An Introduction to Non-Classical Logic

This revised and considerably expanded edition of *An Introduction to Non-Classical Logic* brings together a wide range of topics, including modal, tense, conditional, intuitionist, many-valued, paraconsistent, relevant and fuzzy logics. Part I, on propositional logic, is the old *Introduction*, but contains much new material. Part II is entirely novel, and covers quantification and identity for all the logics in Part I. The material is unified by the underlying theme of world semantics. All of the topics are explained clearly and accessibly, using devices such as tableau proofs, and their relations to current philosophical issues and debates is discussed. Students with a basic understanding of classical logic will find this book an invaluable introduction to an area that has become of central importance in both logic and philosophy. It will also interest people working in mathematics and computer science who wish to know about the area.

GRAHAM PRIEST is Boyce Gibson Professor of Philosophy, University of Melbourne and Arché Professorial Fellow, Departments of Philosophy, University of St Andrews. His most recent publications include *Towards Non-Being* (2005) and *Doubt Truth to be a Liar* (2006).

An Introduction to Non-Classical Logic

From If to Is

Second Edition

GRAHAM PRIEST
University of Melbourne
and
University of St Andrews

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, Delhi

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521670265

First published 2001
Second edition 2008

Printed in the United Kingdom at the University Press, Cambridge

A catalogue record for this publication is available from the British Library

ISBN 978-0-521-85433-7 hardback
ISBN 978-0-521-67026-5 paperback

To all those from whom I have learned

Contents

Cambridge University Press
978-0-521-67026-5 - An Introduction to Non-Classical Logic: From If to Is, Second Edition
Graham Priest
Frontmatter
More information

Cambridge University Press
978-0-521-67026-5 - An Introduction to Non-Classical Logic: From If to Is, Second Edition
Graham Priest
Frontmatter
More information

Preface to the First Edition

Around the turn of the twentieth century, a major revolution occurred in logic. Mathematical techniques of a quite novel kind were applied to the subject, and a new theory of what is logically correct was developed by Gottlob Frege, Bertrand Russell and others. This theory has now come to be called 'classical logic'. The name is rather inappropriate, since the logic has only a somewhat tenuous connection with logic as it was taught and understood in Ancient Greece or the Roman Empire. But it is classical in another sense of that term, namely standard. It is now the logic that people normally learn when they take a first course in formal logic. They do not learn it in the form that Frege and Russell gave it, of course. Several generations of logicians have polished it up since then; but the logic is the logic of Frege and Russell none the less.

Despite this, many of the most interesting developments in logic in the last forty years, especially in philosophy, have occurred in quite different areas: intuitionism, conditional logics, relevant logics, paraconsistent logics, free logics, quantum logics, fuzzy logics, and so on. These are all logics which are intended either to supplement classical logic, or else to replace it where it goes wrong. The logics are now usually grouped under the title 'non-classical logics'; and this book is an introduction to them.

The subject of non-classical logic is now far too big to permit the writing of a comprehensive textbook, so I have had to place some restrictions on what is covered.[1] For a start, the book is restricted to propositional logic. This is not because there are no non-classical logics that are essentially first-order (there are: free logic), but because the major interest in non-classical logics is *usually* at the propositional level. (Often, the quantifier

[1] For a brief introduction and overview of the field, see Priest (2005a).

extensions of these logics are relatively straightforward.) Within propositional logics, I have also restricted the logics considered here to ones which are relevant to the debate about conditionals ('if . . . then . . .' sentences). Again, this is not because this exhausts non-classical propositional logics (there is quantum logic, for example), but because taking the topic of conditionals as a *leitmotiv* gives the material a coherence that it might otherwise lack. And, of course, conditionals are about as central to logic as one can get.

The major semantical technique in non-classical logics is possible-world semantics. Most non-classical logics have such semantics. This is therefore the major semantical technique that I use in the book. In many ways, the book could be thought of as a set of variations on the theme of possible-world semantics. It should be mentioned that many of the systems discussed in the book have semantics other than possible-world semantics – notably, algebraic semantics of some form or other. Those, however, are an appropriate topic for a different book.

Choosing a kind of proof theory presents more options. Logic is about validity, what follows from what. Hence, the most natural proof theories for logic are natural deduction systems and sequent calculi. Most of the systems we will consider here can, in fact, be formulated in these ways. However, I have chosen not to use these techniques, but to use tableau methods instead (except towards the end of the book, where an axiomatic approach becomes necessary). One reason for this choice is that constructing tableau proofs, and so 'getting a feel' for what is, and what is not, valid in a logic, is very easy (indeed, it is algorithmic). Another is that the soundness and, particularly, completeness proofs for logics are very simple using tableaux. Since these areas are both ones where inexperienced students experience difficulty, tableaux have great pedagogical attractions. I first learned to do tableaux for modal logics, in the way that they are presented in the book, from my colleagues Rod Girle and the now greatly missed Ian Hinckfuss. The myriad variations they take on here are my own.

This book is not meant to provide a first course in logic. I assume that readers are familiar with the classical propositional calculus, though I review this material fairly swiftly in chapter 1. (I do not assume that students are familiar with tableaux, however.) Chapter 2 introduces the basic semantic technique of possible worlds, in the form of semantics for basic modal logic. Chapters 3 and 4 extend the techniques to other modal logics.

Chapter 3 looks at other normal systems of modal logic. Chapter 4 looks at non-normal worlds and their uses. Chapter 5 extends the semantic techniques, yet again, to so-called conditional logics. (The material in chapter 5 is significantly harder than anything else before the last couple of chapters of Part I.)

The non-classical logics up to this point are all most naturally thought of as extensions of classical logic. In the subsequent chapters of Part I, the logics are most naturally seen as rivals to it. Chapter 6 deals with intuitionism. Chapter 7 introduces many-valued logics, and the idea that there might be truth-value gaps (sentences that are *neither* true nor false) and gluts (sentences that are *both* true and false). Chapter 8 then describes first degree entailment, a central system of both relevant and paraconsistent logics. The semantic techniques of the final chapters fuse the techniques of both modal and many-valued logic. Non-normal worlds come into their own in chapter 9, where basic relevant logics are considered. Chapter 10 considers relevant logics more generally; and in chapter 11 fuzzy logic comes under the microscope. The chapters are broken up into sections and subsections. Their numeration is self-explanatory.

The major aim of this book is to explain the basic techniques of non-classical logics. However, these techniques do not float in mid-air: they engage with numerous philosophical issues, especially that of conditionality. The meanings of the techniques themselves also raise important philosophical issues. I therefore thought it important to include some philosophical discussion, usually towards the end of each chapter. The discussions are hardly comprehensive – quite the opposite; but they at least serve to elucidate the technical material, and may be used as a springboard for a more extended consideration for those who are so inclined.

Since proofs of soundness and completeness are such an integral part of modern logic, I have included them for the systems considered here, where possible. This technical material is relatively self-contained, however, and, even though the matter in the book is largely cumulative, can be skipped without prejudice by those who have no need, or taste, for it. For this reason, I have relegated the material to separate sections, marked with an asterisk. These sections also take for granted a little more mathematical sophistication on the part of the reader. Towards the end of each chapter there are also sections containing some historical details and giving suggestions for further reading. At the conclusion of each chapter is a section containing a set of problems, exercises and questions. To understand the material in any but a relatively superficial way, there is no substitute for engaging with these. Questions that pertain to the sections marked with an asterisk are themselves marked with an asterisk, and can be ignored without prejudice.

I have taught a course based on the material in this book, or similar material, a number of times over the last ten years. I am grateful to the generations of students whose feedback has helped to improve both the content and the presentation. I have learned more from their questions than they would ever have been aware of. I am particularly grateful to the class of '99, who laboured under a draft of the book, picking up numerous typos and minor errors. I am grateful, too, to Aislinn Batstone, Stephen Read and some anonymous readers for comments which greatly improved the manuscript. I am sure that it could be improved in many other ways. But if one waited for perfection, one would wait for ever.

Preface to the Second Edition

The first edition of *Introduction to Non-Classical Logic* deals with just propositional logics. In 2004, Cambridge University Press and I decided to produce a second volume dealing with quantification and identity in non-classical logics. Late in the piece, it was decided to put the old and the new volumes together, and simply bring out one omnibus volume. The practical decision caused a theoretical problem. Was it the same book as the old *Introduction* or a different one? The answer – as befits a book on non-classical logic – was, of course, both. So the name of the book had to be the same and different. We decided to achieve this seeming impossibility by adding an appropriate sub-title to the book, 'From If to Is'. Though there are many propositional operators and connectives, the conditional, 'if', is perhaps the most vexed. It is, at any rate, the focus around which the old *Introduction* moves. Whether or not 'if' is univocal is a contentious matter; but 'is' is certainly said in many ways. There is the 'is' of predication ('Ponting is Australian'), the 'is' of existence ('There is a spider in the bathtub', 'Socrates no longer is'), and the 'is' of identity ('2 plus 2 is 4'). All of these are in play in first-order logic; they provide the focus around which the new part of the book moves.

On Part I

Though Part I of the present volume is essentially the old *Introduction to Non-Classical Logic*, I have taken the opportunity of revising its contents. With one exception, the revisions simply add new material. Some of the additions are made in the light of what is coming in Part II. Thus, there is a new section on equivalence relations and equivalence classes in the *Mathematical Prolegomenon*. But most of them comprise material that could usefully have been in the old *Introduction*, or that I would have put there had I thought to do so. These are as follows:

- Chapter 3 now contains material on tense logic.
- Chapter 4 contains a section on the modal system *S*0.5, and related systems. This makes the bridge between non-normal logics and the impossible worlds of chapter 9 patent.
- In chapter 7, the section on supervaluations has been extended slightly.
- In chapter 8, a new section on relational semantics and tableaux for $Ł_3$ and RM_3 has been added.
- Chapter 9 now contains a section on systems of 'constructible negation', making a connection with chapter 6 on intuitionist logic. I have renamed this chapter 'Logics with Gaps, Gluts and Worlds' to indicate better its contents. This allowed chapter 10 to be renamed simply 'Relevant Logics'.
- I have added a technical appendix to chapter 11 on fuzzy logic. The Łukasiewicz logic of that chapter is, in fact, a special case of a more general construction. That construction is, perhaps, less likely to be of interest to philosophers. But I think that it is a good idea to have the material there, at least for the sake of reference.
- The appendix, chapter 11a, is a last-minute addition. In a paper I was writing in 2006 I wanted to refer to the general theory of many-valued modal logics. I could not find anything suitable in the literature, so I drafted one. I was persuaded by Stephen Read that this would be a helpful addition to the book.

If it was not already so before, the additions now make it entirely impossible to cover all of the material in Part I in a one-semester course. But it is better to have material there which a teacher can skip over, than no material on a topic which a teacher would like to cover.

The one place where material has not simply been added is in chapter 10 on relevant logic (with a few knock-on consequences in chapter 11). As 10.9 explains, the semantics given in that chapter are not the original Routley–Meyer semantics, but the 'simplified semantics' developed later (by Priest, Sylvan and Restall). It has now turned out that the original simplified semantics completeness proof is incorrect with respect to one of the axioms, $A \to ((A \to B) \to B)$ (A11 in the old *Introduction*) – though this does not affect the tableau completeness proof. In the context of the simplified semantics, the condition C11 of the old *Introduction* is too strong; and the extra strength, resuscitating, as it does, the Disjunctive Syllogism, is not of a desirable kind. The condition can, however, be modified in such a way as to be complete. (See Restall and Roy (200+).) This modification is now employed in chapter 10, occasioning a new section on content inclusion and some more relevant logics whose semantics employ this notion.

In producing the present Part I, I have decided to leave the section and subsection numbering of the old *Introduction* unchanged. It was therefore necessary to accommodate new material in a way that does not disturb the numbering. I use letters to indicate interpolations that would otherwise do so. Thus, subsections between, e.g., 4.3.6 and 4.3.7 are 4.3.6a, 4.3.6b, etc.; and a section between 4.3 and 4.4 is 4.3a, so that its subsections become 4.3a.1, 4.3a.2, etc.

In writing the old *Introduction*, I decided, again as its preface explains, to employ tableaux, as far as possible. Systems of natural deduction have a great deal to recommend them, however. It is therefore very welcome that Fitch-style systems of natural deduction for all the logics of the old *Introduction* have been produced by Tony Roy. These (together with soundness and completeness proofs) can be found in Roy (2006).

Finally, in producing the new Part I, I have taken the opportunity to correct typos, as well as pedagogical and other minor infelicities. A number of people have pointed these out to me; these include Stephan Cursiefen, Rafał Grusczyńiski, Maren Kruger, Jenny Louise, Tanja Osswald, Stephen Read, Wen-fang Wang, and the members of the Arché Logic Group at the University of St Andrews (see below). Kate Manne, Stephen Read and Elia Zardini provided helpful comments on the new material. Finally, correspondence with Petr Hájek was invaluable in writing the appendix to chapter 11. Warm thanks go to all of them.

On Part II

When I wrote the first edition of *Introduction to Non-Classical Logic*, I decided to restrict myself to propositional logic for the reasons explained in its preface. Someone who has mastered that material certainly has a good grasp of what non-classical logics are all about. But it cannot be denied that a book which leaves matters there is leaving the job half done. If any non-classical logic is to be applied, then quantifiers and, probably, identity, are going to be essential. And certainly the philosophical issues surrounding the technical constructions are as acute as anything in the propositional case. Hence it was (in a moment of weakness) that I decided to write a second volume dealing with quantifiers and identity in non-classical logics. That volume would contain details of the behaviour of first-order quantifiers and identity in the logics of the old *Introduction*. As mentioned above, that material eventually became Part II of this volume.

Explaining the techniques of a large number of logics perspicuously and relatively briefly presents various exegetical challenges. So it was with Part I. Part II adds to these. The material in this is, by its nature, more difficult than that in Part I. (Although, by the time a student reaches this material, they are, one would hope, a little more sophisticated, so a little more may be expected of them – or required by them.) Most obviously the semantics of quantifiers are more intricate than those of the connectives. Less obviously, technical results, such as compactness and the Löwenheim–Skolem theorems, assume more importance. This book does not pretend to provide a comprehensive introduction to the metatheory of non-classical logics, important as that topic is. But those who are familiar with some of these matters from classical logic will naturally be curious to know how things stand with respect to the various non-classical logics. Fortunately, then, many of the elementary metatheoretic properties of a logic follow, in a relatively uniform way, from the fact that it has a sound and complete proof system (tableau, axiomatic, or whatever). I have covered the relevant matters for classical logic in chapter 12, and then simply pointed out that essentially the same considerations apply to all the other logics in the book – except for fuzzy logic in chapter 25, where completeness finally fails.

More difficult is the fact that the techniques used permit systematic independent variations. These can be applied in the case of many, if not most, of the logics covered in Part II. The result is a plethora of disparate systems. Attempting to cover all of them in the book would make it far too long, and would, I think, result in the danger of the reader losing the wood for the trees; it would also, I suspect, become tiresome. I therefore decided to explain the relevant variations in detail for certain logics, but to consider their applications to others only when there was some particular point to doing so. Thus, to give one example, the constants employed for the most part in the logics are rigid designators. But all the systems with world semantics can be augmented with non-rigid designators as well. How to do this is explained in the case of modal logic in chapters 16 and 17. I leave it (usually in problems) to those who want other systems of logic with non-rigid designators to extrapolate the techniques for themselves.

As in Part I, I assume that the reader is familiar with the relevant parts of classical logic. There is a review of the material in chapter 12. Free logic is necessary at various places in Part II. Chapter 13 presents this. Perhaps the most important of the aforementioned variations is that between constant domain semantics and variable domain semantics. Chapter 14 explains constant domain modal logic; chapter 15 explains variable domain modal logic. Another important variation is that between necessary identity and contingent identity. Chapter 16 spells out necessary identity in modal logic; chapter 17 spells out contingent identity in the same context. After that, all the fundamental techniques are in place, and the subsequent chapters correspond, one to one, to chapters 4 to 11 of Part I, covering non-normal modal logics, conditional logics, intuitionist logic, many-valued logics, First Degree Entailment, logics with semantics employing worlds and many-values, relevant logics and fuzzy logics. The reader is well advised to be familiar with (or refresh their memory of) the relevant chapter of Part I before passing on to the corresponding chapter of Part II. But, generally speaking, it is unnecessary to master the material after a chapter in Part I to understand material for the corresponding logic in Part II. Thus, for example, it is quite possible to read the material on modal logic in Part I, and then move on directly to the chapters on modal logic in Part II. There are a few notational changes between Part I and Part II. These are very minor, and will not hinder understanding (or usually even be noticed!).

Again as with Part I, the logics of this part inform and are informed by important philosophical considerations. Perhaps the most important of these concern existence and its various machinations. At the appropriate points I have therefore discussed these things. The discussions do little more than raise the relevant issues. But they at least show the reader what is at issue in the technical matters, and provide a certain amount of focusing for the diverse topics. And proofs of theorems and other technical matters are relegated to the starred appendices of each chapter, which can be omitted by uninterested readers.

There is, of course, much more to be said about non-classical logics than can be said here. For example – just to mention a few topics – all the logics in this part can be augmented with function symbols; they can all be extended to second-order logics; and all have algebraic semantics of various kinds. At one time I thought to include some of these topics in this book. But eventually I judged undesirable the additional complexity and length that this would have involved. These topics can be covered in Part III – if anyone should care to write it; it won't be me.

The manuscript of this Part has been much improved by comments and suggestions from a number of people. I taught an honours logic course based on a draft of the manuscript at the University of Melbourne in the first half of 2006, where the students provided helpful feedback. My colleagues Allen Hazen and Greg Restall sat in on the class and provided many helpful suggestions. Kate Manne worked carefully through the whole draft and polished it considerably. Later that year, the Arché Logic Group at the University of St Andrews also worked through the manuscript and made a number of valuable suggestions: Philip Ebert, Andri Hjálmarsson, Ole Hjorthland, Ira Kiourti, Stephen Read, Marcus Rossberg, Andreas Stokke, and, most especially, Elia Zardini. Finally, correspondence with Petr Hájek was invaluable in writing the appendix to chapter 25. To all of them, my warmest thanks. These go, also, to Hilary Gaskin and the staff of Cambridge University Press for all they have done to make this volume possible – indeed, actual.

Book Website

All books contain errors, from the trivial typo, through infelicities of various degrees, to the serious screw-up. I hope that there aren't too many in this book – especially of the last kind! Details of any corrections that I am aware need to be made can be found on the website www.cambridge.org/priest. In due course, the website will also contain solutions to selected exercises.

Mathematical Prolegomenon

In expositions of modern logic, the use of some mathematics is unavoidable. The amount of mathematics used in this text is rather minimal, but it may yet throw a reader who is unfamiliar with it. In this section I will explain briefly three bits of mathematics that will help a reader through the text. The first is some simple set-theoretic notation and its meaning. The second is the notion of proof by induction. The third concerns the notion of equivalence relations and equivalence classes. It is not necessary to master the following before starting the book; the material can be consulted if and when required.

0.1 Set-theoretic Notation

0.1.1 The text makes use of standard set-theoretic notation from time to time (though never in a very essential way). Here is a brief explanation of it.

0.1.2 A set, X, is a collection of objects. If the set comprises the objects $a_1, \ldots, a_n$, this may be written as $\{a_1, \ldots, a_n\}$. If it is the set of objects satisfying some condition, $A(x)$, then it may be written as $\{x : A(x)\}$. $a \in X$ means that a is a *member* of the set X, that is, a is one of the objects in X. $a \notin X$ means that a is not a member of X.

0.1.3 *Examples*: The set of (natural) numbers less than 5 is $\{0, 1, 2, 3, 4\}$. Call this F. The set of even numbers is $\{x : x$ is an even natural number$\}$. Call this E. Then $3 \in F$, and $5 \notin E$.

0.1.4 Sets can have any number of members. In particular, for any a, there is a set whose only member is a, written $\{a\}$. $\{a\}$ is called a *singleton* (and is not to be confused with a itself). There is also a set which has no members, the *empty set*; this is written as ϕ.

0.1.5 *Examples*: {3} is the set containing just the number three. It has one member. It is distinct from 3, which is a number, not a set at all, and so has no members.[2] $3 \notin \phi$.

0.1.6 A set, X, is a *subset* of a set, Y, if and only if every member of X is a member of Y. This is written as $X \subseteq Y$. The empty set is a subset of every set (including itself). $X \subset Y$ means that X is a *proper* subset of Y; that is, everything in X is in Y, but there are some things in Y that are not in X. X and Y are identical sets, $X = Y$, if they have the same members, i.e., if $X \subseteq Y$ and $Y \subseteq X$. Hence, if X and Y are not identical, $X \neq Y$, either there are some members of X that are not in Y, or vice versa (or both).

0.1.7 *Examples*: Let N be the set of all natural numbers, and E be the set of even numbers. Then $\phi \subseteq N$ and $E \subseteq N$. Also, $E \subset N$, since $5 \in N$ but $5 \notin E$. If $X \subseteq N$ and $X \neq E$ then either some odd number is in X, or some even number is not in X (or both).

0.1.8 The *union* of two sets, X, Y, is the set containing just those things that are in X or Y (or both). This is written as $X \cup Y$. So $a \in X \cup Y$ if and only if $a \in X$ or $a \in Y$. The *intersection* of two sets, X, Y, is the set containing just those things that are in both X and Y. It is written $X \cap Y$. So $a \in X \cap Y$ if and only if $a \in X$ and $a \in Y$. The *relative complement* of one set, X, with respect to another, Y, is the set of all things in Y but not in X. It is written $Y - X$. Thus, $a \in Y - X$ if and only if $a \in Y$ but $a \notin X$.

0.1.9 Examples: Let N, E and O be the set of all numbers, all even numbers, and all odd numbers, respectively. Then $E \cup O = N$, $E \cap O = \phi$. Let $T = \{x : x \geq 10\}$. Then $E - T = \{0, 2, 4, 6, 8\}$.

0.1.10 An *ordered pair*, $\langle a, b \rangle$, is a set whose members occur in the order shown, so that we know which is the first and which is the second. Similarly for an ordered triple, $\langle a, b, c \rangle$, quadruple, $\langle a, b, c, d \rangle$, and, in general, n-tuple, $\langle x_1, \ldots, x_n \rangle$. Given n sets $X_1, \ldots, X_n$, their *cartesian product*, $X_1 \times \cdots \times X_n$, is the set of all n-tuples, the first member of which is in X_1, the second of which is in X_2, etc. Thus, $\langle x_1, \ldots, x_n \rangle \in X_1 \times \cdots \times X_n$ if and only if $x_1 \in X_1$ and ... and $x_n \in X_n$. A *relation*, R, between $X_1, \ldots, X_n$ is any subset of $X_1 \times \cdots \times X_n$.

[2] In some reductions of number theory to set theory, 3 is identified with a certain set, and so may have members. But in the most common reduction, 3 has three members, not one.

$\langle x_1, \ldots, x_n \rangle \in R$ is usually written as $Rx_1 \ldots x_n$. If n is 3, the relation is a *ternary* relation. If n is 2, the relation is a *binary* relation, and Rx_1x_2 is usually written as x_1Rx_2. A *function* from X to Y is a binary relation, f, between X and Y, such that for all $x \in X$ there is a unique $y \in Y$ such that xfy. More usually, in this case, we write: $f(x) = y$.

0.1.11 *Examples*: $\langle 2, 3 \rangle \neq \langle 3, 2 \rangle$, since these sets have the same members, but in a different order. Let N be the set of numbers. Then $N \times N$ is the set of all pairs of the form $\langle n, m \rangle$, where n and m are in N. If $R = \{\langle 2, 3 \rangle, \langle 3, 2 \rangle\}$ then $R \subseteq N \times N$ and is a binary relation between N and itself. If $f = \{\langle n, n^2 \rangle : n \in N\}$, then f is a function from numbers to numbers, and $f(n) = n^2$.

0.2 Proof by Induction

0.2.1 The method of proof by induction (or recursion) on the complexity of sentences is used heavily in the asterisked sections of the book. It is also used occasionally in other places, though these can usually be skipped without loss. What this method comes to is this. Suppose that all of the simplest formulas of some formal language (that is, those that do not contain any connectives or quantifiers) have some property, P. (Establishing this fact is usually called the *basis* (or *base*) *case*.) And suppose that whenever one constructs a more complex sentence – that is, one with an extra connective (or quantifier if such things are in the language) – out of formulas that have property P, the resulting formula also has the property P. (Establishing this is usually called the *induction case*.) Then it follows that all the formulas of the language have the property P. Thus, for example, suppose that the simple formulas p and q have property P, and that whenever formulas have that property, so do their negations, conjunctions, etc. Then it follows that $\neg p$, $p \wedge q$, $\neg p \wedge (p \wedge q)$, have the property, as do all sentences that we can construct from p and q using negation and conjunction.

0.2.2 The proof of the induction case normally breaks down into a number of different sub-cases, one for each of the connectives (and quantifiers if present) employed in the construction of more complex formulas. Thus, we assume that A has the property, then show that $\neg A$ has it; we assume that A and B have the property, then show that $A \wedge B$ has it; and so on for every connective (and quantifier). The assumption, in each case, is called the *induction hypothesis*.

0.2.3 Here is a simple example of a proof by induction. We show that every formula of the propositional calculus which is grammatical according to the rules of 1.2.2 has an even number of brackets. (This is a bit like cracking a nut with a sledgehammer; but it illustrates the method clearly.) The symbol ■ marks the end of a proof.

Proof:
BASIS CASE: First, we need to establish that this result holds for all of the simplest formulas, the propositional parameters. All such formulas have no (zero) brackets, and 0 is an even number. Hence, the result holds for propositional parameters.
INDUCTION CASE: Next we must establish that if the result holds for some formulas, and we construct other formulas out of those, the result holds for these too. So suppose that A and B have an even number of brackets. (This is the induction hypothesis.) We need to show that each of $\neg A$, $(A \vee B)$, $(A \wedge B)$, $(A \supset B)$ and $(A \equiv B)$ has an even number of brackets too. There is one case for each of the constructions in question.
For $\neg$: the number of brackets in $\neg A$ is the same as the number of brackets in A. Since this is even (by the induction hypothesis), the result follows. (We did not use the induction hypothesis concerning B in this case, but that does not matter.)
For $\vee$: suppose that the number of brackets in A is a, and the number of brackets in B is b. Then the number of brackets in $(A \vee B)$ is $a + b + 2$ (since the construction introduces two new brackets). But a and b are even, and so $a + b + 2$ is even. Hence, the number of brackets in $(A \vee B)$ is even, as required.
For $\wedge$, $\supset$, and $\equiv$: the arguments are exactly the same as for $\vee$. We have now established the basis case and the induction case. It follows from these that the result holds for all formulas; that is, all grammatical formulas have an even number of brackets. ■

0.3 Equivalence Relations and Equivalence Classes

0.3.1 The notion of an equivalence relation is one that is very useful on a number of occasions, especially when identity comes into play. An equivalence relation on a domain of objects is one, essentially, that chunks the domain into a collection of disjoint (i.e., non-overlapping) classes called *equivalence classes*. Thus, given a class of people, C, 'x has the same height as y' is a relation that partitions them into classes of people with the same height. Suppose that C is:

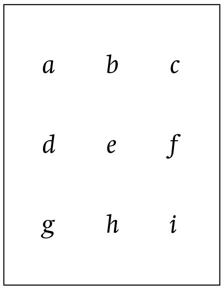

and that a, b, d and e, all have the same height, as do c, f and i, as do g and h. Then the equivalence classes are:

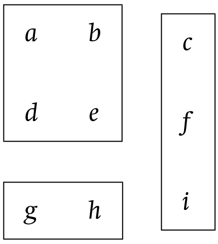

0.3.2 More precisely, if $\sim$ is a binary relation on a collection of objects, C, it is an *equivalence relation* just if it is:

- reflexive: for all $x \in C$, $x \sim x$
- symmetric: for all $x, y \in C$, if $x \sim y$ then $y \sim x$
- transitive: for all $x, y, z \in C$, if $x \sim y$ and $y \sim z$ then $x \sim z$

If $x \in C$, its *equivalence class*, written $[x]$, is defined as $\{w \in C : w \sim x\}$.

0.3.3 The fundamental fact about equivalence classes is that every object in the domain is in exactly one. To see this, note, first, that for any $x \in C$, since $x \sim x$, $x \in [x]$; so x is in *some* equivalence class. Now let $X = [x]$ and $Y = [y]$. Suppose that, for some z, z is in both X and Y. Then $z \sim x$ and $z \sim y$. By symmetry and transitivity, $x \sim y$. For any $w \in X$, $w \sim x$. Since $x \sim y$, $w \sim y$. That is, $w \in Y$. Hence, $X \subseteq Y$. Similarly, $Y \subseteq X$. Hence, $X = Y$.

0.3.4 In constructions employing equivalence classes, it is common to specify a property of a class in terms of one of its members, thus:

$F([x])$ if and only if $G(x)$

Now suppose that $[x] = [y]$. Then the definition will go awry if we can have $G(x)$ but not $G(y)$. In such a definition it is therefore always important to establish that if $x \sim y$, $G(x)$ if and only if $G(y)$.