STUDIES IN LOGIC

AND

THE FOUNDATIONS OF MATHEMATICS

VOLUME 73

Editors

ELSEVIER
AMSTERDAM • LAUSANNE• NEW YORK • OXFORD • SHANNON • SINGAPORE • TOKYO

MODEL THEORY

C. C. CHANG
University of California, Los Angeles

and

H. J. KEISLER
University of Wisconsin, Madison

ELSEVIER
AMSTERDAM • LAUSANNE• NEW YORK • OXFORD • SHANNON • SINGAPORE • TOKYO

MODEL THEORY

ELSEVIER SCIENCE B.V.
Sara Burgerhartstraat 25
P.O. Box 211, 1000 AE Amsterdam, The Netherlands

First edition: 1973
Second edition: 1977
Second impression: 1978
Third edition: 1990
Second impression: 1991
Third impression: 1992
Fourth impression: 1994
Fifth impression: 1998

ISBN 0 444 88054 2

This book is printed on acid-free paper.

Printed in The Netherlands

PREFACE

Model theory is the branch of mathematical logic which deals with the connection between a formal language and its interpretations, or models. In this book we shall present the model theory of first order predicate logic, which is the simplest language that has applications to the main body of mathematics. Most of the techniques in model theory were originally developed and are still best explained in terms of first order logic.

The early pioneers in the development of model theory were Löwenheim (1915), Skolem (1920), Gödel (1930), Tarski (1931), and Malcev (1936). The subject became a separate branch of mathematical logic with the work of Henkin, Robinson, and Tarski in the late 1940's and early 1950's. Since that time it has been an active area of research.

Looking over the subject as it stands today, we feel that it can best be analyzed on the basis of a few general methods of constructing models. While the methods in their pure form are quite simple, they can be iterated and combined in a great variety of ways to yield practically all the deeper results of the theory. For this reason we have organized the book on the following plan. As a rule, we introduce a method in the first section of a chapter and then give some applications of it in the remaining sections. The basic methods of constructing models are: Constants (Section 2.1), Elementary chains (Section 3.1), Skolem functions (Section 3.3), Indiscernibles (Section 3.3), Ultraproducts (Section 4.1), and Special models (Section 5.1). In the last two chapters, 6 and 7, we present some more advanced topics which combine several of these methods. We believe that this book covers most of first-order model theory and many of its applications to algebra and set theory.

Up to now no book of this sort has been written. This has made it difficult for students and outsiders to learn about large areas of the subject. It has been necessary for them to chase down an almost unlimited number of widely scattered articles, some of which are hard to read. We do not claim to have compiled all the results in first-order model theory, but we have tried to include the important results which are indispensable for further work in this area. In addition we have included some of the more recent results which are stimulating present and probably future research. In this category are the Keisler-Shelah isomorphism theorem, the Morley categoricity theorem, the work of Ax-Kochen and Ershov in field theory, and the results of Rowbottom, Gaifman and Silver on large cardinals and the constructible universe.

First-order model theory is a prerequisite for the other types of model theory and such applications as nonstandard analysis. Other logics whose model theories have been investigated are infinitary logic, logic with additional quantifiers, many-valued logic, many sorted logic, intuitionistic logic, modal logic, second-order logic. In recent years model theory for infinitary logic has made rapid progress. Model theory for second-order logic is largely beyond present methods but has a great deal of potential importance. We hope that the availability of this book will contribute to future research in all kinds of model theory and to the discovery of more applications.

This book grew out of a number of graduate courses in model theory that we have taught at UCLA and Wisconsin. The idea of writing a textbook of this sort arose in 1963 as we were completing our earlier monograph, *Continuous Model Theory*. Some lecture notes by Keisler in 1963–64 were tried out from time to time and the present form of the book gradually evolved from them. The actual writing of the book began in early 1965. In the intervening period as the book took shape it was tested in classes, expanded in scope, and almost completely rewritten during the logic year 1967–68 at UCLA. Major changes were again made in 1971–72.

We owe a debt to the many mathematicians whose work forms the subject of this book. A tribute is due to Alfred Tarski who was the motivating and influencing force in the shaping of the theory. On a more personal level, we both received our Ph. D. degrees under his direction at the University of California, Berkeley. Space does not permit us to list the names of all the colleagues and students who at various times have read or used our manuscripts and have made many constructive suggestions and criticisms.

For the amusement of all those who gave us help, we dedicate our book to all model theorists who have never dedicated a book to themselves.

We have been supported during the writing by the Departments of Math-

ematics at the University of California, Los Angeles, and the University of Wisconsin, Madison, by the National Science Foundation under several research grants, and by a Fullbright grant to Chang in 1966–67 and a Sloan Fellowship to Keisler in 1966–67, 1968–69.

Invaluable assistance in the proof reading and preparation of the manuscript was rendered by Jerry Gold. Perry Smith has spent many hours helping us with the page proofs. We are grateful to Sister Kathleen Sullivan for preparing the index. We wish to thank Mrs. Gerry Formanack for her excellent typing of the manuscript.

University of California, Los Angeles — C. C. CHANG
University of Wisconsin, Madison — H. J. KEISLER

April 1973

PREFACE TO THE SECOND EDITION

The field of model theory has developed rapidly since the publication of the first edition of this book in 1973. There is an up-to-date survey of the subject and extensive references to the literature in the *Handbook of Mathematical Logic*.

Only minor changes have been made from the first edition of this book. We have added a few pages at the end of Appendix B discussing the current state of the list of open problems. A few of the problems have been completely solved, and partial results have been obtained on several others.

Throughout the book, errors, misprints, and ambiguities have been corrected. We are grateful to the many colleagues who have pointed out errors and offered suggestions. We especially wish to thank S. C. Kleene, who suggested over one hundred corrections after teaching a course from the book this spring.

University of California, Los Angeles — C. C. CHANG
University of Wisconsin, Madison — H. J. KEISLER

September 1976

PREFACE TO THE THIRD EDITION

It has been thirteen years since the Second Edition of this book was written, and as one would expect, the subject of model theory has changed radically. Model theory is now dominated by new areas which were in their infancy in 1976 and have blossomed into thriving fields in their own right. Among these fields are classification (or stability) theory, nonstandard analysis, model-theoretic algebra, recursive model theory, abstract model theory, and model theories for a host of nonfirst order logics. Model-theoretic methods have also had a major impact on set theory, recursion theory, and proof theory.

In spite of the changes in the field, this book still serves well as a beginning graduate textbook and reference work. Classical first order model theory as developed here remains a prerequisite for all of the newer branches of model theory, and many newer books have relied on this book for the necessary background.

In preparing this Third Edition, we have been careful to preserve the usefulness of the book as a first textbook in model theory. We have made no attempt to cover the whole field, but have added new topics which now belong in a first graduate course. Four new sections have been added. These sections have been placed at the end of the original chapters to minimize changes in the numbering of results. Throughout the book, new exercises have been added, usually at the end of the original exercise lists. We made a number of updates, improvements, and corrections in the main text, have updated the appendix on the current status of the open problems, and have added a list of additional references.

The new Section 2.4 introduces recursively saturated models, which have led to the simplification of many arguments in model theory by replacing large saturated or special models by countable models. As an illustration, we have replaced the proof of the Vaught two-cardinal theorem in Section 3.2 by a simpler proof using recursively saturated models.

The new Section 2.5 presents Lindström's celebrated characterization of first order logic. This result has gained importance as the launching point for the subject of abstract model theory.

Because of the growing importance of model-theoretic algebra, our treatment of model completeness which had been in Section 3.1 was greatly expanded and moved to the new Section 3.5.

The new Section 4.4 on nonstandard universes was added to provide an interface which is needed to apply results from model theory to nonstandard analysis.

We wish to thank our many colleagues and students who have given us invaluable help and encouragement on this textbook. We have been supported by several National Science Foundation grants and the Vilas Trust Fund at the University of Wisconsin.

Madison, Wisconsin, 1989 H.J. KEISLER

HOW TO USE THIS BOOK AS A TEXT

This book is written at a level appropriate to first year graduate students in mathematics. The only prerequisite is some exposure to elementary logic including the notion of a formal proof. It would be helpful if the student has had undergraduate-level courses in set theory and modern algebra. All the set theory needed for the book is presented in the Appendix which the student can use to fill in any gaps in his knowledge. The first four chapters proceed at a leisurely pace. The last three chapters proceed more rapidly and require more sophistication on the part of the student.

There is ample material for a full-year graduate course in Model Theory, and there is enough flexibility so that a variety of shorter courses can be made up. Chapter 1 contains introductory material and from Chapter 2 on there is at least one interesting theorem in every section.

The core of the subject which must be in any model theory course is composed of Sections 1.1, 1.3, 1.4, 2.1, 3.1, 4.1. The sections which are next in priority are 1.2, 2.2, 3.3, 4.3, 5.1, 6.1.

To help the instructor to make up a course, we give below a table showing the dependence of the sections in the first five chapters. This table applies only to the text itself and not to the exercises, which may depend on any earlier section.

1.4:	1.3	5.1:	3.1
1.5:	1.4	5.2:	5.1
2.1:	1.4	5.3:	5.1
2.2:	2.1	5.4:	5.1
2.3:	2.2	5.5:	5.1
2.4:	2.3	6.1:	4.3, 5.1
2.5:	2.4	6.2:	4.1
3.1:	2.1	6.3:	5.5, 6.2
3.2:	2.4, 3.1	6.4:	4.2, 4.4
3.3:	2.2, 3.1	6.5:	6.4
3.4:	3.2, 3.3	7.1:	2.3, 3.2, 3.3, 5.1
3.5:	2.4, 3.1	7.3:	3.2, 3.3, 4.2
4.1:	3.1	7.4:	7.3
4.2:	4.1		
4.3:	4.1		
4.4:	4.3		

Any hereditary set in the above partial ordering of sections can be used as a course. A very short course could consist of the core Sections 1.1, 1.3, 1.4, 2.1, 3.1, and 4.1. A one quarter course might consist of the above core plus Sections 2.2, 2.3, 2.4, 3.2, and 3.3 or 3.5; this would give a fairly complete picture of countable models. An alternative one quarter course which emphasizes ultraproducts and saturated models would add to the core the Sections 4.3, 5.1, and either 4.2, 4.4, or 6.1. All of Chapters 1 through 4 plus Sections 5.1, 6.1, and 7.1 would make an appropriate one semester course.

The exercises range from extremely easy to impossibly difficult. Exercises of more than routine difficulty are indicated by a single star; a few of the more difficult ones have double stars. Quite often improvements of the basic theorems proved in the text are put in the exercises. Some exercises are major theorems in their own right and we have included them to broaden the coverage. In order to gain an understanding of the field the student should try to do at least a third of the exercises.

At the end of the book we have included a list of unsolved problems in classical model theory. We feel that the solution of any of them would be a substantial contribution and worthy of publication. Not all of the problems originated with us.

We have collected all the historical remarks on the results in the text, the exercises, and the open problems in a separate section entitled Historical

Notes. In all probability there will be some omissions and errors for which we apologize in advance. In many cases students can find suggestions for further study in these notes.

Two final remarks on typography. The word 'iff' is used in all definitions that require it and is to mean 'if and only if'. The end of each proof is indicated by the symbol ⊣, which is meant to suggest the reverse of the common yield sign of first-order logic.

A MAPPING FROM THE SECOND TO THE THIRD EDITION

Here is a table of results and exercises in the Second Edition which have new numbers in the Third Edition. Most of these changes are in the old part of Section 3.1 on model completeness, which has been replaced by the new Section 3.5, and in Section 3.2, which has been simplified using recursively saturated models. E stands for "Exercise".

OLD NUMBER	NEW NUMBER
3.1.7	3.5.1
3.1.8	3.5.3
3.1.9	3.5.11
3.1.10	3.1.7
3.1.11	3.1.8
3.1.12	3.5.8
3.1.13	3.1.9
3.1.14	3.5.10
3.1.15	3.1.10
3.1.16	3.1.11
3.2.7	2.4.3
3.2.8	2.4.2
3.2.9	3.2.8
3.2.10	E3.2.11
3.2.11	3.2.7
3.2.12 – 3.2.16	3.2.9 – 3.2.13

E2.1.25	E2.1.20
E2.1.16 – E2.1.19	E2.1.21 – E2.1.24
E3.1.2	E3.5.2
E3.1.11	E3.5.6
E3.1.12	3.5.9
E3.1.14	E3.5.7
E3.1.16	after 3.5.18
E3.1.17	3.5.18
E3.1.18	E3.5.9
E3.2.5	2.4.3
E3.2.6	E2.4.6
E3.2.7	E2.4.8
E3.2.11	E2.4.7

CONTENTS

Chapter 1. Introduction 1
  1.1. What is model theory? 1
  1.2. Model theory for sentential logic 4
  1.3. Languages, models and satisfaction 18
  1.4. Theories and examples of theories 36
  1.5. Elimination of quantifiers 49

Chapter 2. Models constructed from constants 61
  2.1. Completeness and compactness 61
  2.2. Refinements of the method. Omitting types and interpolation theorems 77
  2.3. Countable models of complete theories 96
  2.4. Recursively saturated models 109
  2.5. Lindström's characterization of first order logic . . . 127

Chapter 3. Further model-theoretic constructions 136
  3.1. Elementary extensions and elementary chains . . . 136
  3.2. Applications of elementary chains 147
  3.3. Skolem functions and indiscernibles 163
  3.4. Some examples 178
  3.5. Model completeness 186

Chapter 4. Ultraproducts 211
  4.1. The fundamental theorem 211
  4.2. Measurable cardinals 227
  4.3. Regular ultrapowers 248
  4.4. Nonstandard universes 262

Chapter 5. Saturated and special models . . . 292
5.1. Saturated and special models . . . 292
5.2. Preservation theorems . . . 306
5.3. Applications of special models to the theory of definability . 323
5.4. Applications to field theory . . . 342
5.5. Application to Boolean algebras . . . 372

Chapter 6. More about ultraproducts and generalizations . . 384
6.1. Ultraproducts which are saturated . . . 384
6.2. Direct products, reduced products, and Horn sentences . 405
6.3. Direct products, reduced products, and Horn sentences (continued) . . . 420
6.4. Limit ultrapowers and complete extensions . . . 447
6.5. Iterated ultrapowers . . . 463

Chapter 7. Selected topics . . . 482
7.1. Categoricity in power . . . 482
7.2. An extension of Ramsey's theorem and applications; some two-cardinal theorems . . . 509
7.3. Models of large cardinality . . . 534
7.4. Large cardinals and the constructible universe . . . 558

Appendix A. Set theory . . . 579

Appendix B. Open problems in classical model theory . . . 597

Historical notes . . . 603

References . . . 623

Additional references . . . 634

Index of definitions . . . 641

Index of symbols . . . 649

CHAPTER 1

INTRODUCTION

1.1. What is model theory?

Model theory is the branch of mathematical logic which deals with the relation between a formal language and its interpretations, or models. We shall concentrate on the model theory of first-order predicate logic, which may be called 'classical model theory'.

Let us now take a short introductory tour of model theory. We begin with the models which are structures of the kind which arise in mathematics. For example, the cyclic group of order 5, the field of rational numbers, and the partially-ordered structure consisting of all sets of integers ordered by inclusion, are models of the kind we consider. At this point we could, if we wish, study our models at once without bringing the formal language into the picture. We then would be in the area known as universal algebra, which deals with homomorphisms, substructures, free structures, direct products, and the like. The line between universal algebra and model theory is sometimes fuzzy; our own usage is explained by the equation

$$\text{universal algebra} + \text{logic} = \text{model theory.}$$

To arrive at model theory, we set up our formal language, the first-order logic with identity. We specify a list of symbols and then give precise rules by which sentences can be built up from the symbols. The reason for setting up a formal language is that we wish to use the sentences to say things about the models. This is accomplished by giving a basic *truth definition*, which specifies for each pair consisting of a sentence and a model one of the truth values *true* or *false*. The truth definition is the bridge connecting the formal language with its interpretation by means of models. If the truth value 'true' goes with the sentence φ and model $\mathfrak{A}$, we say that

φ is *true in* $\mathfrak{A}$ and also that $\mathfrak{A}$ is a *model of* φ. Otherwise we say that φ is *false in* $\mathfrak{A}$ and that $\mathfrak{A}$ is not a model of φ. Moreover, we say that $\mathfrak{A}$ is a *model of a set* Σ of sentences iff $\mathfrak{A}$ is a model of each sentence in the set Σ.

What kinds of theorems are proved in model theory? We can already give a few examples. Perhaps the earliest theorem in model theory is Löwenheim's theorem (Löwenheim, 1915): If a sentence has an infinite model, then it has a countable model. Another classical result is the compactness theorem, due to Gödel (1930) and Malcev (1936): if each finite subset of a set Σ of sentences has a model, then the whole set Σ has a model. As a third example, we may state a more recent result, due to Morley (1965). Let us say that a set Σ of sentences is *categorical in power* α iff there is, up to isomorphism, exactly one model of Σ of power α. Morley's theorem states that, if Σ is categorical in one uncountable power, then Σ is categorical in every uncountable power.

These theorems are typical results of model theory. They say something negative about the 'power of expression' of first-order predicate logic. Thus Löwenheim's theorem shows that no consistent sentence can imply that a model is uncountable. Morley's theorem shows that first-order predicate logic cannot, as far as categoricity is concerned, tell the difference between one uncountable power and another. And the compactness theorem has been used to show that many interesting properties of models cannot be expressed by a set of first-order sentences – for instance, there is no set of sentences whose models are precisely all the finite models.

The three theorems we have stated also say something positive about the existence of models having certain properties. Indeed, in almost all of the deeper theorems in model theory the key to the proof is to construct the right kind of a model. For instance, look again at Löwenheim's theorem. To prove that theorem, we must begin with an uncountable model of a given sentence and construct from it a countable model of the sentence. Likewise, to prove the compactness theorem we must construct a single model in which each sentence of Σ is true. Even Morley's theorem depends vitally on the construction of a model. To prove it we begin with the assumption that Σ has two different models of one uncountable power and construct two different models of every other uncountable power.

There are a small number of extremely important ways in which models have been constructed. For example, for various purposes they can be constructed from individual constants, from functions, from Skolem terms, or from unions of chains. These constructions give the subject of model theory unity. To a large extent, we have organized this book according to these ways of constructing models.

Another point which gives model theory unity is the distinction between *syntax* and *semantics*. Syntax refers to the purely formal structure of the language – for instance, the length of a sentence and the collection of symbols occurring in a sentence, are syntactical properties. Semantics refers to the interpretation, or meaning, of the formal language – the truth or falsity of a sentence in a model is a semantical property. As we shall soon see, much of model theory deals with the interplay of syntactical and semantical ideas.

We now turn to a brief historical sketch. The mathematical world was forced to observe that a theory may have more than one model in the 19th century, when Bolyai and Lobachevsky developed non-Euclidean geometry, and Riemann constructed a model in which the parallel postulate was false but all the other axioms were true. Later in the 19th century, Frege formally developed the predicate logic, and Cantor developed the intuitive set theory in which our models live.

Model theory is a young subject. It was not clearly visible as a separate area of research in mathematics until the early 1950's. However, its historical roots go back to the older subjects of logic, universal algebra, and set theory – and some of the early work, such as Löwenheim's theorem, is now classified as model theory. Other important early developments which contributed to the theory are: the extension of Löwenheim's theorem by Skolem (1920) and Tarski; the completeness theorem of Gödel (1930) and its generalization by Malcev (1936); the characterization of definable sets of real numbers, the rigorous definition of the truth of a sentence in a model, and the study of relational systems by Tarski (1931, 1933, 1935a); the construction of a nonstandard model of number theory by Skolem (1934); and the study of equational classes initiated by Birkhoff (1935). Model theory owes a great deal to general methods which were originally developed for special purposes in older branches of mathematics. We shall come across many instances of this in our book; to mention just one, the important notion of a saturated model (Chapter 5) goes back to the η_α-structures in the theory of simple order, due to Hausdorff (1914). The subject grew rapidly after 1950, stimulated by the papers of Henkin (1949), Tarski (1950), and Robinson (1950). The phrase 'theory of models' is due to Tarski (1954). Today the literature in the subject is quite extensive. There is a rather complete bibliography in Addison, Henkin and Tarski (1965). In recent years, the theory of models has been applied to obtain significant results in other fields, notably set theory, algebra and analysis. However, until now only a tiny part of the potential strength of model theory has been used in such applications. It will be interesting to see what happens when (and if) the full strength is used.

1.2. Model theory for sentential logic

In our introduction, Section 1.1, we gave a general idea of the flavor of model theory, but we were not yet ready to give many details. We shall now come down to earth and give a rigorous treatment of model theory for a very simple formal language, sentential logic (also known as propositional calculus). We shall quickly develop this 'toy' model theory along lines parallel to the much deeper model theory for predicate logic. The basic ideas are the decision procedure via truth tables, due to Post (1921), and Lindenbaum's theorem with the compactness theorem which follows. This section will give a preview of what lies ahead in our book.

We are assuming (see Preface) that the reader is already thoroughly familiar with sentential, and even predicate, logic. Thus we shall feel free to proceed at a fairly rapid pace. Nevertheless, we shall start from scratch, in order to show what sentential logic looks like when it is developed in the spirit of model theory.

Classical sentential logic is designed to study a set $\mathcal{S}$ of simple statements, and the compound statements built up from them. At the most intuitive level, an intended interpretation of these statements is a 'possible world', in which each statement is either true or false. We wish to replace these intuitive interpretations by a collection of precise mathematical objects which we may use as our models. The first thing which comes to mind is a function F which associates with each simple statement S one of the truth values 'true' or 'false'. Stripping away the inessentials, we shall instead take a model to be a subset A of $\mathcal{S}$; the idea is that $S \in A$ indicates that the simple statement S is true, and $S \notin A$ indicates that the simple statement S is false.

1.2.1. By a *model A for $\mathcal{S}$* we simply mean a subset A of $\mathcal{S}$.

Thus the set of all models has the power $2^{|\mathcal{S}|}$. Several relations and operations between models come to mind; for example, $A \subset B$, $\mathcal{S} - A$, and the intersection $\bigcap_{i \in I} A_i$ of a set $\{A_i : i \in I\}$ of models. Two distinguished models are the empty set $\emptyset$ and the set $\mathcal{S}$ itself.

We now set up the sentential logic as a formal language. The symbols of our language are as follows:
connectives $\wedge$ (and), $\neg$ (not);
parentheses), (;
a nonempty set $\mathcal{S}$ of sentence symbols.
Intuitively, the sentence symbols stand for simple statements, and the connectives $\wedge$, $\neg$ stand for the words used to combine simple statements into compound statements. Formally, the *sentences* of $\mathcal{S}$ are defined as follows:

1.2.2.
(i). Every sentence symbol S is a sentence.
(ii). If φ is a sentence then $(\neg\varphi)$ is a sentence.
(iii). If φ, ψ are sentences, then $(\varphi \wedge \psi)$ is a sentence.
(iv). A finite sequence of symbols is a sentence only if it can be shown to be a sentence by a finite number of applications of (i)-(iii).

Our definition of sentence of $\mathcal{S}$ may be restated as a recursive definition based on the length of a finite sequence of symbols:

A single symbol is a sentence iff it is a sentential symbol; a sequence φ of symbols of length $n > 1$ is a sentence iff there are sentences ψ and θ of length less than n such that φ is either $(\neg\psi)$ or $(\psi \wedge \theta)$.

Alternatively, our definition may be restated in set-theoretical terms:

The set of all sentences of $\mathcal{S}$ is the least set Σ of finite sequences of symbols of $\mathcal{S}$ such that each sentence symbol S belongs to Σ and, whenever ψ, θ are in Σ, then $(\neg\psi)$, $(\psi \wedge \theta)$ belong to Σ.

No matter how we may think of sentences, the important thing is that *properties of sentences can only be established through an induction based on* 1.2.2. More precisely, to show that every sentence φ has a given property P, we must establish three things: (1) Every sentence symbol S has the property P; (2) if φ is $(\neg\psi)$ and ψ has the property P, then φ has the property P; (3) if φ is $(\psi \wedge \theta)$ and ψ, θ have the property P, then φ has the property P. (The reader may check his understanding of this point by proving through induction that every sentence φ has the same number of right parentheses as it has left parentheses.)

How many sentences of $\mathcal{S}$ are there? This depends on the number of sentence symbols $S \in \mathcal{S}$. Each sentence is a finite sequence of symbols. If the set $\mathcal{S}$ is finite or countable, then there are countably many sentences of $\mathcal{S}$. Of course, not every finite sequence of symbols is a sentence; for instance, $(S_0 \wedge (\neg S_5))$ is a sentence, but $\wedge \wedge)S_3$ and $S_0 \wedge \neg S_5$ are not. If the set $\mathscr{S}$ of sentence symbols has uncountable cardinal α, then the set of sentences of $\mathscr{S}$ also has power α.

Let us pause briefly to explain the role of the Greek letters φ, ψ, Σ, etc. In the above paragraphs we have used the lower case Greek letters $\varphi, \psi, \theta, \ldots$ as names for arbitrary finite sequences of symbols of $\mathscr{S}$. These letters were needed in order to write down the definition of a sentence. From now on, we shall be much more interested in sentences than in arbitrary finite sequences of symbols. We shall hereafter use the lower case Greek letters $\varphi, \psi, \theta, \ldots$ as names for arbitrary sentences of $\mathscr{S}$. The situation is similar to elementary arithmetic, where we study natural numbers 0, 1, 2, 3, ..., but much of the time we write down letters like $m, n, x, y, \ldots$ as names for arbitrary natural numbers. Just as in arithmetic where we write things like $m = x+y$, we shall now write, for example, $\varphi = (\psi \wedge \theta)$ to express the fact that φ and $(\psi \wedge \theta)$ are the same sentence. In the above paragraphs we also used capital Greek letters $\Sigma, \Gamma, \ldots$ as names for arbitrary sets of finite sequences of symbols of $\mathscr{S}$; hereafter we shall use the capital Greek letters as names for arbitrary sets of sentences of $\mathscr{S}$. The symbols $\varphi, \psi, \theta, \ldots, \Sigma, \Gamma, \ldots$ are *not* in our list of formal symbols of our language – they are merely informal symbols which we use to talk more easily about $\mathscr{S}$.

We shall introduce abbreviations to our language in the usual way, in order to make sentences more readable. The symbols $\vee$ (or), $\rightarrow$ (implies), and $\leftrightarrow$ (if and only if) are abbreviations defined as follows:

$$(\varphi \vee \psi) \quad \text{for} \quad (\neg((\neg \varphi) \wedge (\neg \psi))),$$
$$(\varphi \rightarrow \psi) \quad \text{for} \quad ((\neg \varphi) \vee \psi),$$
$$(\varphi \leftrightarrow \psi) \quad \text{for} \quad ((\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)).$$

Of course, $\vee$, $\rightarrow$ and $\leftrightarrow$ could just as well have been included in our list of symbols as three more connectives. However, there are certain advantages to keeping our list of symbols short. For instance, 1.2.2 and proofs by induction based on it are shorter this way. At the other extreme, we could have managed with only a single connective, whose English translation is 'neither ... nor ...'. We did not do this because 'neither ... nor ...' is a rather unnatural connective.

Another abbreviation which we shall adopt is to leave out unnecessary parentheses. For instance, we shall never bother to write outer parentheses in a sentence – thus $\neg S$ is our abbreviation for the sentence $(\neg S)$. We shall follow the commonly accepted usage in dropping other parentheses. Thus $\neg$ is considered more binding than $\wedge$ and $\vee$, which in turn are more binding than $\rightarrow$ and $\leftrightarrow$. For instance, $\neg \varphi \vee \psi \rightarrow \theta \wedge \varphi$ means $((\neg \varphi) \vee \psi) \rightarrow (\theta \wedge \varphi)$.

Hereafter we shall use the single symbol $\mathscr{S}$ to denote both the set of sentence symbols and the language built on these symbols. There is no fear of confusion in this double usage since the language is determined uniquely, modulo the connectives, by the sentence symbols.

We are now ready to build a bridge between the language $\mathscr{S}$ and its models, with the definition of the truth of a sentence in a model. We shall express the fact that a sentence φ is true in a model A succinctly by the special notation

$$A \models \varphi.$$

The relation $A \models \varphi$ is defined as follows:

1.2.3.

(i). If φ is a sentence symbol S, then $A \models \varphi$ holds if and only if $S \in A$.

(ii). If φ is $\psi \wedge \theta$, then $A \models \varphi$ if and only if both $A \models \psi$ and $A \models \theta$.

(iii). If φ is $\neg \psi$, then $A \models \varphi$ iff it is not the case that $A \models \psi$.

When $A \models \varphi$, we say that φ is *true in A*, or that φ *holds in A*, or that A is a *model of* φ. When it is not the case that $A \models \varphi$, we say that φ is *false in A*, or that φ *fails in A*. The above definition of the relation $A \models \varphi$ is an example of a recursive definition based on 1.2.2. The proof that the definition is unambiguous for each sentence φ is, of course, a proof by induction based on 1.2.2.

An especially important kind of sentence is a *valid sentence*. A sentence φ is called *valid*, in symbols $\models \varphi$, iff φ holds in all models for $\mathscr{S}$, that is, iff $A \models \varphi$ for all A. Some notions closely related to validity are mentioned in the exercises.

At first glance, it seems that for $\mathscr{S}$ infinite we have to examine uncountably many different infinite models A in order to find out whether a sentence φ is valid. This is because validity is a semantical notion, defined in terms of models. However, as the reader surely knows, there is a simple and uniform test by which we can find out in only finitely many steps whether or not a given sentence φ is valid.

This decision procedure for validity is based on a syntactical notion, the notion of a tautology. Let φ be a sentence such that all the sentence symbols which occur in φ are among the $n+1$ symbols $S_0, S_1, \ldots, S_n$. Let $a_0, a_1, \ldots, a_n$ be a sequence made up of the two letters t, f. We shall call such a sequence an *assignment*.

1.2.4. The *value* of a sentence φ for the assignment $a_0, \ldots, a_n$ is defined recursively as follows:

(i). If φ is the sentence symbol S_m, $m \leqslant n$, then the value of φ is a_m.
(ii). If φ is $\neg\psi$, then the value of φ is the opposite of the value of ψ.
(iii). If φ is $\psi \wedge \theta$, then the value of φ is t if the values of ψ and θ are both t, and otherwise the value of φ is f.

Note how similar Definitions 1.2.3 and 1.2.4 are. The only essential difference is that 1.2.3 involves an infinite model A, while 1.2.4 involves only a finite assignment $a_0, \ldots, a_n$.

1.2.5. Let φ be a sentence and let $S_0, \ldots, S_n$ be all the sentence symbols occurring in φ. φ is said to be a *tautology*, in symbols $\vdash \varphi$, iff φ has the value t for every assignment $a_0, \ldots, a_n$.

We shall use both of the symbols $\models$, $\vdash$ in many ways throughout this book. To keep things straight, remember this: $\models$ is used for semantical ideas, and $\vdash$ is used for syntactical ideas.

The value of a sentence φ for an assignment $a_0, \ldots, a_n$ may be very easily computed. We first find the values of the sentence symbols occurring in φ and then work our way through the smaller sentences used in building up the sentence φ. A table showing the value of φ for each possible assignment $a_0, \ldots, a_n$ is called a *truth table* of φ. We shall assume that truth tables are already quite familiar to the reader, and that he knows how to construct a truth table of a sentence. Truth tables provide a simple and purely mechanical procedure to determine whether a sentence φ is a tautology – simply write down the truth table for φ and check to see whether φ has the value t for every assignment.

PROPOSITION 1.2.6. *Suppose that all the sentence symbols occurring in φ are among $S_0, S_1, \ldots, S_n$. Then the value of φ for an assignment $a_0, a_1, \ldots, a_n, \ldots, a_{n+m}$ is the same as the value of φ for the assignment $a_0, a_1, \ldots, a_n$.*

We now prove the first of a series of theorems which state that a certain syntactical condition is equivalent to a semantical condition.

THEOREM 1.2.7 (Completeness Theorem). $\vdash \varphi$ *if and only if* $\models \varphi$; *in words, a sentence is a tautology if and only if it is valid.*

PROOF. Let φ be a sentence and let all the sentence symbols in φ be among $S_0, \ldots, S_n$. Consider an arbitrary model A. For $m = 0, 1, \ldots, n$, put $a_m = $ t if $S_m \in A$, and $a_m = $ f if $S_m \notin A$. This gives us an assignment $a_0, a_1, \ldots, a_n$. We claim:

(1) $A \models \varphi$ if and only if the value of φ for the assignment $a_0, a_1, \ldots, a_n$ is t.
This can be readily proved by induction. It is immediate if φ is a sentence symbol S_m. Assuming that (1) holds for $\varphi = \psi$ and for $\varphi = \theta$, we see at once that (1) holds for $\varphi = \neg\psi$ and $\varphi = \psi \wedge \theta$.

Now let $S_0, \ldots, S_n$ be all the sentence symbols occurring in φ. If φ is a tautology, then by (1), φ is valid. Since every assignment $a_0, a_1, \ldots, a_n$ can be obtained from some model A, it follows from (1) that, if φ is valid, then φ is a tautology. ⊣

Our decision procedure for $\vdash \varphi$ now can be used to decide whether φ is valid. Several times we shall have an occasion to use the fact that a particular sentence is a tautology, or is valid. We shall never take the trouble actually to give the proof that a sentence of $\mathscr{S}$ is valid, because the proof is always the same – we simply look at the truth table.

Let us now introduce the notion of a formal deduction in our logic $\mathscr{S}$. The *Rule of Detachment* (or *Modus Ponens*) states:

From ψ and $\psi \rightarrow \varphi$ infer φ.

We say that φ is *inferred from* ψ, θ *by detachment* iff θ is the sentence $\psi \rightarrow \varphi$.

Now consider a finite or infinite set Σ of $\mathscr{S}$.

A sentence φ is *deducible from* Σ, in symbols $\Sigma \vdash \varphi$, iff there is a finite sequence $\psi_0, \psi_1, \ldots, \psi_n$ of sentences such that $\varphi = \psi_n$ and each sentence ψ_m is either a tautology, belongs to Σ, or is inferred from two earlier sentences of the sequence by detachment. The sequence $\psi_0, \psi_1, \ldots, \psi_n$ is called a *deduction of* φ from Σ. Note that φ is deducible from the empty set of sentences if and only if φ is a tautology.

We shall say that Σ is *inconsistent* iff we have $\Sigma \vdash \varphi$ for all sentences φ. Otherwise, we say that Σ is *consistent*. Finally, we say that Σ is *maximal consistent* iff Σ is consistent, but the only consistent set of sentences which includes Σ is Σ itself. The proposition below contains facts which can be found in most elementary logic texts.

PROPOSITION 1.2.8.

(i). *If Σ is consistent and Γ is the set of all sentences deducible from Σ, then Γ is consistent.*

(ii). *If Σ is maximal consistent and $\Sigma \vdash \varphi$, then $\varphi \in \Sigma$.*
(iii). *Σ is inconsistent if and only if $\Sigma \vdash S \wedge \neg S$ (for any $S \in \mathcal{S}$).*
(iv) (Deduction Theorem). *If $\Sigma \cup \{\psi\} \vdash \varphi$, then $\Sigma \vdash \psi \to \varphi$.*

LEMMA 1.2.9 (Lindenbaum's Theorem). *Any consistent set Σ of sentences can be enlarged to a maximal consistent set Γ of sentences.*

PROOF. Let us arrange all the sentences of $\mathcal{S}$ in a list, $\varphi_0, \varphi_1, \varphi_2, \ldots, \varphi_\alpha, \ldots$. The order in which we list them is immaterial, as long as the list associates in a one–one fashion an ordinal number with each sentence. We shall form an increasing chain

$$\Sigma = \Sigma_0 \subset \Sigma_1 \subset \Sigma_2 \subset \ldots \subset \Sigma_\alpha \subset \ldots$$

of consistent sets of sentences. If $\Sigma \cup \{\varphi_0\}$ is consistent, define $\Sigma_1 = \Sigma \cup \{\varphi_0\}$. Otherwise define $\Sigma_1 = \Sigma$. At the αth stage, we define $\Sigma_{\alpha+1} = \Sigma_\alpha \cup \{\varphi_\alpha\}$ if $\Sigma_\alpha \cup \{\varphi_\alpha\}$ is consistent, and otherwise define $\Sigma_{\alpha+1} = \Sigma_\alpha$. At limit ordinals α take unions, $\Sigma_\alpha = \bigcup_{\beta<\alpha}\Sigma_\beta$. Now let Γ be the union of all the sets Σ_α.

We claim that Γ is consistent. Suppose not. Then there is a deduction $\psi_0, \psi_1, \ldots, \psi_p$ of the sentence $S \wedge \neg S$ from Γ (see Proposition 1.2.8). Let $\theta_1, \ldots, \theta_q$ be all the sentences in Γ which are used in this deduction. We may choose α so that all of $\theta_1, \ldots, \theta_q$ belong to Σ_α. But this means that Σ_α is inconsistent (again see Proposition 1.2.8), which is a contradiction.

Having shown that Γ is consistent, we next claim that Γ is maximal consistent. For suppose Δ is consistent and $\Gamma \subset \Delta$. Let $\varphi_\alpha \in \Delta$. Then $\Sigma_\alpha \cup \{\varphi_\alpha\}$ is consistent, and hence $\Sigma_{\alpha+1} = \Sigma_\alpha \cup \{\varphi_\alpha\}$. Thus $\varphi_\alpha \in \Gamma$, and hence $\Delta = \Gamma$. ⊣

LEMMA 1.2.10. *Suppose Γ is a maximal consistent set of sentences in $\mathcal{S}$. Then:*
(i). *For each sentence φ, exactly one of the sentences $\varphi, \neg\varphi$ belongs to Γ.*
(ii). *For each pair of sentences φ, ψ, $\varphi \wedge \psi$ belongs to Γ if and only if both φ and ψ belong to Γ.*

We leave the proof as an exercise.

Now consider a set Σ of sentences of $\mathcal{S}$. We shall say that A is a *model of* Σ, $A \models \Sigma$, iff every sentence $\varphi \in \Sigma$ is true in A. Σ is said to be *satisfiable* iff it has at least one model. We now prove the most important theorem of sentential logic, which is a criterion for a set Σ to be satisfiable.

THEOREM 1.2.11 (Extended Completeness Theorem). *A set Σ of sentences of $\mathcal{S}$ is consistent if and only if Σ is satisfiable.*

PROOF. Assume first that Σ is satisfiable, and let $A \models \Sigma$. We show that every sentence deducible from Σ holds in A. Let $\psi_0, \psi_1, \ldots, \psi_n$ be a deduction of ψ_n from Σ. Let $m \leqslant n$. If $\psi_m \in \Sigma$ or if ψ_m is a tautology, then ψ_m holds in A. If ψ_m is inferred from two sentences ψ_p, $\psi_p \to \psi_m$ which hold in A, then ψ_m must clearly hold in A. It follows by induction on m that each of the sentences $\psi_0, \psi_1, \ldots, \psi_n$ holds in A. Since $S \wedge \neg S$ does not hold in A, it is not deducible from Σ, so Σ is consistent.

Now assume that Σ is consistent. By Lindenbaum's theorem we enlarge Σ to a maximal consistent set Γ.

We now construct a model of Σ. Let A be the set of all sentence symbols $S \in \mathcal{S}$ such that $S \in \Gamma$. We show by induction that, for each sentence φ,

(1) $\qquad \varphi \in \Gamma$ if and only if $A \models \varphi$.

By definition, (1) holds when φ is a sentence symbol S_n. Lemma 1.2.10(i) guarantees that, if (1) holds when $\varphi = \psi$, then (1) holds when $\varphi = \neg\psi$. Lemma 1.2.10(ii) guarantees that, if (1) holds when $\varphi = \psi$ and when $\varphi = \theta$, then (1) holds when $\varphi = \psi \wedge \theta$. From (1) it follows that $A \models \Gamma$, and, since $\Sigma \subset \Gamma$, $A \models \Sigma$. ⊣

We can obtain a purely semantical corollary. Σ is said to be *finitely satisfiable* iff every finite subset of Σ is satisfiable.

COROLLARY 1.2.12 (Compactness Theorem). *If Σ is finitely satisfiable, then Σ is satisfiable.*

PROOF. Suppose Σ is not satisfiable. Then by the extended completeness theorem Σ is inconsistent. Hence, $\Sigma \vdash S \wedge \neg S$. In the deduction of the sentence $S \wedge \neg S$ from Σ only a finite set Σ_0 of sentences of Σ is used. It follows that $\Sigma_0 \vdash S \wedge \neg S$, so Σ_0 is inconsistent. Then Σ_0 is not satisfiable, so Σ is not finitely satisfiable. ⊣

Note that the converse of the compactness theorem is trivially true, i.e., every satisfiable set of sentences is finitely satisfiable.

We say that φ is a *consequence* of Σ, in symbols $\Sigma \models \varphi$, iff every model of Σ is a model of φ. The reader is asked to prove Exercises 1.2.3–1.2.6 as well as the following:

COROLLARY 1.2.13

(i). *$\Sigma \vdash \varphi$ if and only if $\Sigma \models \varphi$.*

(ii). *If $\Sigma \models \varphi$, then there is a finite subset Σ_0 of Σ such that $\Sigma_0 \models \varphi$.*

We shall conclude our model theory for sentential logic with a few applications of the compactness theorem. In these applications, the true spirit of model theory will appear, but at a very rudimentary level. Since we shall often wish to combine a finite set of sentences into a single sentence, we shall use expressions like

$$\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_n$$

and

$$\varphi_1 \vee \varphi_2 \vee \ldots \vee \varphi_n.$$

In these expressions the parentheses are assumed, for the sake of definiteness, to be associated to the right; for instance,

$$\varphi_1 \wedge \varphi_2 \wedge \varphi_3 = \varphi_1 \wedge (\varphi_2 \wedge \varphi_3).$$

First we introduce a bit more terminology. A set Γ of sentences is called a *theory*. A theory is said to be *closed* iff every consequence of Γ belongs to Γ. A set Δ of sentences is said to be a *set of axioms for* a theory Γ iff Γ and Δ have the same consequences. A theory is called *finitely axiomatizable* iff it has a finite set of axioms. Since we may form the conjunction of a finite set of axioms, a finitely axiomatizable theory actually always has a single axiom. The set $\bar{\Gamma}$ of all consequences of Γ is the unique closed theory which has Γ as a set of axioms.

PROPOSITION 1.2.14. *Δ is a set of axioms for a theory Γ if and only if Δ has exactly the same models as Γ.*

COROLLARY 1.2.15. *Let Γ_1 and Γ_2 be two theories such that the set of all models of Γ_2 is the complement of the set of all models of Γ_1. Then Γ_1 and Γ_2 are both finitely axiomatizable.*

PROOF. The set $\Gamma_1 \cup \Gamma_2$ is not satisfiable, so it is not finitely satisfiable. Thus we may choose finite sets $\Delta_1 \subset \Gamma_1$, $\Delta_2 \subset \Gamma_2$ such that $\Delta_1 \cup \Delta_2$ is not satisfiable. If $A \models \Delta_1$, then A is not a model of Γ_2, and consequently $A \models \Gamma_1$. It follows by Proposition 1.2.14 that Δ_1 is a finite set of axioms for Γ_1. Similarly Δ_2 is a finite set of axioms for Γ_2. ⊣

The next group of theorems shows connections between mathematical operations on models and syntactical properties of sentences. The first result of this group concerns positive sentences. A sentence φ is said to be *positive* iff φ is built up from sentence symbols using only the two connectives $\wedge$, $\vee$. For example, $(S_0 \wedge (S_2 \vee S_3)) \vee S_{16}$ is positive, while $\neg S_4$ and $S_3 \leftrightarrow S_3$ are not positive. A set Σ of sentences is called *increasing* iff $A \models \Sigma$ and $A \subset B$ implies $B \models \Sigma$.

THEOREM 1.2.16.

(i). *$A \subset B$ if and only if every positive sentence which holds in A holds in B.*

(ii). *A consistent theory Γ is increasing if and only if Γ has a set of positive axioms.*

(iii). *A sentence φ is increasing if and only if either φ is equivalent to a positive sentence, φ is valid, or $\neg \varphi$ is valid.*

PROOF. (i). The fact that, if $A \subset B$, then every positive sentence which holds in A holds in B, is proved by induction. First, every sentence symbol which holds in A holds in B, because of 1.2.3(i) and $A \subset B$. Using 1.2.3(ii) and Exercise 1.2.2, it can be checked that, if the condition 'if φ holds in A, then φ holds in B' is true when $\varphi = \psi$, and when $\varphi = \theta$, then it is also true when $\varphi = \psi \wedge \theta$ and when $\varphi = \psi \vee \theta$. Hence that condition is true for every positive sentence φ.

Suppose that every positive sentence which holds in A holds in B. In particular, for each $S \in \mathscr{S}$, if $A \models S$, then $B \models S$. Thus, if $S \in A$, then $S \in B$, so $A \subset B$. This proves (i).

(ii). Now let Γ be a consistent increasing theory. Let Δ be the set of all positive consequences of Γ. Suppose $B \models \Delta$. Let Σ be the set of all sentences $\neg \varphi$ such that φ is positive and $B \models \neg \varphi$. Let $\neg \varphi_1, \ldots, \neg \varphi_n \in \Sigma$. Then the sentence $\varphi_1 \vee \ldots \vee \varphi_n$ is a positive sentence which fails in B. Hence $\varphi_1 \vee \ldots \vee \varphi_n$ does not belong to Δ and is not a consequence of Γ. Thus the set $\Gamma \cup \{\neg \varphi_1, \ldots, \neg \varphi_n\}$ is satisfiable, and the set $\Gamma \cup \Sigma$ is finitely satisfiable. By the compactness theorem, $\Gamma \cup \Sigma$ has a model, say A. Now for every positive sentence φ which fails in B, $\neg \varphi \in \Sigma$, so φ fails in A. Thus every positive sentence holding in A holds in B, and by (i), $A \subset B$. Since $A \models \Gamma$ and Γ is increasing, we have $B \models \Gamma$. We conclude that every model of Δ is a model of Γ. But $\Delta \subset \bar{\Gamma}$, and therefore Δ is a set of positive axioms for Γ.

Conversely, if Γ has a set of positive axioms, then it follows from (i) that Γ is increasing.

(iii). Let φ be an increasing sentence. We may assume further that φ is satisfiable. If Γ is the set of all consequences of φ, then by (ii) Γ has a positive set Δ of axioms. Now $\varphi \in \Gamma$, so $\Delta \vDash \varphi$, and by Corollary 1.2.13 there is a finite subset $\{\psi_1, \ldots, \psi_n\}$ of Δ such that $\{\psi_1, \ldots, \psi_n\} \vDash \varphi$. If $n = 0$, then φ is valid. Let $n > 0$. Each ψ_m is in Δ and thus in Γ, so each ψ_m is a consequence of φ. It follows that φ is equivalent to the positive sentence $\psi_1 \wedge \ldots \wedge \psi_n$.

Conversely, it follows from (i) that every positive sentence is increasing. Obviously, every valid sentence and every refutable sentence are also increasing. ⊣

A completely trivial fact which is analogous to part (i) of the above theorem is: $A = B$ if and only if every sentence which holds in A holds in B. We shall see later on in this book that the situation is very different in predicate logic, where a maximal consistent theory ordinarily does not even come close to characterizing a single model. This is one thing which makes model theory for predicate logic so much more interesting and difficult than model theory for sentential logic.

We now turn to another kind of sentence. By a *conditional sentence* we mean a sentence $\varphi_1 \wedge \ldots \wedge \varphi_n$, where each φ_i is of one of the following three kinds:

(1) S,

(2) $\neg S_1 \vee \neg S_2 \vee \ldots \vee \neg S_p$,

(3) $\neg S_1 \vee \neg S_2 \vee \ldots \vee \neg S_p \vee T$.

A set Σ of sentences is said to be *preserved under finite intersections* iff $A \vDash \Sigma$ and $B \vDash \Sigma$ implies $A \cap B \vDash \Sigma$. Σ is said to be *preserved under arbitrary intersections* iff for every nonempty set $\{A_i : i \in I\}$ of models of Σ the intersection $\bigcap_{i \in I} A_i$ is also a model of Σ.

Lemma 1.2.17. *A theory Γ is preserved under finite intersections if and only if Γ is preserved under arbitrary intersections.*

Proof. Let Γ be preserved under finite intersections, let $\{A_i : i \in I\}$ be a nonempty set of models of Γ, and let $B = \bigcap_{i \in I} A_i$. Let Σ be the set of all sentences of the form S or $\neg S$ which hold in B. We show that $\Gamma \cup \Sigma$ is satisfiable. Let Σ_0 be an arbitrary finite subset of Σ, and let the negative sentences in Σ_0 be $\neg S_1, \ldots, \neg S_p$. If $p = 0$, all the sentences in Σ_0 are positive, and each of the models A_i is a model of Σ_0, because $B \subset A_i$. Let $p > 0$ and choose models $A_{i_1}, \ldots, A_{i_p}$ from among the A_i such that $S_1 \notin A_{i_1}, \ldots, S_p \notin A_{i_p}$. Then $A = A_{i_1} \cap \ldots \cap A_{i_p}$ is a model of Σ_0; since Γ is preserved under finite intersections, A is also a model of Γ. We have shown that $\Gamma \cup \Sigma$ is finitely satisfiable. By the compactness theorem, $\Gamma \cup \Sigma$ has a model. But the only model of Σ is B, so B is a model of Γ. ⊣

In view of the above lemma, we may as well simply say from now on that Γ is *preserved under intersections*, since it makes no difference whether we say finite or arbitrary intersections.

Theorem 1.2.18.

(i). *A theory Γ is preserved under intersections if and only if Γ has a set of conditional axioms.*

(ii). *A sentence φ is preserved under intersections if and only if φ is equivalent to a conditional sentence.*

Proof. (i). We leave to the reader the proof that every conditional sentence (and hence every set of conditional sentences) is preserved under intersections.

Conversely, let Γ be preserved under intersections. Consider the set Δ of all conditional consequences of Γ. It suffices to show that every model of Δ is a model of Γ. Let B be an arbitrary model of Δ. For each $T \in \mathscr{S} - B$, let Σ_T be the set of all sentences of the form

$$S_1 \wedge \ldots \wedge S_p \wedge \neg T$$

which hold in B. We also let the sentence $\neg T$ itself be in Σ_T. We first note that the conjunction of finitely many sentences in Σ_T is again equivalent to a sentence in Σ_T. Consider a sentence $\varphi \in \Sigma_T$. Then $\neg \varphi$ is clearly equivalent to a conditional sentence ψ either of the form S or of the form

$$\neg S_1 \vee \ldots \vee \neg S_p \vee T.$$

But ψ fails in B, so ψ does not belong to Δ. This means that ψ, and hence $\neg \varphi$, is not a consequence of Γ, and it follows that $\Gamma \cup \{\varphi\}$ is satisfiable. Since Σ_T is, up to equivalence, closed under finite conjunction, we see that $\Gamma \cup \Sigma_T$ is finitely satisfiable. Applying the Compactness Theorem, we may choose a model A_T of $\Gamma \cup \Sigma_T$.

For each $T \in \mathscr{S} - B$, we have $T \notin A_T$ and $B \subset A_T$. Thus, if $\mathscr{S} - B$ is not empty, then

$$B = \bigcap_{T \notin B} A_T.$$

Since each A_T is a model of Γ and Γ is closed under intersections, we have $B \models \Gamma$. In the remaining case $B = \mathscr{S}$, we let Σ be the set of all sentences of the form

$$S_1 \wedge \ldots \wedge S_p.$$

Arguing as before, we find that $\Gamma \cup \Sigma$ is finitely satisfiable and thus has a model. But B is the only model of Σ, so again B is a model of Γ.

We have now shown that every model of Δ is a model of Γ, and it follows that Δ is a set of conditional axioms for Γ.

(ii). This follows from (i) by an argument similar to the last part of the proof of Theorem 1.2.16. ⊣

We conclude with a table which summarizes the semantical and syntactical notions that we have shown to be equivalent (some of these are done in the exercises).

TABLE 1.2.1

Syntax	Semantics
φ is a tautology, $\vdash \varphi$	φ is valid, $\models \varphi$
Σ is consistent	Σ is satisfiable
φ is inconsistent	φ is not satisfiable
φ is deducible from Σ, $\Sigma \vdash \varphi$	φ is a consequence of Σ, $\Sigma \models \varphi$
φ is equivalent to a positive sentence	φ is increasing, and not valid or refutable
φ is equivalent to a conditional sentence	φ is preserved under intersections

EXERCISES

1.2.1. Let A be a model such that $S, T \in A$ and $U, V \in \mathscr{S} - A$. Which of the following sentences are true in A?

$$U, S, T \wedge U, \neg\neg\neg U, S \rightarrow V, S \wedge (S \vee U \leftrightarrow (V \rightarrow T)).$$

1.2.2. Show that, if $\varphi = \psi \vee \theta$, then $A \models \varphi$ if and only if $A \models \psi$ or $A \models \theta$ or both. Concoct similar rules for $A \models \psi \rightarrow \theta$ and $A \models \psi \leftrightarrow \theta$.

1.2.3. A sentence φ is *satisfiable* iff it has at least one model. Show that φ is satisfiable if and only if $\neg \varphi$ is not valid.

1.2.4. A sentence φ is a *consequence of* another sentence ψ, in symbols $\psi \models \varphi$, iff every model of ψ is a model of φ. Show that $\psi \models \varphi$ if and only if $\models \psi \rightarrow \varphi$.

1.2.5. Two sentences φ and ψ are (*semantically*) *equivalent* iff they have exactly the same models. Show that φ and ψ are equivalent if and only if each one is a consequence of the other, and also if and only if $\models \varphi \leftrightarrow \psi$.

1.2.6. Prove that if φ is satisfiable and $\mathscr{S}$ is countable, then the set of all models of φ has the cardinal number of the continuum.

1.2.7* (Interpolation Theorem). Assume that $\varphi \models \psi$. Show that either (i) φ is refutable, (ii) ψ is valid, or (iii) there exists a sentence θ such that $\varphi \models \theta$, $\theta \models \psi$, and every sentence symbol which occurs in θ also occurs in both φ and ψ.

1.2.8. Prove Proposition 1.2.6.

1.2.9*

(i). For every finite set K of models, there is a set Σ of sentences such that K is the set of all models of Σ.

(ii). Give an example of a set Σ of sentences such that the set of all models of Σ is countably infinite.

(iii). Give an example of a countable set of models which cannot be represented as the set of all models of some set of sentences.

In (ii) and (iii), assume that $\mathscr{S}$ is countable.

1.2.10. If $\Sigma \vdash \varphi$ for all $\varphi \in \Gamma$ and if $\Sigma \cup \Gamma \vdash \theta$, then $\Sigma \vdash \theta$.

1.2.11. Prove that the set of all non-models of Σ is empty or of power $2^{|\mathscr{S}|}$.

1.2.12. Show that no positive sentence is valid and no positive sentence is refutable.

1.2.13. A theory Γ is said to be *complete* iff for every sentence φ, exactly one of $\Gamma \models \varphi$, $\Gamma \models \neg \varphi$ holds. For any set Σ of sentences, the following are equivalent:

(i). The set of consequences of Σ is maximal consistent.
(ii). Σ is a complete theory.
(iii). Σ has exactly one model.
(iv). There is a model A such that for all φ, $\Sigma \models \varphi$ iff $A \models \varphi$.

1.2.14. Let Γ be a consistent theory and let B be a model for $\mathscr{S}$. Prove that

B is a model of the set of all positive consequences of Γ if and only if there is a model A of Γ such that $A \subset B$.

1.2.15. Show that every conditional sentence is preserved under intersections.

1.2.16. State and prove the analogue of Exercise 1.2.14 for intersections and conditional sentences.

1.2.17*. Formulate and prove a result like Theorem 1.2.18 for unions of sets of models.

1.2.18. A set Σ of sentences is said to be *independent* iff, for each $\sigma \in \Sigma$, σ is not a consequence of $\Sigma - \{\sigma\}$. Prove that if $\mathcal{S}$ is countable, then every theory Γ in $\mathcal{S}$ has an independent set of axioms.

[*Hint*: Show that Γ has a set of axioms $\Sigma = \{\sigma_1, \sigma_2, \sigma_3, ...\}$ such that, for each n, $\vdash \sigma_{n+1} \rightarrow \sigma_n$ but not $\vdash \sigma_n \rightarrow \sigma_{n+1}$. Then consider the set $\{\sigma_1, \sigma_1 \rightarrow \sigma_2, \sigma_2 \rightarrow \sigma_3, ...\}$.]

1.2.19**. Prove without any restriction on the cardinality of $\mathcal{S}$ that every theory in $\mathcal{S}$ has an independent set of axioms. (The case where $|\mathcal{S}| = \omega_1$ is very much easier than the general case, but still a challenge.)

1.3. Languages, models and satisfaction

We begin here the development of first-order languages in a way parallel to the treatment of sentential logic in Section 1.2. First, we shall define the notions of a first-order predicate language $\mathcal{L}$ and of a model for $\mathcal{L}$. We introduce some basic relations between models – reductions and expansions, isomorphisms, submodels and extensions. We shall then develop the syntax of the language $\mathcal{L}$, defining the sets of terms, formulas and sentences, and presenting the axioms and rules of inference. Finally, we give the key definition of a sentence being true in a model for the language $\mathcal{L}$. The precise formulation of this definition is much more of a challenge in first-order logic than it was for sentential logic. At the end of this section, we state the completeness and compactness theorems (Theorems 1.3.20–1.3.22), but the proofs of these theorems are deferred until the next chapter.

We first establish a uniform notation and set of conventions for such languages and their models. A *language* $\mathcal{L}$ is a collection of symbols. These symbols are separated into three groups, *relation symbols, function symbols* and (*individual*) *constant symbols*. The relation and function symbols of $\mathcal{L}$ will be denoted by capital Latin letters P, F, with subscripts. Lower case Latin letters c, with subscripts, range over the constant symbols of $\mathcal{L}$. If $\mathcal{L}$ is a finite set, we may display the symbols of $\mathcal{L}$ as follows:

$$\mathcal{L} = \{P_0, ..., P_n, F_0, ..., F_m, c_0, ..., c_q\}.$$

Each relation symbol P of $\mathcal{L}$ is assumed to be an n-placed relation for some integer $n \geqslant 1$, depending on P. Similarly, each function symbol F of $\mathcal{L}$ is an m-placed function symbol, where $m \geqslant 1$ and m depends on F. Note that we do not allow 0-placed relation or function symbols. When dealing with several languages at the same time, we use the letters $\mathcal{L}$, $\mathcal{L}'$, $\mathcal{L}''$, etc. If the symbols of the language are quite standard, as for example $+$ for addition, $\leqslant$ for an order relation, etc., we shall simply write

$$\mathcal{L} = \{\leqslant\}, \quad \mathcal{L} = \{\leqslant, +, \cdot, 0\}, \quad \mathcal{L} = \{+, \cdot, -, 0, 1\}, \quad \text{etc.,}$$

for such languages. The number of places of the various kinds of symbols is understood to follow the standard usage. The *power*, or *cardinal* of the language $\mathcal{L}$, denoted by $\|\mathcal{L}\|$, is defined as

$$\|\mathcal{L}\| = \omega \cup |\mathcal{L}|.$$

We say that a language $\mathcal{L}$ is countable or uncountable depending on whether $\|\mathcal{L}\|$ is countable or uncountable.

We occasionally pass from a given language $\mathcal{L}$ to another language $\mathcal{L}'$ which has all the symbols of $\mathcal{L}$ plus some additional symbols. In such cases we use the notation $\mathcal{L} \subset \mathcal{L}'$ and say that the language $\mathcal{L}'$ is an *expansion* of $\mathcal{L}$, and that $\mathcal{L}$ is a *reduction* of $\mathcal{L}'$. In the special case where all the symbols in $\mathcal{L}'$ but not in $\mathcal{L}$ are constant symbols, $\mathcal{L}'$ is said to be a *simple expansion* of $\mathcal{L}$. Since $\mathcal{L}$ and $\mathcal{L}'$ are just sets of symbols, the expansion $\mathcal{L}'$ may be written as $\mathcal{L}' = \mathcal{L} \cup X$, where X is the set of new symbols.

Turning now to the models for a given language $\mathcal{L}$, we first point out that the situation here is more complicated than for the sentential logic $\mathcal{S}$ in Section 1.2. There, each $S \in \mathcal{S}$ could take on at most two values, true or false. Thus the set of intended interpretations for $\mathcal{S}$ has rather simple properties, as the reader discovered. This time, each n-placed relation symbol has as its intended interpretations all n-placed relations among the objects, each m-placed function symbol has as its intended interpretations all m-placed functions from objects to objects, and, finally, each constant symbol has as intended interpretations fixed or constant objects. Therefore, a 'possible world', or model for $\mathcal{L}$ consists, first of all, of a *universe* A, a nonempty set. In this universe, each n-placed P corresponds to an n-placed *relation* $R \subset A^n$ on A, each m-placed F corresponds to an m-placed *function* $G : A^m \to A$ on A, and each constant symbol c corresponds to a *constant* $x \in A$. This correspondence is given by an *interpretation* function $\mathscr{I}$ mapping the symbols of $\mathscr{L}$ to appropriate relations, functions and constants in A. A *model for* $\mathscr{L}$ is a pair $\langle A, \mathscr{I} \rangle$. We use Gothic letters to range over models. Thus we write $\mathfrak{A} = \langle A, \mathscr{I} \rangle$, $\mathfrak{B} = \langle B, \mathscr{J} \rangle$, $\mathfrak{C} = \langle C, \mathscr{K} \rangle$, etc., with appropriate subscripts and superscripts. We shall try to be quite consistent in this respect, so that the universes of the models $\mathfrak{B}'$, $\mathfrak{B}''$, $\mathfrak{B}_i$, $\mathfrak{B}_j$, etc., are precisely the sets B', B'', B_i, B_j, etc. The relations, functions and constants of $\mathfrak{A}$ are, respectively, the images under $\mathscr{I}$ of the relation symbols, function symbols and constant symbols of $\mathscr{L}$.

Note that in a given universe A there are many different permissible interpretations of the symbols of $\mathscr{L}$. Suppose $\mathfrak{A} = \langle A, \mathscr{I} \rangle$, $\mathfrak{A}' = \langle A', \mathscr{I}' \rangle$ are models for $\mathscr{L}$ and R, R' are relations of $\mathfrak{A}$, $\mathfrak{A}'$, respectively. We say that R' is the *corresponding relation to* R if they are the interpretations of the same relation symbol in $\mathscr{L}$, i.e.

$$\mathscr{I}(P) = R \quad \text{and} \quad \mathscr{I}'(P) = R' \quad \text{for some } P \in \mathscr{L}.$$

We introduce similar conventions as regards the functions and constants. When

$$\mathscr{L} = \{P_0, \ldots, P_n, F_0, \ldots, F_m, c_0, \ldots, c_q\},$$

we write the models for $\mathscr{L}$ in displayed form as

$$\mathfrak{A} = \langle A, R_0, \ldots, R_n, G_0, \ldots, G_m, x_0, \ldots, x_q \rangle.$$

When the symbols of $\mathscr{L}$ are familiar, we shall agree to use, for instance,

$$\mathfrak{A} = \langle A, \leqslant, +, \cdot \rangle$$

for models of the language $\mathscr{L} = \{\leqslant, +, \cdot\}$. We may resort to

$$\mathfrak{A} = \langle A, \leqslant_{\mathfrak{A}}, +_{\mathfrak{A}}, \cdot_{\mathfrak{A}} \rangle, \qquad \mathfrak{B} = \langle B, \leqslant_{\mathfrak{B}}, +_{\mathfrak{B}}, \cdot_{\mathfrak{B}} \rangle, \text{ etc.,}$$

if the context of the discussion requires it.

If we start with a model $\mathfrak{A}$ for the language $\mathscr{L}$ we can always expand it to a model for the language $\mathscr{L}' = \mathscr{L} \cup X$ by giving appropriate interpretations for the symbols in X. If $\mathscr{I}'$ is any interpretation for the symbols of X in $\mathfrak{A}$, and X is disjoint from $\mathscr{L}$, then $\mathfrak{A}' = \langle A, \mathscr{I} \cup \mathscr{I}' \rangle$ is a model for $\mathscr{L}'$. In this case we say that $\mathfrak{A}'$ is an *expansion* of $\mathfrak{A}$ to $\mathscr{L}'$, and $\mathfrak{A}$ is the *reduct* of $\mathfrak{A}'$ to $\mathscr{L}$. Sometimes we use the shorter notation $(\mathfrak{A}, \mathscr{I}')$ for $\mathfrak{A}'$. Clearly, there are many ways a model $\mathfrak{A}$ for $\mathscr{L}$ can be expanded to a model $\mathfrak{A}'$ for $\mathscr{L}'$. On the other hand, given a model $\mathfrak{A}'$ for $\mathscr{L}'$, it has only one reduction $\mathfrak{A}$ to $\mathscr{L}$. Namely, we form $\mathfrak{A}$ by restricting the interpretation function $\mathscr{I}'$ on $\mathscr{L} \cup X$ to $\mathscr{L}$. The processes of expansion and reduction do not change the universe of the model.

The *cardinal*, or *power*, of the model $\mathfrak{A}$ is the cardinal $|A|$. $\mathfrak{A}$ is said to be finite, countable or uncountable if $|A|$ is finite, countable or uncountable. Note that on a finite universe A, while there can be only finitely many different relations, functions and constants, the number of different interpretation functions $\mathscr{I}$ can be very large and depends on $|\mathscr{L}|$.

We next introduce some simple but basic notions and operations on models. The reader should go through the exercises at the end of this section in order to be familiar with them.

Two models $\mathfrak{A}$ and $\mathfrak{A}'$ for $\mathscr{L}$ are *isomorphic* iff there is a 1-1 function f mapping A onto A' satisfying:

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

$$R(x_1 \ldots x_n) \quad \text{if and only if} \quad R'(f(x_1) \ldots f(x_n))$$

for all $x_1, \ldots, x_n$ in A.

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

$$f(G(x_1 \ldots x_m)) = G'(f(x_1) \ldots f(x_m)),$$

for all $x_1, \ldots, x_m$ in A.

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

$$f(x) = x'.$$

A function f that satisfies the above is called *an isomorphism of* $\mathfrak{A}$ *onto* $\mathfrak{A}'$, or *an isomorphism between* $\mathfrak{A}$ *and* $\mathfrak{A}'$. We use the notation $f : \mathfrak{A} \cong \mathfrak{A}'$ to denote that f is an isomorphism of $\mathfrak{A}$ onto $\mathfrak{A}'$, and we use $\mathfrak{A} \cong \mathfrak{A}'$ for $\mathfrak{A}$ is isomorphic to $\mathfrak{A}'$. For convenience we use $\cong$ to denote the *isomorphism relation* between models for $\mathscr{L}$. It is quite clear that $\cong$ is an equivalence relation. Furthermore, it preserves powers, that is, if $\mathfrak{A} \cong \mathfrak{B}$, then $|A| = |B|$. Indeed, unless we wish to consider the particular structure of each element of A or B, for all practical purposes $\mathfrak{A}$ and $\mathfrak{B}$ are the same if they are isomorphic.

A model $\mathfrak{A}'$ is called a *submodel* of $\mathfrak{A}$ if $A' \subset A$ and:

(i). Each n-placed relation R' of $\mathfrak{A}'$ is the restriction to A' of the corresponding relation R of $\mathfrak{A}$, i.e., $R' = R \cap (A')^n$.

(ii). Each m-placed function G' of $\mathfrak{A}'$ is the restriction to A' of the corresponding function G of $\mathfrak{A}$, i.e., $G' = G|(A')^m$.

(iii). Each constant of $\mathfrak{A}'$ is the corresponding constant of $\mathfrak{A}$.

We use $\mathfrak{A}' \subset \mathfrak{A}$ to denote that $\mathfrak{A}'$ is a submodel of $\mathfrak{A}$, and the symbol $\subset$ for the submodel relation between models for $\mathscr{L}$. The reader should show that $\subset$ is a partial-order relation and that, if $\mathfrak{A} \subset \mathfrak{B}$, then $|A| \leqslant |B|$. We say that $\mathfrak{B}$ is an *extension* of $\mathfrak{A}$ if $\mathfrak{A}$ is a submodel of $\mathfrak{B}$.

Combining the above two notions, we say that $\mathfrak{A}$ is *isomorphically embedded in* $\mathfrak{B}$ if there is a model $\mathfrak{C}$ and an isomorphism f such that $f : \mathfrak{A} \cong \mathfrak{C}$ and $\mathfrak{C} \subset \mathfrak{B}$. In this case we call the function f an *isomorphic embedding of* $\mathfrak{A}$ *in* $\mathfrak{B}$. If $\mathfrak{A}$ is isomorphically embedded in $\mathfrak{B}$, then $\mathfrak{B}$ is isomorphic to an extension of $\mathfrak{A}$.

To formalize a language $\mathscr{L}$, we need the following *logical symbols* (see the corresponding development for $\mathscr{S}$ in Section 1.2.):

parentheses	), (;
variables	$v_0, v_1, \ldots, v_n, \ldots$;
connectives	$\wedge$ (and), $\neg$ (not);
quantifier	$\forall$ (for all);

and one binary relation symbol $\equiv$ (identity).

We assume, of course, that no symbol in $\mathscr{L}$ occurs in the above list. Certain strings of symbols from the above list and from $\mathscr{L}$ are called *terms*. They are defined as follows:

1.3.1.

(i). A variable is a term.

(ii). A constant symbol is a term.

(iii). If F is an m-placed function symbol and $t_1, \ldots, t_m$ are terms, then $F(t_1 \ldots t_m)$ is a term.

(iv). A string of symbols is a term only if it can be shown to be a term by a finite number of applications of (i)–(iii).

The *atomic formulas* of $\mathscr{L}$ are strings of the form given below:

1.3.2.

(i). $t_1 \equiv t_2$ is an atomic formula, where t_1 and t_2 are terms of $\mathscr{L}$.

(ii). If P is an n-placed relation symbol and $t_1, \ldots, t_n$ are terms, then $P(t_1 \ldots t_n)$ is an atomic formula.

Finally, the *formulas* of $\mathscr{L}$ are defined as follows:

1.3.3.

(i). An atomic formula is a formula.

(ii). If φ and ψ are formulas, then $(\varphi \wedge \psi)$ and $(\neg \varphi)$ are formulas.

(iii). If v is a variable and φ is a formula, then $(\forall v)\varphi$ is a formula.

(iv). A sequence of symbols is a formula only if it can be shown to be a formula by a finite number of applications of (i)–(iii).

Just as in the case of $\mathscr{S}$, we may put definitions 1.3.1 and 1.3.3 in a set-theoretical setting. Namely, the set of terms of $\mathscr{L}$ is the least set T such that

T contains all constant symbols and all variables v_n, $n = 0, 1, 2, \ldots$, and, whenever F is an m-placed function symbol and $t_1, \ldots, t_m \in T$, then $F(t_1 \ldots t_m) \in T$.

Similarly, the set of formulas of $\mathscr{L}$ is the least set Φ such that

every atomic formula belongs to Φ and, whenever $\varphi, \psi \in \Phi$ and v is a variable, then $(\varphi \wedge \psi)$, $(\neg \varphi)$, $(\forall v)\varphi$ all belong to Φ.

Note that we have tacitly used the letters t (with subscripts) to range over terms, v to range over variables, and φ, ψ to range over formulas. Again, we emphasize that *properties of terms and formulas of* $\mathscr{L}$ *can only be established by an induction based on definitions* 1.3.1 *and* 1.3.3.

We can now introduce the abbreviations $\vee$, $\rightarrow$, $\leftrightarrow$ as in Section 1.2. Furthermore, we adopt all the conventions introduced earlier. The new symbol $\exists$ (there exists) is introduced as an abbreviation defined as

$$(\exists v)\varphi \quad \text{for} \quad \neg(\forall v) \neg \varphi.$$

Some new conventions are the following:

$$\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_n \quad \text{for} \quad (\varphi_1 \wedge (\varphi_2 \wedge \ldots \wedge \varphi_n));$$
$$\varphi_1 \vee \varphi_2 \vee \ldots \vee \varphi_n \quad \text{for} \quad (\varphi_1 \vee (\varphi_2 \vee \ldots \vee \varphi_n));$$
$$(\forall x_1 x_2 \ldots x_n)\varphi \quad \text{for} \quad (\forall x_1)(\forall x_2) \ldots (\forall x_n)\varphi;$$
$$(\exists x_1 x_2 \ldots x_n)\varphi \quad \text{for} \quad (\exists x_1)(\exists x_2) \ldots (\exists x_n)\varphi.$$

At this point we assume that the reader has enough experience in first-order predicate logic to continue the development on his own. In particular, we leave it to him to decide on the notions of *subformulas*, of *free* and *bound* occurrences of a variable in a formula, and to give a proper definition (based on definitions 1.3.1 and 1.3.3) of *substitution* of a term for a variable in a formula.

We now come to an extremely important convention of notation. To make sure that the reader does not miss it, we enclose it in a box: