended Omitting Types Theorem). *Let T be a consistent theory in a countable language $\mathscr{L}$, and for each $r < \omega$ let $\Sigma_r(x_1 \ldots x_{n_r})$ be a set of formulas in n_r variables. If T locally omits each Σ_r, then T has a countable model which omits each Σ_r.*

PROOF. Similar to the proof of the omitting types theorem. The only difference is that for each r the n_r-tuples of new constants are arranged in a list:

$$s^r_r, s^r_{r+1}, s^r_{r+2}, \ldots.$$

The theories T_m are built up so that for each $r = 0, 1, \ldots, m$, there is a formula $\sigma \in \Sigma_r$ such that $(\neg \sigma(s^r_m)) \in T_{m+1}$. ⊣

Here is a first application of the extended omitting types theorem. It uses the notion of an *elementary extension* which plays an important role in the rest of this book.

2.2.16. $\mathfrak{B}$ is said to be an *elementary extension* of $\mathfrak{A}$, $\mathfrak{A} \prec \mathfrak{B}$, iff

(i). $\mathfrak{B}$ is an extension of $\mathfrak{A}$, $\mathfrak{A} \subset \mathfrak{B}$.

(ii). For any formula $\varphi(x_1 \ldots x_n)$ of $\mathscr{L}$ and any $a_1, \ldots, a_n \in A$, $a_1, \ldots, a_n$ satisfies φ in $\mathfrak{A}$ if and only if it satisfies φ in $\mathfrak{B}$.

When $\mathfrak{B}$ is an elementary extension of $\mathfrak{A}$ we also say that $\mathfrak{A}$ is an *elementary submodel* of $\mathfrak{B}$.

A mapping $f: A \to B$ is said to be an *elementary embedding* of $\mathfrak{A}$ into $\mathfrak{B}$, in symbols $f: \mathfrak{A} \prec \mathfrak{B}$, iff for all formulas $\varphi(x_1 \ldots x_n)$ of $\mathscr{L}$ and n-tuples $a_1, \ldots, a_n \in A$, we have

$$\mathfrak{A} \models \varphi[a_1 \ldots a_n] \text{ if and only if } \mathfrak{B} \models \varphi[fa_1 \ldots fa_n].$$

An elementary embedding of $\mathfrak{A}$ into $\mathfrak{B}$ is thus the same thing as an isomorphism of $\mathfrak{A}$ onto an elementary submodel of $\mathfrak{B}$.

The following analogue of Proposition 2.1.8 is often useful.

PROPOSITION 2.2.17. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be models for $\mathscr{L}$ and let $f: A \to B$. Then the following are equivalent:*

(a). *f is an elementary embedding of $\mathfrak{A}$ into $\mathfrak{B}$.*

(b). *There is an elementary extension $\mathfrak{C} \succ \mathfrak{A}$ and an isomorphism $g: \mathfrak{C} \cong \mathfrak{B}$ such that $g \supset f$.*

(c). *$(\mathfrak{B}, fa)_{a \in A}$ is a model of the elementary diagram of $\mathfrak{A}$.*

Proposition 2.2.17 shows that the following three conditions are equivalent:

(a′) *$\mathfrak{A}$ is elementarily embeddable in $\mathfrak{B}$.*

(b′) *$\mathfrak{B}$ is isomorphic to an elementary extension of $\mathfrak{A}$.*

(c′) *$\mathfrak{B}$ can be expanded to a model of the elementary diagram of $\mathfrak{A}$.*

In the special case that $A \subset B$ and f is the indentity mapping from A into B, Proposition 2.2.17 shows that $\mathfrak{A}$ is an elementary submodel of $\mathfrak{B}$ if and only if $\mathfrak{B}_A$ is a model of the elementary diagram of $\mathfrak{A}$.

Let us now consider the theory ZF, Zermelo–Fraenkel set theory. A model $\mathfrak{B} = \langle B, F \rangle$ of ZF is said to be an *end extension* of a model $\mathfrak{A} = \langle A, E \rangle$ of ZF iff $\mathfrak{B}$ is a proper extension of $\mathfrak{A}$ and no member of A gets a new element, that is,

$$\text{if } a \in A \text{ and } b \in B, \text{ then } bFa \text{ implies } b \in A.$$

THEOREM 2.2.18. *Every countable model $\mathfrak{A} = \langle A, E \rangle$ of ZF has an end elementary extension.*

PROOF. Let $\mathscr{L}$ be the language with the symbol ϵ, a constant symbol $\bar{a}$ for each $a \in A$, and a new constant symbol c. Let T be the theory with the axioms

$$\mathrm{Th}((\mathfrak{A}, a)_{a \in A}),$$

$$c \notin \bar{a}, \quad \text{where} \quad a \in A.$$

T is consistent because every finite subset of T has a model of the form $(\mathfrak{A}, a, c)_{a \in A}$. For each $a \in A$, let $\Sigma_a(x)$ be the set of formulas

$$\Sigma_a(x) = \{x \epsilon \bar{a}\} \cup \{x \neq \bar{b} : bEa\}.$$

It suffices to show that T locally omits each set $\Sigma_a(x)$. For then T has a model $(\mathfrak{B}, a, c)_{a\in A}$ which omits each $\Sigma_a(x)$. We may also assume that $A \subset B$. $\mathfrak{B}$ is an elementary extension of $\mathfrak{A}$ because $\mathrm{Th}((\mathfrak{A}, a)_{a\in A}) \subset T$, whence $(\mathfrak{A}, a)_{a\in A} \equiv (\mathfrak{B}, a)_{a\in A}$. $\mathfrak{B}$ is a proper extension because $c \in B\setminus A$. Finally, $\mathfrak{B}$ is an end extension because it omits each $\Sigma_a(x)$.

To see that T locally omits each $\Sigma_a(x)$, we note that a formula $\varphi(x, c)$ of $\mathscr{L}$ is consistent with T if and only if

$$(\mathfrak{A}, a)_{a\in A} \models (\forall y)(\exists z)(\exists x)[z \notin y \wedge \varphi(x, z)].$$

Suppose $\varphi(x, c)$ is consistent with T, but $\varphi(x, c) \wedge \neg\, x \in \bar{a}$ is not. Then $\varphi(x, c) \wedge x \in \bar{a}$ is consistent with T. Using the axiom of replacement in ZF, we see in turn that the following sentences hold in $(\mathfrak{A}, a)_{a\in A}$:

$$(\forall y)(\exists z)(\exists x)[z \notin y \wedge \varphi(x, z) \wedge x \in \bar{a}]$$

$$(\exists x)(\forall y)(\exists z)[z \notin y \wedge \varphi(x, z) \wedge x \in \bar{a}].$$

Then for some $b \in A$, $\varphi(\bar{b}, c) \wedge \bar{b} \in \bar{a}$ is consistent with T, whence $\varphi(x, c) \wedge x \equiv \bar{b}$ is consistent with T. Thus T locally omits $\Sigma_a(x)$. ⊣

The omitting types theorem as it stands is false for uncountable languages. For example, let T be the theory with the axioms

$$c_\alpha \not\equiv c_\beta, \quad \alpha < \beta < \omega_1$$

in the language $\mathscr{L}$ with constants

$$\{c_\alpha : \alpha < \omega_1\} \cup \{d_n : n < \omega\}.$$

Let $\Gamma(x)$ be the set of formulas

$$\Gamma(x) = \{x \not\equiv d_n : n < \omega\}.$$

Then T locally omits $\Gamma(x)$. However no model of T omits $\Gamma(x)$ because every model of T is uncountable but each model which omits $\Gamma(x)$ is countable.

A more complicated counterexample where the theory T is complete has been given by Fuhrken (1962).

However, the omitting types theorem can be generalized to uncountable languages if we define the notion of 'locally omits' in the proper way. Let T be a theory and $\Sigma(x_1 \dots x_n)$ a set of formulas in a language $\mathscr{L}$ of power α. We say that T *α-realizes* Σ iff there is a set $\Phi(x_1 \dots x_n)$ of fewer than α formulas of $\mathscr{L}$ such that:

(i). Φ is consistent with T,

(ii). $T \cup \Phi(x_1 \dots x_n) \models \Sigma(x_1 \dots x_n)$,

that is, in any model $\mathfrak{A}$ of T, any n-tuple which realizes Φ realizes Σ. T is said to *α-omit* $\Sigma(x_1 \dots x_n)$ iff T does not α-realize $\Sigma(x_1 \dots x_n)$. Note that if Σ has power less than α, then T α-realizes Σ trivially. Thus only sets of formulas of power α can ever be α-omitted.

THEOREM 2.2.19 (α-Omitting Types Theorem). *Let T be a consistent theory in a language $\mathscr{L}$ of power α and let $\Sigma(x_1 \dots x_n)$ be a set of formulas of $\mathscr{L}$. If T α-omits Σ, then T has a model of power $\leqslant \alpha$ which omits Σ.*

The proof is like the proof of the omitting types theorem. An important problem is to find a useful sufficient condition for a theory in an uncountable language to have a model which omits a countable set of formulas. The α-omitting types theorem is of no help here since a countable set of formulas is never α-omitted when $\alpha > \omega$.

We now turn to the interpolation theorems of Craig and Lyndon.

THEOREM 2.2.20 (Craig Interpolation Theorem). *Let φ, ψ be sentences such that $\varphi \models \psi$. Then there exists a sentence θ such that:*

(i). *$\varphi \models \theta$ and $\theta \models \psi$.*

(ii). *Every relation, function or constant symbol (excluding identity) which occurs in θ also occurs in both φ and ψ.*

The sentence θ will be called a *Craig interpolant* of φ, ψ. The identity symbol is allowed to occur in θ. The following example shows why this is necessary.

EXAMPLE 2.2.21. In each of the following, φ and ψ are sentences such that the identity symbol occurs in at most one of them, and $\varphi \models \psi$; however, φ, ψ have no Craig interpolant in which the identity symbol does not occur:

(i). φ is $(\exists x)(P(x) \wedge \neg P(x))$, ψ is $(\exists x)Q(x)$;

(ii). φ is $(\exists x)Q(x)$, ψ is $(\exists x)(P(x) \vee \neg P(x))$;

(iii). φ is $(\forall xy)(x \equiv y)$, ψ is $(\forall xy)(P(x) \leftrightarrow P(y))$.

We shall see in an exercise, however, that in the Craig interpolation theorem, if the identity symbol occurs in neither φ nor ψ, and if not $\models \neg \varphi$ and not $\models \psi$, then φ and ψ have a Craig interpolant in which the identity symbol does not occur.

PROOF OF THEOREM 2.2.20. We assume that there is no Craig interpolant θ of φ and ψ, and prove that it is not the case that $\varphi \models \psi$. To do this we construct a model of $\varphi \wedge \neg \psi$. We may assume without loss of generality that $\mathscr{L}$ is the language of all symbols which occur in either φ or ψ or both. Let $\mathscr{L}_1$ be the language of all symbols of φ, $\mathscr{L}_2$ the language of all symbols of ψ, and $\mathscr{L}_0$ the language of all symbols occurring in both φ and ψ. Thus

$$\mathscr{L}_1 \cap \mathscr{L}_2 = \mathscr{L}_0, \qquad \mathscr{L}_1 \cup \mathscr{L}_2 = \mathscr{L}.$$

Form an expansion $\mathscr{L}'$ of $\mathscr{L}$ by adding a countable set C of new constant symbols and let

$$\mathscr{L}_0' = \mathscr{L}_0 \cup C, \qquad \mathscr{L}_1' = \mathscr{L}_1 \cup C, \qquad \mathscr{L}_2' = \mathscr{L}_2 \cup C.$$

The proof will resemble the proofs of the completeness and omitting types theorems, but the notion of a consistent theory will be replaced by the more general notion of an inseparable pair of theories.

Consider a pair of theories T in $\mathscr{L}_1'$ and U in $\mathscr{L}_2'$. A sentence θ of $\mathscr{L}_0'$ is said to *separate* T and U iff

$$T \models \theta \quad \text{and} \quad U \models \neg \theta.$$

T and U are said to be *inseparable* iff no sentence θ of $\mathscr{L}_0'$ separates them. To begin with, we see that

(1) $\{\varphi\}$ and $\{\neg \psi\}$ are inseparable.

For, if $\theta(c_1 \dots c_n)$ separates $\{\varphi\}$ and $\{\neg\psi\}$ and $u_1 \dots, u_n$ are variables not occurring in $\theta(c_1 \dots c_n)$, then $(\forall u_1 \dots u_n)\theta(u_1 \dots u_n)$ is a Craig interpolant of φ and ψ, contrary to our assumption.

Now let

$$\varphi_0, \varphi_1, \varphi_2, \dots, \psi_0, \psi_1, \psi_2, \dots$$

be enumerations of all sentences of $\mathscr{L}_1'$ and of $\mathscr{L}_2'$, respectively. We shall construct two increasing sequences of theories,

$$\{\varphi\} = T_0 \subset T_1 \subset T_2 \subset \dots,$$
$$\{\neg\psi\} = U_0 \subset U_1 \subset U_2 \subset \dots$$

in $\mathscr{L}_1'$ and $\mathscr{L}_2'$, respectively, such that:

(2). T_m and U_m are inseparable finite sets of sentences.

(3). If $T_m \cup \{\varphi_m\}$ and U_m are inseparable, then $\varphi_m \in T_{m+1}$.
If T_{m+1} and $U_m \cup \{\psi_m\}$ are inseparable, then $\psi_m \in U_{m+1}$.

(4). If $\varphi_m = (\exists x)\sigma(x)$ and $\varphi_m \in T_{m+1}$, then $\sigma(c) \in T_{m+1}$ for some $c \in C$.
If $\psi_m = (\exists x)\delta(x)$ and $\psi_m \in U_{m+1}$, then $\delta(d) \in U_{m+1}$ for some $d \in C$.

Given T_m and U_m, the theories T_{m+1} and then U_{m+1} are constructed in the obvious way. For (4), use constants c and d which do not occur in T_m, U_m, φ_m or ψ_m. Then inseparability will be preserved. Let

$$T_\omega = \bigcup_{m<\omega} T_m, \qquad U_\omega = \bigcup_{m<\omega} U_m.$$

Then T_ω and U_ω are inseparable. It follows that T_ω and U_ω are each consistent. We must show that $T_\omega \cup U_\omega$ is consistent. We show first that:

(5). T_ω is a maximal consistent theory in $\mathscr{L}_1'$, and U_ω is a maximal consistent theory in $\mathscr{L}_2'$.

To show this, suppose $\varphi_m \notin T_\omega$ and $(\neg\varphi_m) \notin T_\omega$. Since $T_m \cup \{\varphi_m\}$ is separable from U_m, there exists $\theta \in \mathscr{L}_0'$ such that

$$T_\omega \models \varphi_m \to \theta, \qquad U_\omega \models \neg\theta.$$

We see by the same argument that there exists $\theta' \in \mathscr{L}_0'$ such that

$$T_\omega \models \neg\varphi_m \to \theta', \qquad U_\omega \models \neg\theta'.$$

But then

$$T_\omega \models \theta \vee \theta', \qquad U_\omega \models \neg(\theta \vee \theta'),$$

contradicting the inseparability of T_ω and U_ω. This shows that T_ω is maximal consistent in $\mathscr{L}_1'$. The maximality of U_ω is similar.

Our next observation is:

(6). $T_\omega \cap U_\omega$ is a maximal consistent theory in $\mathscr{L}_0'$.

To prove (6), let σ be a sentence of $\mathscr{L}_0'$. By (5), either $\sigma \in T_\omega$ or $(\neg\sigma) \in T_\omega$, and either $\sigma \in U_\omega$ or $(\neg\sigma) \in U_\omega$. By inseparability, we cannot have $\sigma \in T_\omega$ and $(\neg\sigma) \in U_\omega$, or vice versa. Therefore either $T_\omega \cap U_\omega \models \sigma$ or $T_\omega \cap U_\omega \models \neg\sigma$.

We are now ready to construct a model. Let $\mathfrak{B}_1' = (\mathfrak{B}_1, b_0, b_1, \dots)$ be a model of T_ω. Using (4) and (5), we see that the submodel $\mathfrak{A}_1' = (\mathfrak{A}_1, b_0, b_1, \dots)$ with universe $A_1 = \{b_0, b_1, \dots\}$ is also a model of T_ω. Similarly, U_ω has a model $\mathfrak{A}_2' = (\mathfrak{A}_2, d_0, d_1, \dots)$ with universe $A_2 = \{d_0, d_1, \dots\}$. By (6), the $\mathscr{L}_0'$ reducts of $\mathfrak{A}_1'$ and $\mathfrak{A}_2'$ are isomorphic, with b_n corresponding to d_n. We may therefore take $b_n = d_n$ for each n, whence $\mathfrak{A}_1$ and $\mathfrak{A}_2$ have the same $\mathscr{L}_0$ reduct. Let $\mathfrak{A}$ be the model for $\mathscr{L}$ with $\mathscr{L}_1$ reduct $\mathfrak{A}_1$ and $\mathscr{L}_2$ reduct $\mathfrak{A}_2$. Since $\varphi \in T_\omega$ and $(\neg\psi) \in U_\omega$, $\mathfrak{A}$ is a model of $\varphi \wedge \neg\psi$. ⊣

We give two applications of the Craig interpolation theorem. The first application deals with ways of defining a relation. Let P and P' be two new n-placed relation symbols, not in the language $\mathscr{L}$. Let $\Sigma(P)$ be a set of sentences of the language $\mathscr{L} \cup \{P\}$, and let $\Sigma(P')$ be the corresponding set of sentences of $\mathscr{L} \cup \{P'\}$ formed by replacing P everywhere by P'. We say that $\Sigma(P)$ *defines P implicitly* iff

$$\Sigma(P) \cup \Sigma(P') \models (\forall x_1 \ldots x_n)[P(x_1 \ldots x_n) \leftrightarrow P'(x_1 \ldots x_n)].$$

Equivalently, if $(\mathfrak{A}, R)$ and $(\mathfrak{A}, R')$ are models of $\Sigma(P)$, then $R = R'$. $\Sigma(P)$ is said to *define P explicitly* iff there exists a formula $\varphi(x_1 \ldots x_n)$ of $\mathscr{L}$ such that

$$\Sigma(P) \models (\forall x_1 \ldots x_n)[P(x_1 \ldots x_n) \leftrightarrow \varphi(x_1 \ldots x_n)].$$

It is obvious that, if $\Sigma(P)$ defines P explicitly, then $\Sigma(P)$ defines P implicitly. Thus, to show that $\Sigma(P)$ does not define P explicitly, it suffices to find two models $(\mathfrak{A}, R)$ and $(\mathfrak{A}, R')$ of $\Sigma(P)$, with the same reduct $\mathfrak{A}$ to $\mathscr{L}$, such that $R \neq R'$. This is a useful classical method known as Padoa's method. We now prove the converse of Padoa's method.

THEOREM 2.2.22 (Beth's Theorem). *$\Sigma(P)$ defines P implicitly if and only if $\Sigma(P)$ defines P explicitly.*

PROOF. We prove only the 'hard' direction. Suppose that $\Sigma(P)$ defines P implicitly. Add new constants $c_1, \ldots, c_n$ to $\mathscr{L}$. Then

$$\Sigma(P) \cup \Sigma(P') \models P(c_1 \ldots c_n) \rightarrow P'(c_1 \ldots c_n).$$

By the compactness theorem, there exist finite subsets $\Delta \subset \Sigma(P)$, $\Delta' \subset \Sigma(P')$ such that

$$\Delta \cup \Delta' \models P(c_1 \ldots c_n) \rightarrow P'(c_1 \ldots c_n).$$

Let $\psi(P)$ be the conjunction of all $\sigma(P) \in \Sigma(P)$ such that either $\sigma(P) \in \Delta$ or $\sigma(P') \in \Delta'$. Then

$$\psi(P) \wedge \psi(P') \models P(c_1 \ldots c_n) \rightarrow P'(c_1 \ldots c_n).$$

Rearranging to get all symbols P on one side and all symbols P' on the other,

$$\psi(P) \wedge P(c_1 \ldots c_n) \models \psi(P') \rightarrow P'(c_1 \ldots c_n).$$

Then, by the Craig interpolation theorem, there is a sentence $\theta(c_1 \ldots c_n)$ of $\mathscr{L} \cup \{c_1 \ldots c_n\}$ such that

(1) $$\psi(P) \wedge P(c_1 \ldots c_n) \models \theta(c_1 \ldots c_n),$$

(2) $$\theta(c_1 \ldots c_n) \models \psi(P') \rightarrow P'(c_1 \ldots c_n).$$

But any model $(\mathfrak{A}, R')$ for $\mathscr{L} \cup \{P', c_1, \ldots, c_n\}$ is also a model for $\mathscr{L} \cup \{P, c_1, \ldots, c_n\}$ when we interpret P by R'. Thus (2) implies

(3) $$\theta(c_1 \ldots c_n) \models \psi(P) \rightarrow P(c_1 \ldots c_n).$$

Now (1) and (3) yield

(4) $$\psi(P) \models P(c_1 \ldots c_n) \leftrightarrow \theta(c_1 \ldots c_n).$$

Since $c_1, \ldots, c_n$ do not occur in $\psi(P)$ (which is built from $\Sigma(P)$), we have

$$\psi(P) \models \forall x_1 \ldots x_n[P(x_1 \ldots x_n) \leftrightarrow \theta(x_1 \ldots x_n)],$$

where $x_1, \ldots, x_n$ are variables not occurring in $\theta(c_1 \ldots c_n)$. Therefore

$$\Sigma(P) \models \forall x_1 \ldots x_n[P(x_1 \ldots x_n) \leftrightarrow \theta(x_1 \ldots x_n)]. \dashv$$

THEOREM 2.2.23 (Robinson Consistency Theorem). *Let $\mathscr{L}_1$ and $\mathscr{L}_2$ be two languages and let $\mathscr{L} = \mathscr{L}_1 \cap \mathscr{L}_2$. Suppose T is a complete theory in $\mathscr{L}$, and $T_1 \supset T$, $T_2 \supset T$ are consistent theories in $\mathscr{L}_1$, $\mathscr{L}_2$, respectively. Then $T_1 \cup T_2$ is consistent in the language $\mathscr{L}_1 \cup \mathscr{L}_2$.*

PROOF. Suppose $T_1 \cup T_2$ is inconsistent. Then there exist finite subsets $\Sigma_1 \subset T_1$, $\Sigma_2 \subset T_2$ such that $\Sigma_1 \cup \Sigma_2$ is inconsistent. Let σ_1 be the conjunction of Σ_1 and σ_2 the conjunction of Σ_2. It follows that $\sigma_1 \models \neg \sigma_2$. By the Craig interpolation theorem, there is a sentence θ such that $\sigma_1 \models \theta$, $\theta \models \neg \sigma_2$, and every relation, function or constant symbol occurring in θ occurs in both σ_1 and σ_2. Consequently, θ is a sentence of $\mathscr{L}_1 \cap \mathscr{L}_2 = \mathscr{L}$. Now returning to T_1 and T_2, we find that $T_1 \models \theta$. Since T_1 is consistent, $T_1 \not\models \neg \theta$, so $T \not\models \neg \theta$. Moreover, $T_2 \models \neg \theta$, and, by the consistency of T_2, $T_2 \not\models \theta$; so $T \not\models \theta$. But this contradicts the hypothesis that T is a complete theory in $\mathscr{L}$. $\dashv$

The Lyndon interpolation theorem is an improvement of the Craig interpolation theorem, but it holds only for languages which have no function or constant symbols. In order to state it, we need the notions of a positive and a negative occurrence of a symbol in a formula.

In the following discussion we shall consider only formulas which are built up using the connectives $\wedge$, $\vee$, $\neg$, and the quantifiers $\forall$, $\exists$. We do not allow the connectives $\rightarrow$, $\leftrightarrow$.

[Strictly speaking, the language $\mathscr{L}$ was defined in Section 1.2 so that the only connectives are $\wedge$ and $\neg$, and the only quantifier is $\forall$. The other connectives and $\exists$ were introduced as abbreviations. Thus we now wish to avoid using the abbreviations $\rightarrow$, $\leftrightarrow$.]

We now shall consider more closely the ways in which a symbol can occur in a sentence. Let s be a symbol of $\mathscr{L}$, and let φ be a sentence of $\mathscr{L}$. Then s is said to occur *positively* in φ iff s has an occurrence in φ which is within the scope of an even number of negation symbols. The symbol s occurs *negatively* in φ iff s has an occurrence in φ which is within the scope of an odd number of negation symbols. Remembering that s may have several different occurrences in φ, we see that there are four possibilities:

s does not occur in φ;
s occurs positively in φ;
s occurs negatively in φ;
s occurs both positively and negatively in φ.

The reason we do not want to use the abbreviations $\rightarrow$ and $\leftrightarrow$ is that they contain 'hidden' negation symbols. For example, the sentence $P(c) \rightarrow Q(c)$ is an abbreviation of $\neg(P(c) \wedge \neg Q(c))$, so P occurs negatively but not positively in it, and the constant c occurs both positively and negatively in it.

On the other hand, the abbreviations

$$\varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi), \qquad (\exists x)\varphi = \neg(\forall x)\neg\varphi$$

will not cause any trouble in deciding whether a symbol s occurs positively or negatively, because they introduce exactly two 'hidden' negation symbols about φ, ψ, and two is an even number.

THEOREM 2.2.24 (Lyndon Interpolation Theorem). *Let φ, ψ be sentences of $\mathscr{L}$ such that $\varphi \models \psi$. Then there is a sentence θ of $\mathscr{L}$ such that:*

(i). *$\varphi \models \theta$ and $\theta \models \psi$.*

(ii). *Every relation symbol (excluding equality) which occurs positively in θ occurs positively in both φ and ψ.*

(iii). *Same as* (ii) *for 'negatively'.*

The following simple example shows that we cannot find an interpolant θ which satisfies (ii) and (iii) for constant symbols:

$$(\exists x)(x \equiv c \wedge \neg R(x)) \models \neg R(c).$$

Note that c is positive on the left, negative on the right, but must occur in any interpolant.

PROOF OF THEOREM 2.2.24. The proof is obtained by making only a very few changes in the proof of the Craig interpolation theorem. We begin by assuming that there is no sentence θ such that (i)–(iii) hold, and prove that $\varphi \wedge \neg\psi$ has a model. Form the expansion $\mathscr{L}' = \mathscr{L} \cup C$ as before.

A formula is said to be in *negation normal form* (nnf) iff it is built up from atomic formulas and their negations using $\wedge$, $\vee$, $\exists$, $\forall$. Every formula is equivalent to an nnf formula. We assume that φ and ψ are nnf formulas. Let σ^* denote the nnf of $\neg\sigma$.

This time, the notion of an inseparable pair of theories is defined as follows. Let Φ be the set of all nnf sentences σ of $\mathscr{L}'$ such that every relation symbol which occurs positively (or negatively) in σ also occurs positively (negatively) in φ. The set Ψ is defined similarly with respect to ψ. Let $\Psi^* = \{\sigma^* : \sigma \in \Psi\}$. Two theories $T \subset \Phi$ and $U \subset \Psi^*$ are said to be *inseparable* iff there is no sentence $\theta \in \Phi \cap \Psi$ such that $T \models \theta$ and $U \models \neg\theta$. Using this notion we can apply the construction given in the proof of the Craig interpolation theorem to obtain a model of $\varphi \wedge \psi^*$.

This time we enumerate the sets of sentences Φ and Ψ instead of the languages $\mathscr{L}_1'$ and $\mathscr{L}_2'$, and then construct T_n and U_n as before. Some changes are needed in the rest of the proof because the sets Φ and Ψ are not necessarily closed under negation. Instead of proving that T_ω is maximal consistent, show that if $\sigma \vee \theta \in T_\omega$ then either $\sigma \in T_\omega$ or $\theta \in T_\omega$, and similarly for U_ω. Then show that T_ω and U_ω have the same equations and inequalities, and that the set Δ for all atomic and negated atomic sentences in $T_\omega \cup U_\omega$ is consistent. Finally, let $\mathfrak{A}$ be a model of Δ whose universe is the set of all constants, and prove by induction on complexity of formulas that $\mathfrak{A}$ is a model of both T_ω and U_ω and therefore a model of $\varphi \wedge \psi^*$. ⊣

A suggestion for further reading: The book "Building Models by Games" by Hodges [1985] gives an interesting treatment of a wide variety of applications of the Henkin construction in model theory.

EXERCISES

2.2.1. Let T be a complete theory in a countable language, and let $\Gamma_1(x_1), \Gamma_2(x_2), \Gamma_3(x_3), \ldots$ be a countable set of sets of formulas such that each $\Gamma_n(x_n)$ is consistent with T. Prove that T has a countable model which realizes each set $\Gamma_n(x_n)$.

2.2.2. Let T be a complete theory. Show that T has a model $\mathfrak{A}$ such that every set of formulas $\Gamma(x_1, x_2, \ldots)$ which is consistent with T is realized in $\mathfrak{A}$.

2.2.3. Let $\mathfrak{A} = \langle A, \leqslant, +, \cdot, 0, 1\rangle$ be an ordered field. An element $a \in A$ is said to be finite iff there is an $n < \omega$ such that $-n \leqslant a \leqslant n$. Suppose that for any formula $\varphi(x)$, if $\mathfrak{A} \models (\exists x)\varphi(x)$, then there is a finite $a \in A$ such that $\mathfrak{A} \models \varphi[a]$. Show that $\mathfrak{A}$ is elementarily equivalent to an Archimedean-ordered field.

2.2.4. Let T be a theory in a countable $\mathscr{L}$ and let $\Sigma(x)$ and $\Delta(y)$ be two sets of formulas of $\mathscr{L}$ which are consistent with T. Suppose that for every formula $\varphi(x, y)$ of $\mathscr{L}$ there exists $\sigma(x) \in \Sigma(x)$ such that for all $\delta_1(y), \ldots, \delta_n(y) \in \Delta(y)$: if $\{\varphi, \delta_1, \ldots, \delta_n\}$ is consistent with T, then $\{\varphi, \delta_1, \ldots, \delta_n, \neg\sigma\}$ is consistent with T. Prove that T has a model realizing $\Delta(y)$ and omitting $\Sigma(x)$.

2.2.5. Let T be a complete theory in a countable language $\mathscr{L}$. Suppose that for each $n < \omega$, T has a model $\mathfrak{A}_n$ omitting the set of formulas $\Sigma_n(x)$. Prove that T has a model $\mathfrak{A}$ which omits each $\Sigma_n(x)$.

2.2.6. Let $\mathscr{L}$ be a countable language and let $\mathscr{L}' = \mathscr{L} \cup \{P_0, P_1, \ldots\}$ be a countable expansion of $\mathscr{L}$. Let T' be a maximal consistent theory in $\mathscr{L}'$ and $\Gamma(x)$ a set of formulas of $\mathscr{L}$. Suppose that for each n, the restriction of T' to $\mathscr{L} \cup \{P_0, P_1, \ldots, P_n\}$ has a model which omits $\Gamma(x)$. Prove that T' has a model omitting $\Gamma(x)$.

2.2.7. Prove that there is an ordinal $\alpha < \omega_1$ such that every formula φ of ω-logic which has a proof has a proof of length less than α.

2.2.8. Show that the compactness theorem fails for ω-logic.

2.2.9. Show that the Löwenheim–Skolem–Tarski theorem fails for models of T which omit Σ.

2.2.10*. A model $\mathfrak{B}$ of Peano arithmetic is said to be an *end extension* of $\mathfrak{A}$ iff $\mathfrak{B}$ is a proper extension of $\mathfrak{A}$ and, for all $b \in B$ and $a \in A$, if $b < a$, then $b \in A$. Prove that every countable model of Peano arithmetic has an end elementary extension.

2.2.11*. Prove the following *Restricted Omitting Types Theorem*. Let $\mathscr{L}$ be a countable language and let T be a consistent $\forall\exists$ theory in $\mathscr{L}$, that is, a theory whose axioms are sentences of the form

$$(\forall y_1 \ldots y_p)(\exists z_1 \ldots z_q)\varphi$$

where φ has no quantifiers. For each $n < \omega$, let $\Sigma_n(x_1 \ldots x_k)$ be a set of universal formulas of $\mathscr{L}$. Suppose that for each n and each existential formula $\theta(x_1 \ldots x_k)$ consistent with T, there is a formula $\sigma(x_1 \ldots x_k) \in \Sigma_n(x_1 \ldots x_k)$ such that $\theta \wedge \neg\sigma$ is consistent with T. Prove that T has a countable model which omits each $\Sigma_n(x_1 \ldots x_k)$.

[*Hint*: The proof is similar to that of the Extended Omitting Types Theorem.]

2.2.12*. Deduce the Craig interpolation theorem from the Robinson consistency theorem.

2.2.13. Let Σ, Γ be sets of sentences of $\mathscr{L}$ such that $\Sigma \cup \Gamma$ is inconsistent. Then there exists a sentence θ of $\mathscr{L}$ such that:

(i). $\Sigma \models \theta$ and $\Gamma \models \neg\theta$.

(ii). Every relation, function or constant symbol which occurs in θ occurs in some member of Σ and in some member of Γ.

2.2.14

(i). Show that the Robinson consistency theorem fails if T is not assumed to be complete.

(ii). Show that the Robinson consistency theorem holds if the hypothesis that T is complete is replaced by the hypothesis that T is consistent, and for $i = 1, 2$, T contains every consequence of T_i in $\mathscr{L}$.

2.2.15. Prove the Craig interpolation theorem for formulas $\varphi(x_1 \ldots x_n)$, $\psi(x_1 \ldots x_n)$. It can be deduced easily from the Craig interpolation theorem for sentences.

2.2.16. Assume $\mathscr{L}$ has no function or constant symbols. Suppose that a set of sentences $\Sigma(P)$ of $\mathscr{L} \cup \{P\}$ defines P implicitly. Then there is a formula $\varphi(x_1 \ldots x_n)$ of $\mathscr{L}$ such that:

(i). $\Sigma(P) \vdash P(x_1 \ldots x_n) \leftrightarrow \varphi(x_1 \ldots x_n)$.

(ii). Any symbol of $\mathscr{L}$ which occurs in φ occurs both positively and negatively in $\Sigma(P)$.

2.2.17. Let $\mathscr{L}'$ be an expansion of the language $\mathscr{L}$ and let P be an n-placed relation symbol in $\mathscr{L}' \setminus \mathscr{L}$. Let T be a theory in $\mathscr{L}'$. Suppose that for any model $\mathfrak{A}$ for $\mathscr{L}$ and any two expansions $\mathfrak{A}'$, $\mathfrak{A}''$ of $\mathfrak{A}$ to models of T, the relations of $\mathfrak{A}'$ and $\mathfrak{A}''$ corresponding to P are the same. Prove that there exists a formula $\theta(x_1 \ldots x_n)$ of $\mathscr{L}$ such that

$$T \vdash P(x_1 \ldots x_n) \leftrightarrow \theta(x_1 \ldots x_n).$$

2.2.18. Let $\mathscr{L}'$ be an expansion of $\mathscr{L}$ and let T' be a theory in $\mathscr{L}'$. Suppose that each model $\mathfrak{A}$ for $\mathscr{L}$ has at most one expansion to a model T'. Prove that there is a theory T in $\mathscr{L}$ such that the models of T are exactly the reducts of the models of T' to $\mathscr{L}$.

2.2.19*. Show that the Lyndon interpolation theorem remains true when we add the conclusion:

(iv). If φ is a universal sentence, then so is θ.

Alternatively, it holds when we add:

(iv′). If ψ is an existential sentence, then so is θ.

However, the theorem becomes false if we add both the extra conclusions (iv), (iv′) at the same time.

2.2.20. Show that the Craig and Lyndon interpolation theorems hold with the following additional conclusion:

(iv). If not $\models \neg\varphi$, not $\models \psi$, and the identity symbol occurs in neither φ nor ψ, then the identity symbol does not occur in θ.

2.2.21*. Show that there is a model $\mathfrak{A}$ of Peano arithmetic which has an infinite element x such that no $y < x$ realizes the same complete type as x in $\mathfrak{A}$.

2.2.22*. Show that Peano arithmetic has two models $\mathfrak{A}$ and $\mathfrak{B}$ such that $\langle A, + \rangle \cong \langle B, + \rangle$ but not $\mathfrak{A} \cong \mathfrak{B}$. [*Hint*: Use Beth's Theorem.]

2.2.23*. Let T be a complete theory in a countable language and let $\Gamma(x)$ be a type over T which is consistent with T and locally omitted by T. Prove that T has a model in which infinitely many elements realize $\Gamma(x)$.

2.2.24*. Let S be a set of fewer than 2^ω types $\Gamma(x)$ which are maximal consistent with T and locally omitted by T. Prove that T has a countable model which simultaneously omits each $\Gamma(x) \in S$. [*Hint*: Represent the Henkin construction by a binary tree.]

2.3. Countable models of complete theories

In this section, we assume that $\mathscr{L}$ is a countable language. We shall embark on a thorough study of countable models of a complete theory. This study will give insight into what can be expected in general. Our study will center on two kinds of countable models, the atomic models, which are 'small', and the countably saturated models, which are 'large'. We begin with the atomic models.

Consider a complete theory T in $\mathscr{L}$. A formula $\varphi(x_1 \ldots x_n)$ is said to be *complete* (*in* T) iff for every formula $\psi(x_1 \ldots x_n)$ exactly one of

$$T \models \varphi \rightarrow \psi, \qquad T \models \varphi \rightarrow \neg\psi$$

holds. A formula $\theta(x_1 \ldots x_n)$ is said to be *completable* (*in* T) iff there is a complete formula $\varphi(x_1 \ldots x_n)$ with $T \models \varphi \rightarrow \theta$. If $\theta(x_1 \ldots x_n)$ is not completable it is said to be *incompletable*.

A theory T is said to be *atomic* iff every formula of $\mathscr{L}$ which is consistent with T is completable in T. A model $\mathfrak{A}$ is said to be an *atomic model* iff every n-tuple $a_1, \ldots, a_n \in A$ satisfies a complete formula in $\mathrm{Th}(\mathfrak{A})$.

In this and the next chapter we shall frequently pause to illustrate our definitions with examples. We shall sometimes make assertions about the examples without proofs. These proofs usually involve a combination of standard algebraic results and the theorems in the first three chapters of this book. In Section 3.4, we return to the examples and supply proofs.

2.3.1. Examples

(1). Let T be a complete theory and let $c_0, c_1, c_2, \ldots$ be constant symbols of $\mathscr{L}$. Then any formula of $\mathscr{L}$ of the form

$$x_0 \equiv c_0 \wedge x_1 \equiv c_1 \wedge \ldots \wedge x_n \equiv c_n$$

is complete in T. If $\mathfrak{A}$ is a model of T such that every element of A is a constant, then $\mathfrak{A}$ is an atomic model.

(2). The standard model of number theory is an atomic model.

(3). Let T be the theory of real closed ordered fields. The ordered field of real algebraic numbers is the unique atomic model of T. For example, the ordered field of real numbers is not atomic.

(4). Every finite model is atomic.

(5). Every model of pure identity theory is atomic. This gives an example of uncountable atomic models.

(6). Every dense linear ordering without endpoints is atomic.

(7). The following theory T is a complete theory which has no completable formulas and no atomic models. The language $\mathscr{L}$ has unary relation symbols $P_0(x), P_1(x), \ldots$. The axioms of T are all sentences of the form

$$(\exists x)(P_{i_1}(x) \wedge \ldots \wedge P_{i_m}(x) \wedge \neg P_{j_1}(x) \wedge \ldots \wedge \neg P_{j_n}(x)),$$

where the $i_1, \ldots, i_m, j_1, \ldots, j_n$ are all distinct.

Our first theorem about atomic models is an application of the extended omitting types theorem.

THEOREM 2.3.2 (Existence Theorem for Atomic Models). *Let T be a complete theory. Then T has a countable atomic model if and only if T is atomic.*

PROOF. First assume that T has an atomic model $\mathfrak{A}$. Let $\varphi(x_1 \dots x_n)$ be consistent with T. Then, since T is complete,

$$T \models (\exists x_1 \dots x_n)\varphi(x_1 \dots x_n).$$

Let $a_1, \dots, a_n \in A$ satisfy φ, and let $\psi(x_1 \dots x_n)$ be a complete formula satisfied by $a_1, \dots, a_n$. Then we cannot have $T \models \psi \rightarrow \neg\varphi$, so we must have $T \models \psi \rightarrow \varphi$. Hence φ is completable and T is atomic.

Now assume T is atomic. For each $n < \omega$, let $\Gamma_n(x_1 \dots x_n)$ be the set of all negations of complete formulas $\psi(x_1 \dots x_n)$ in T. Then every formula $\varphi(x_1 \dots x_n)$ which is consistent with T is completable, and hence $\varphi \wedge \neg\gamma$ is consistent with T for some $\gamma \in \Gamma_n$. Therefore T locally omits each set $\Gamma_n(x_1 \dots x_n)$. By the extended omitting types theorem, T has a countable model $\mathfrak{A}$ which omits each Γ_n. Then each $a_1, \dots, a_n \in A$ satisfies a complete formula, whence $\mathfrak{A}$ is an atomic model. ⊣

Returning to our examples, we see that complete number theory and the theory of real closed ordered fields are atomic, because they have atomic models.

THEOREM 2.3.3 (Uniqueness Theorem for Atomic Models). *If $\mathfrak{A}$ and $\mathfrak{B}$ are countable atomic models and $\mathfrak{A} \equiv \mathfrak{B}$, then $\mathfrak{A} \cong \mathfrak{B}$.*

PROOF. If $\mathfrak{A}$ or $\mathfrak{B}$ is finite, then $\mathfrak{A} \cong \mathfrak{B}$ is trivial. Let $\mathfrak{A}$ and $\mathfrak{B}$ be infinite and well-order the sets A and B with order type ω. The proof will be our first example of a *back and forth construction*. We shall see many other proofs of this type later.

Let a_0 be the first element of A and let $\varphi_0(x_0)$ be a complete formula satisfied by a_0 in $\mathfrak{A}$. Since $\mathfrak{A} \models (\exists x_0)\varphi_0(x_0)$, $\mathfrak{B} \models (\exists x_0)\varphi_0(x_0)$. Thus we may choose $b_0 \in B$, which satisfies $\varphi_0(x_0)$. Now let b_1 be the first element of $B\setminus\{b_0\}$, and let $\varphi_1(x_0 x_1)$ be a complete formula satisfied by b_0, b_1 in $\mathfrak{B}$. Then both $\mathfrak{A}$ and $\mathfrak{B}$ satisfy

$$\forall x_0(\varphi_0(x_0) \rightarrow (\exists x_1)\varphi_1(x_0 x_1)),$$

because φ_0 is complete. Therefore there exists $a_1 \in A$ such that a_0, a_1 satisfy $\varphi_1(x_0 x_1)$. Next, let a_2 be the first element of $A \setminus \{a_0, a_1\}$, and so on. Going back and forth ω times, we obtain sequences

$$a_0, a_1, a_2, \dots, b_0, b_1, b_2, \dots.$$

By going back and forth we used up all of A and B, so

$$A = \{a_0, a_1, \dots\}, \qquad B = \{b_0, b_1, \dots\}.$$

Moreover, for each n the n-tuples $a_0, \dots, a_{n-1}$ and $b_0, \dots, b_{n-1}$ satisfy the same complete formula. It follows that the mapping $a_m \rightarrow b_m$ is an isomorphism of $\mathfrak{A}$ onto $\mathfrak{B}$. ⊣

Our third result on atomic models shows that they should be thought of as 'small' models of T. First, we need to define the notion of a prime model.

$\mathfrak{A}$ is said to be a *prime model* iff $\mathfrak{A}$ is elementarily embedded in every model of Th($\mathfrak{A}$). $\mathfrak{A}$ is said to be *countably prime* iff $\mathfrak{A}$ is elementarily embedded in every countable model of Th($\mathfrak{A}$).

THEOREM 2.3.4. *The following are equivalent:*

(i). *$\mathfrak{A}$ is a countable atomic model.*
(ii). *$\mathfrak{A}$ is a prime model.*
(iii). *$\mathfrak{A}$ is a countably prime model.*

PROOF. First assume that $\mathfrak{A}$ is a countable atomic model and let $T = $ Th($\mathfrak{A}$). The proof that $\mathfrak{A}$ is prime is one-half of the 'back and forth' construction. Let $A = \{a_0, a_1, a_2, \dots\}$ and let $\mathfrak{B}$ be any model of T. Let $\varphi_0(x_0)$ be a complete formula satisfied by a_0. Then $T \models (\exists x_0)\varphi_0$, so we may choose $b_0 \in B$ which satisfies $\varphi_0(x_0)$. Now let $\varphi_1(x_0 x_1)$ be a complete formula satisfied by a_0, a_1. Then $T \models \varphi_0(x_0) \rightarrow (\exists x_1)\varphi_1(x_0 x_1)$. Choose $b_1 \in B$ so that b_0, b_1 satisfies φ_1, and so forth. The function $a_m \rightarrow b_m$ is an elementary embedding of $\mathfrak{A}$ into $\mathfrak{B}$.

Now assume $\mathfrak{A}$ is prime. Then $\mathfrak{A}$ is elementarily embedded in every countable model of T, so $\mathfrak{A}$ is countably prime.

Assume $\mathfrak{A}$ is countably prime. Let $a_1, \dots, a_n \in A$ and let $\Gamma(x_1 \dots x_n)$ be the set of all formulas $\gamma(x_1 \dots x_n)$ of $\mathscr{L}$ satisfied by $a_1, \dots, a_n$. For any countable model $\mathfrak{B}$ of T, we have some elementary embedding $f: \mathfrak{A} \prec \mathfrak{B}$, whence $fa_1, \dots, fa_n$ satisfies $\Gamma(x_1 \dots x_n)$ in $\mathfrak{B}$. Thus Γ is realized in every countable model of T. By the omitting types theorem, Γ is locally realized by T. Thus there is a formula $\varphi(x_1 \dots x_n)$ consistent with T such that $T \models \varphi \rightarrow \gamma$ for all $\gamma \in \Gamma$. But, for each formula $\psi(x_1 \dots x_n)$, either $\psi \in \Gamma$ or $(\neg\psi) \in \Gamma$. Thus φ is complete in T. We cannot have $T \models \varphi \rightarrow \neg\varphi$, so $\varphi \in \Gamma$. Therefore $\varphi(x_1 \dots x_n)$ is a complete formula satisfied by $a_1, \dots, a_n$ in $\mathfrak{A}$, and $\mathfrak{A}$ is atomic. ⊣

We now turn to the study of 'large' countable models. Given a model $\mathfrak{A}$ and a subset $Y \subset A$, the expanded model $(\mathfrak{A}, a)_{a \in Y}$ will be denoted by $\mathfrak{A}_Y$, and its language by $\mathscr{L}_Y$.

A model $\mathfrak{A}$ is said to be *ω-saturated* iff for every finite set $Y \subset A$, every set of formulas $\Gamma(x)$ of $\mathscr{L}_Y$ consistent with $\mathrm{Th}(\mathfrak{A}_Y)$ is realized in $\mathfrak{A}_Y$. A model is said to be *countably saturated* iff it is countable and ω-saturated.

To gain some intuition, we shall list some examples of countably saturated models. Note that if $\mathfrak{A}$ is ω-saturated, then so is $\mathfrak{A}_Y$ for every finite subset $Y \subset A$.

2.3.5. Examples

(1). Every countable infinite model of pure identity theory is countably saturated.

(2). The ordering of the rational numbers is countably saturated.

(3). Let T be a theory in the language with only the constant symbols $c_0, c_1, \ldots$, and axioms $c_i \not\equiv c_j$, $i < j < \omega$. There are countably many countable models of T up to isomorphism; for each $\alpha \leqslant \omega$, there is a model with exactly α elements which are not constants. The model with zero nonconstants is the atomic model. The model with ω nonconstants is the countably saturated model.

(4). Let T be the theory of algebraically closed fields of characteristic zero. Again there are countably many countable models; for each $\alpha \leqslant \omega$, there is a model of transcendence degree α over the rationals. The model of degree zero, i.e. the field of algebraic numbers, is the atomic model of T. The model of transcendence degree ω is the countably saturated model.

(5). Every finite model is countably saturated.

We need some additional notation for sets of formulas. Remember that a type in the variables $x_1, \ldots, x_n$ is a maximal consistent set $\Gamma(x_1 \ldots x_n)$ of formulas. The set T' of sentences which belong to Γ is a maximal consistent theory; we call T' the *theory of* Γ. If $T \subseteq \Gamma$, Γ is called a *type of* T. Given a model $\mathfrak{A}$ of T and an n-tuple $a_1, \ldots, a_n \in A$, the set of all formulas $\gamma(x_1 \ldots x_n)$ of $\mathscr{L}$ satisfied by $a_1, \ldots, a_n$ is a type of T, called the *type of* $a_1, \ldots, a_n$. By a *type of* $\mathfrak{A}$ we mean a type of $\mathrm{Th}(\mathfrak{A})$.

Consider a set of formulas $\Sigma(x_1 \ldots x_n)$ of $\mathscr{L}$. A formula $\varphi(x_1 \ldots x_n)$ is said to be a *consequence* of Σ, in symbols $\Sigma \models \varphi$, iff for every model $\mathfrak{A}$ and every n-tuple $a_1, \ldots, a_n \in A$, if $a_1, \ldots, a_n$ satisfies Σ, then it satisfies φ. That is,

$$\mathfrak{A} \models \Sigma[a_1 \ldots a_n] \quad \text{implies} \quad \mathfrak{A} \models \varphi[a_1 \ldots a_n].$$

We let $\Sigma(c_1 \ldots c_n)$ denote the set of all consequences in $\mathscr{L} \cup \{c_1, \ldots, c_n\}$ of the set

$$\{\sigma(c_1 \ldots c_n) : \sigma(x_1 \ldots x_n) \in \Sigma\}.$$

The notation $\Sigma(c_1 \ldots c_m x_{m+1} \ldots x_n)$ is defined in a similar way.

Let $\mathscr{L}' = \mathscr{L} \cup \{c_1, \ldots, c_m\}$ be a finite simple expansion of $\mathscr{L}$. There is a natural one-to-one correspondence between the types $\Sigma(x_1 \ldots x_n)$ of $\mathscr{L}$ and the types $\Gamma(x_{m+1} \ldots x_n)$ of $\mathscr{L}'$. If $\Sigma(x_1 \ldots x_n)$ is a type of $\mathscr{L}$, then

$$\Sigma' = \Sigma(c_1 \ldots c_m x_{m+1} \ldots x_n)$$

is a type of $\mathscr{L}'$. On the other hand, if $\Gamma(x_{m+1} \ldots x_n)$ is a type of $\mathscr{L}'$, then

$$\Sigma(x_1 \ldots x_n) = \{\sigma(x_1 \ldots x_n) : \sigma(c_1 \ldots c_m x_{m+1} \ldots x_n) \in \Gamma\}$$

is the unique type of $\mathscr{L}$ such that $\Sigma' = \Gamma$. (We leave the verification of this as an exercise.)

One might wonder why we used only sets of formulas in one free variable in the definition of an ω-saturated model. At first sight, it may appear that we would obtain a stronger notion by considering sets of formulas with finitely many free variables. The next proposition shows that we do not obtain a stronger notion in this way.

Proposition 2.3.6. *Let $\mathfrak{A}$ be an ω-saturated model. Then for each finite $Y \subset A$, each set of formulas $\Gamma(x_1 \ldots x_n)$ of $\mathscr{L}_Y$ consistent with* $\mathrm{Th}(\mathfrak{A}_Y)$ *is realized in $\mathfrak{A}_Y$.*

Proof. We argue by induction on n. The result holds for $n = 1$ by definition. Assume the result for $n-1$ and let $\Gamma(x_1 \ldots x_n)$ be consistent with $\mathrm{Th}(\mathfrak{A}_Y)$. We may assume that Γ is closed under finite conjunctions. Let

$$\Gamma'(x_1 \ldots x_{n-1}) = \{(\exists x_n)\gamma(x_1 \ldots x_n) : \gamma \in \Gamma\}.$$

Then Γ' is consistent with $\mathrm{Th}(\mathfrak{A}_Y)$. By inductive hypothesis, there is an $(n-1)$-tuple $a_1, \ldots, a_{n-1}$ realizing Γ' in $\mathfrak{A}_Y$. Let $Y' = Y \cup \{a_1, \ldots, a_{n-1}\}$, Then Y' is still finite. Moreover, the set $\Gamma(c_1 \ldots c_{n-1} x_n)$ is consistent with $\mathrm{Th}(\mathfrak{A}_{Y'})$ because for each $\gamma_1, \ldots, \gamma_m \in \Gamma$, $(\exists x_n)(\gamma_1 \wedge \ldots \wedge \gamma_m) \in \Gamma'$. Since $\mathfrak{A}$ is ω-saturated, there exists $a_n \in A$ realizing $\Gamma(c_1 \ldots c_{n-1} x_n)$ in $\mathfrak{A}_{Y'}$. Then $a_1, \ldots, a_n$ realizes Γ in $\mathfrak{A}_Y$. ⊣

Our three theorems below on countably saturated models will closely parallel our three theorems for atomic models. We shall prove an existence theorem, a uniqueness theorem, and a theorem showing that countably saturated models are 'large'.

THEOREM 2.3.7 (Existence Theorem for Countably Saturated Models). *Let T be a complete theory. Then T has a countably saturated model if and only if for each $n < \omega$, T has only countably many types in n variables.*

PROOF. Suppose first that T has a countably saturated model $\mathfrak{A}$. By Proposition 2.3.6, every type of T in n variables is realized in $\mathfrak{A}$. But no n-tuple can realize two different types in n variables. Therefore T has only countably many types.

Now suppose that for each n, T has only countably many types in n variables. Add a countable set $C = \{c_1, c_2, \dots\}$ of new constant symbols to $\mathscr{L}$, forming $\mathscr{L}'$. For each finite subset

$$Y = \{d_1, \dots, d_n\} \subset C,$$

the types $\Gamma(x)$ of T in $\mathscr{L}_Y$ are in one-to-one correspondence with the types $\Sigma(x_1 \dots x_n x)$ of T in $\mathscr{L}$. Therefore T has only countably many types $\Gamma(x)$ in $\mathscr{L}_Y$. Also, there are only countably many finite subsets $Y \subset C$. Let

$$\Gamma_1(x), \Gamma_2(x), \dots$$

be an enumeration of all types of T in all expansions $\mathscr{L}_Y$, Y a finite subset of C. Let

$$\varphi_1, \varphi_2, \dots$$

be an enumeration of all sentences of $\mathscr{L}'$. We form an increasing sequence

$$T = T_0 \subset T_1 \subset T_2 \subset \dots$$

of theories of $\mathscr{L}'$ such that for each $m < \omega$:

(1). T_m is a consistent theory which contains only finitely many constants from C.

(2). Either $\varphi_m \in T_{m+1}$ or $(\neg \varphi_m) \in T_{m+1}$.

(3). If $\varphi_m = (\exists x)\psi(x)$ is in T_{m+1}, then $\psi(c) \in T_{m+1}$ for some $c \in C$.

(4). If $\Gamma_m(x)$ is consistent with T_{m+1}, then $\Gamma_m(d) \subset T_{m+1}$ for some $d \in C$.

The construction of T_m is straightforward. The union $T_\omega = \bigcup_{n<\omega} T_n$ is a maximal consistent theory in $\mathscr{L}'$. Using (3) we see that T_ω has a model $\mathfrak{A}' = (\mathfrak{A}, a_1, a_2, \dots)$ such that $A = \{a_1, a_2, \dots\}$. Thus $\mathfrak{A}$ is a countable model of T.

It remains to prove that $\mathfrak{A}$ is ω-saturated. Let $Y \subset A$ be finite and let $\Sigma(x)$ be consistent with $\mathrm{Th}(\mathfrak{A}_Y)$. Extend $\Sigma(x)$ to a type $\Gamma(x)$ in $\mathrm{Th}(\mathfrak{A}_Y)$.

For some m, $\Gamma(x) = \Gamma_m(x)$. $\Gamma_m(x)$ is consistent with T_ω and hence with T_{m+1}. Then by (4), $\Gamma_m(c_i) \subset T_{m+1}$ for some $c_i \in C$, and it follows that a_i realizes $\Gamma(x)$ in $\mathfrak{A}_Y$. ⊣

COROLLARY 2.3.8. *If T is a complete theory with only countably many nonisomorphic countable models, then T has a countably saturated model.*

PROOF. Each type of T is realized in some countable model of T, and each countable model realizes only countably many types. Therefore T has countably many types. ⊣

THEOREM 2.3.9 (Uniqueness Theorem for Countably Saturated Models). *If $\mathfrak{A}$ and $\mathfrak{B}$ are countably saturated models and $\mathfrak{A} \equiv \mathfrak{B}$, then $\mathfrak{A}$ is isomorphic to $\mathfrak{B}$.*

PROOF. The proof uses a back and forth construction which closely parallels the proof of the uniqueness theorem for atomic models. The only difference is that instead of working with complete formulas we work with types. Using countable saturation of $\mathfrak{A}$ and $\mathfrak{B}$, we obtain two sequences

$$a_0, a_1, \dots, \qquad b_0, b_1, \dots,$$

such that

$$A = \{a_0, a_1, \dots\}, \qquad B = \{b_0, b_1, \dots\},$$

and, for each n, a_n realizes the same type in $(\mathfrak{A}, a_0, \dots, a_{n-1})$ as b_n realizes in $(\mathfrak{B}, b_0, \dots, b_{n-1})$. Then

$$(\mathfrak{A}, a_0, a_1, \dots) \equiv (\mathfrak{B}, b_0, b_1, \dots),$$

whence $\mathfrak{A} \cong \mathfrak{B}$ by the mapping $a_n \to b_n$. ⊣

The 'dual' of a prime model is a countably universal model. A model $\mathfrak{A}$ is said to be *countably universal* iff $\mathfrak{A}$ is countable and every countable model $\mathfrak{B} \equiv \mathfrak{A}$ is elementarily embedded in $\mathfrak{A}$. The next theorem shows that countably saturated models are 'large'.

THEOREM 2.3.10. *Every countably saturated model is countably universal.*

PROOF. Let $\mathfrak{B}$ be a countable model and $\mathfrak{A}$ a countably saturated model, $\mathfrak{A} \equiv \mathfrak{B}$. Let $B = \{b_0, b_1, \dots\}$. Using one half of the back and forth construction and the saturation of $\mathfrak{A}$, we obtain a sequence $a_0, a_1, a_2, \dots$ in A such that

$$(\mathfrak{B}, b_0, b_1, \dots) \equiv (\mathfrak{A}, a_0, a_1, \dots).$$

Then the mapping $b_n \to a_n$ is an elementary embedding of $\mathfrak{B}$ into $\mathfrak{A}$. ⊣