Chapter 1

Introduction for teachers

Readers with no previous knowledge of formal logic will find it more useful to begin with Chapter 2.

Purpose and intended audience

1.1. *CNL* (*Classical and Nonclassical Logics*) is intended as an introduction to mathematical logic. However, we wish to immediately caution the reader that the topics in this book are

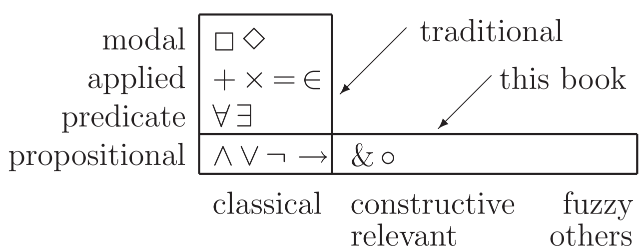

not the same as those in a conventional introduction to logic. *CNL* should only be adopted by teachers who are aware of the differences and are persuaded of this book's advantages. Chiefly, *CNL* trades some depth for breadth:

- A traditional introduction to logic covers *classical logic only*, though possibly at several levels — propositional, predicate, modal, etc.

- *CNL* is *pluralistic*, in that it covers classical and several non-classical logics — constructive, quantitative, relevant, etc. — though almost solely at the propositional level.

Of course, a logician needs both depth and breadth, but both cannot be acquired in the first semester. The depth-first approach is prevalent in textbooks, perhaps merely because classical logic developed a few years before other logics. I am convinced that a breadth-first approach would be better for the students, for reasons discussed starting in 1.9.

1.2. *Intended audience.* This is an introductory textbook. No previous experience with mathematical logic is required. Some experience with algebraic computation and abstract thinking is expected, perhaps at the precalculus level or slightly higher. The exercises in this book are mostly computational and require little originality; thus *CNL* may be too elementary for a graduate course. Of course, the book may be used at different levels by different instructors.

CNL was written for classroom use; I have been teaching undergraduate classes from earlier versions for several years. However, its first few chapters include sufficient review of prerequisite material to support the book's use also for self-guided study. Those chapters have some overlap with a "transition to higher mathematics" course; *CNL* might serve as a resource in such a course.

I would expect *CNL* to be used mainly in mathematics departments, but it might be adopted in some philosophy departments as well. Indeed, some philosophers are very mathematically inclined; many of this book's mathematical theorems originated on the chalkboards of philosophy departments. Colleagues have also informed me that this book will be of some interest to students of computer science, but I am unfamiliar with such connections and have not pursued them in this book.

1.3. *In what sense is this new?* This book is a work of exposition and pedagogy, not of research. All the main theorems of this book have already appeared, albeit in different form, in research journals or advanced monographs. But those articles and books were written to be read by experts. I believe that the present work is a substantially new *selection and reformulation of results*, and that it will be more accessible to beginners.

1.4. *Avoidance of algebra.* Aside from its pluralistic approach (discussed at much greater length later in this chapter), probably *CNL*'s most unusual feature is its attempt to avoid higher algebra.

In recent decades, mathematical logic has been freed from its philosophical and psychological origins; the current research literature views different logics simply as different kinds of algebraic structures. That viewpoint may be good for research, but it is not a good prescription for motivating undergraduate students, who know little of higher algebra.

CNL attempts to use as little algebra as possible. For instance, we shall use topologies instead of Heyting algebras; they are more concrete and easier to define. (See the remark in 4.6.i.)

1.5. *Rethinking of terminology.* I have followed conventional terminology for the most part, but I have adopted new terminology whenever a satisfactory word or phrase was not available in the literature. Of course, what is "satisfactory" is a matter of opinion.

It is my opinion that there are far too many objects in mathematics called "regular," "normal," etc. Those words are not descriptive — they indicate only that some standard is being adhered to, without giving the beginner any assistance whatsoever in identifying and assimilating that standard. Whenever possible, I have attempted to replace such terms with phrases that are more descriptive, such as "truth-preserving" and "tautology-preserving."

A more substantive, and perhaps more controversial, example of rejected terminology is "intuitionistic logic." That term has been widely used for one particular logic since it was introduced in the early 20th century by Brouwer, Heyting, and Kolmogorov. To call it anything else is to fight a strong tradition. But the word "intuitionistic" has connotations of subjectivity and mysticism that may drive away some scientifically inclined students. There is nothing subjective, mystical, or unscientific about this interesting logic, which we develop in Chapters 10, 22, 27, 28, and part of 29.

Moreover, not all mathematicians share the same intuition. Indeed, aside from logicians, most mathematicians today are schooled only in classical logic and find all other logics to be nonintuitive. It is only a historical accident that Brouwer, Heyting and Kolmogorov appropriated the word "intuitionistic" for their system. The term "BHK logic," used in some of the literature, is less biased, but it too is descriptive only to someone who already knows the subject.

A more useful name is "constructive logic," because BHK logic is to a large extent the reasoning system of constructive mathematics (discussed in 2.42–2.46). Mathematicians may not be entirely in agreement about the importance of constructivism, but at least there is consensus on what the term "constructive" means. Its meaning in mathematics is quite close to its meaning outside mathematics, and thus should be more easily grasped by beginning students.

1.6. *What is not covered.* This book is intended as an introductory textbook, not an encyclopedia — it includes enough different logics to illustrate some basic ideas of the subject, but it does not include all major logics. Derivations in *CNL* follow only the Hilbert style, because in my opinion that is easiest for beginners to understand. The treatment of quantifiers consists of only a few pages (sections 5.40–5.51), and that treatment is informal, not axiomatic. Omitted entirely (or mentioned in just a sentence or two) are $\Box$, $\Diamond$, formal predicate logic, Gentzen sequents, natural deduction, modal logics, Gödel's Incompleteness Principles, recursive functions, Turing machines, linear logic, quantum logic, substructures logics, nonmonotonic logics, and many other topics.

Topics in the book

1.7. *Order of topics.* I have tried to arrange the book methodically, so that topics within it are not hard to find; but I have also provided frequent cross-referencing, to facilitate reading the book in other orders than mine.

Chapter 2 gives an overview of, and informal introduction to, the subject of logic. The chapter ends with a detailed discussion (2.42–2.46) of constructivism and Jarden's Proof, surely the simplest example of the fact that a different philosophy can require a different logic.

Chapters 3 and 4 give a brief introduction to naive set theory and general topology. Chapter 5 gives a more detailed introduction to informal classical logic, along with comments about how it compares with nonclassical logics and with ordinary nonmathematical English. Particular attention is given to the ambiguities of English.

Chapters 2–5 may be considered "prerequisite" material, in the sense that their content is not part of logic but will be used to develop logic. Different students will need different parts of this prerequisite material; by including it I hope to make the book accessible to a wide variety of students. Admittedly, these introductory chapters take up an unusually large portion of the book, but they are written mostly in English; the remainder of the book is written in the more concise language of mathematics.

Finally, in Chapter 6 we begin formal logic. This chapter presents and investigates a formal language that will be used throughout the remainder of the book. Among the terms defined in this chapter are "formula," "rank of a formula," "variable sharing," "generalization," "specialization," and "order preserving" and "order reversing."

There are several feasible strategies for ordering the topics after formal language. The most obvious would be to present various logics one by one — e.g., classical logic, then constructive logic, then relevant logic, etc. This strategy would juxtapose related results — e.g., constructive semantics with constructive syntactics — and perhaps it is the most desirable approach for a reference book. But I have instead elected to cover all of semantics before beginning any syntactics. This approach is better for the beginning student because semantics is more elementary and concrete than syntactics, and because this approach juxtaposes related *techniques* — e.g., constructive semantics and relevant semantics.

Semantics is introduced in Chapter 7, which defines "valuation," "interpretation," and "tautology." Then come some examples of interpretations — numerically valued in Chapter 8, set-valued in Chapter 9, and topological in Chapter 10. In the presentation of these examples, one recurring theme is the investigation of relevance: If A and B are formulas that are unrelated in the sense that they share no propositional variable symbols, and $A \to B$ is a tautology in some interpretation, does it follow that $\overline{A}$ or B are tautologies? Our conclusions are summarized in one column of the table in 2.37.

The aforementioned chapters deal with examples of semantic systems, one at a time. Chapter 11, though not lacking in examples, presents more abstract results. Sections 11.2–11.7 give shortcuts that are often applicable in verifying that a formula is tautologous. Sections 11.8–11.12 give sufficient conditions for one interpretation to be an extension of another. Sections 11.13–11.17 show that, under mild assumptions, the Dugundji formula in n symbols is tautological for interpretations with fewer than n semantic values, but not for interpretations with n or more semantic values; as a corollary we see that (again under mild assumptions) an infinite semantics cannot be replaced by a finite semantics.

Syntactics is introduced in Chapter 12, which defines "axiom," "assumed inference rule," "derivation," "theorem," etc. The chapters after that will deal with various syntactic logics, but in what order should those be presented? My strategy is as follows.

The logics of greatest philosophical interest in this book are classical, constructive (intuitionist), relevant, and fuzzy (Zadeh and Łukasiewicz), shown in the upper half of the diagram below. These logics have a substantial overlap, which I call *basic logic*;[1] it appears at the bottom of the diagram. To reduce repetition, our syntactic development will begin with basic logic and then

[1] That's my own terminology; be cautioned that different mathematicians use the word "basic" in different ways.

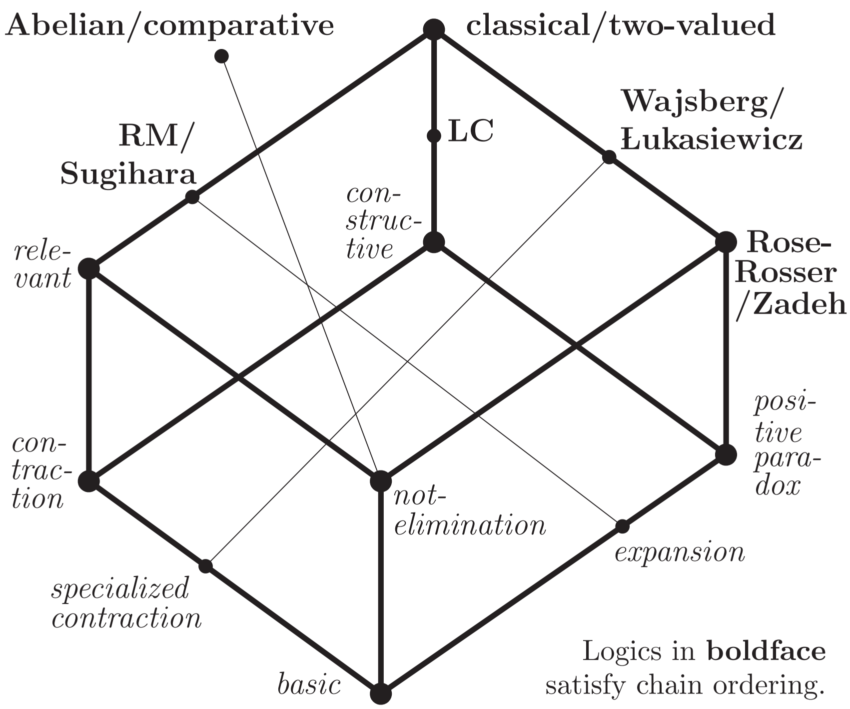

gradually add more ingredients.

Chapter 13 introduces the assumptions of basic implication,

$$\{A, A \to B\} \vdash B, \qquad \vdash (G \to H) \to [(I \to G) \to (I \to H)],$$
$$\vdash C \to C, \qquad \vdash [D \to (E \to F)] \to [E \to (D \to F)],$$

and investigates their consequences. One elementary but important consequence is the availability of detachmental corollaries; that is, $\vdash A \to B \;\Rightarrow\; A \vdash B$. Chapter 14 adds the remaining assumptions of basic logic,

$$\{A, B\} \vdash A \wedge B, \qquad \vdash A \to (A \vee B), \qquad \vdash (A \wedge B) \to A,$$
$$\vdash (A \to \overline{B}) \to (B \to \overline{A}), \qquad \vdash B \to (A \vee B), \qquad \vdash (A \wedge B) \to B,$$
$$\vdash [(A \to B) \wedge (A \to C)] \to [A \to (B \wedge C)],$$
$$\vdash [(B \to A) \wedge (C \to A)] \to [(B \vee C) \to A],$$
$$\vdash [A \wedge (B \vee C)] \to [(A \wedge B) \vee C],$$

and then investigates their consequences. One consequence is a substitution principle: if $\mathcal{S}$ is an order preserving or order reversing function, then $\vdash A \to B$ implies, respectively, $\vdash \mathcal{S}(A) \to \mathcal{S}(B)$ or $\vdash \mathcal{S}(B) \to \mathcal{S}(A)$.

Next come several short chapters, each investigating a different one-axiom extension of basic logic:

Chapter		axiom added to basic logic
15	Contraction	$\vdash (A \to (A \to B)) \to (A \to B)$,
16	Expansion and positive paradox	$\vdash (A \to B) \to (A \to (A \to B))$, $\vdash A \to (B \to A)$
17	Explosion	$\vdash (A \wedge \overline{A}) \to B$ (conjunctive) or $\vdash A \to (\overline{A} \to B)$ (implicative)
18	Fusion	$\vdash [(A \& B) \to X] \leftrightarrow [A \to (B \to X)]$,
19	Not-elimination	$\vdash \overline{\overline{A}} \to A$,
20	Relativity	$\vdash ((A \to B) \to B) \to A$.

Those extensions are considered independently of one another (i.e., results of one of those chapters may not be assumed in another of those chapters), with this exception:

$$\text{relativity} \quad \Rightarrow \quad \text{not-elimination} \quad \Rightarrow \quad \text{fusion.}$$

Anticipating the discussion below, we mention a few more one-axiom extensions :

24.5	Implicative disjunction	$\vdash ((A \to B) \to B) \to (A \vee B)$
15.3.a	Specialized contraction	$\vdash (A \to (A \to \overline{A})) \to (A \to \overline{A})$
8.37.f	Centering	$\vdash \ (A \to A) \ \leftrightarrow \ \overline{A \to A}$

The preceding chapters have shown that various expressions are derivable. Chapter 21 introduces soundness, a new tool that will finally enable us to show that certain expressions are *not* derivable, a fact that we have only hinted at in earlier chapters. A pairing of an interpretation (semantic) with an axiomatization (syntactic) is *sound* if

$$\left\{ \begin{array}{c} \text{theorems of the} \\ \text{axiomatization} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{tautologies of the} \\ \text{interpretation} \end{array} \right\}.$$

If equality holds, we have *completeness*, but that's much harder to establish and doesn't come until the end of the book. Soundness is introduced at this point because we can put it to good use in the next few chapters.

The next few chapters investigate our "major" logics:

Chapter		assumptions: basic logic plus ...
22	Constructive	positive paradox, contraction, explosion;
23	Relevant	not-elimination and contraction;
24	Fuzzy	positive paradox, implicative disjunction, specialized contraction, not-elimination;
25	Classical	all of the above;
26	Abelian	relativity and centering.

These chapters include, among other things, several deduction principles, converses to the detachmental corollary procedure; see 2.37.

Chapter 27 proves the propositional version of Harrop's admissibility rule for constructive logic. The proof is via Meyer's "metavaluation," a computational device that is a sort of mixture of semantics and syntactics. Two corollaries are the Disjunction Property and Mints's Admissibility Rule. The latter is the most elementary example of an admissibility rule that is not also a derivable inference rule. (That admissibility and derivability are the same in classical logic is proved in 29.15.)

Finally, we prove some completeness pairings. Proofs for constructive implication and relevant implication, in Chapter 28, use Kripke-style "multiple worlds" interpretations. Proofs for classical, fuzzy, and constructive logic are presented in Chapter 29, all using what I call "maximal Z-unproving sets" — i.e., sets $\mathcal{S}$ that are maximal for the property that $\mathcal{S} \nvdash Z$.

1.8. *What to cover; what to skip.* That's up to the individual instructor's own taste, but here are a few suggestions and hints based on my own teaching experience.

Different students are motivated in different ways, so the first few chapters cover several different kinds of introductions — historical, linguistic, etc. In my own lectures, I skip large parts of the first few chapters, leaving those parts as recommended reading for students with weak backgrounds. But my lectures generally *do* include Jarden's example (2.45), sets of sets (3.3–3.6), Russell's Paradox (3.11–3.12), Venn diagrams (3.47–3.55.a), topologies and interiors (4.1–4.18), and quantifiers (5.40–5.51).

I do *not* lecture on every section in the book. I merely hit the highlights and some crucial parts; I expect the students to read the rest on their own. I permit my students to use their textbooks and notes on all tests and quizzes after midsemester, because I do not expect students to memorize vast collections of formulas. I encourage students to write their own summaries, or "crib sheets," for use during those tests and quizzes; preparing such summaries is an excellent way to study.

I tend to spend most of my semester on those parts of the book that have exercises. The exercises may seem repetitious and mechanical to advanced readers, but they are crucial for bringing the beginner into frequent and close contact with basic ideas. I assign many of the exercises as homework, and give tests made of similar problems. Note that the chapters vary greatly in length,[2] and the exercises do not appear at a constant rate, so a syllabus cannot be planned by a rule like "a chapter every week."

Though there are occasional exceptions, the general trend in the book is from elementary topics to topics of greater mathematical sophistication. Graduate students might skip the first few chapters of the book; at the other extreme, my undergraduate classes generally do not reach the last few chapters. Still, I do state some of the results of those chapters. Even without their proofs, the results on admissibility and completeness play an essential role in tying the book together conceptually, so they are mentioned frequently throughout the book.

One simple way to abridge the book is to leave out all the

---

[2]It has been my experience that students sometimes become confused about which formula schemes or inference rules are permitted as justifications in any particular homework problem. I have arranged the book so that I can answer, "anything from Chapters such-and-such, plus anything from the beginning of the current chapter to our current location in that chapter." This consideration outweighed the unpleasantness of having chapters with vastly different lengths.

topological material, and all the material on constructive logic that is unavoidably topological — but retain the nontopological parts of constructive logic. This means skipping Chapters 4 and 10, as well as some parts of Chapters 22 and 29. Chapter 27 does not depend on topology, so the instructor might decide near the end of the semester whether to cover it or jettison it, depending on how much time is left. Its proofs are slightly tedious, but they do make good exercises, and the conclusion of that chapter is one of the juiciest bits of the book.

Why pluralism?

1.9. It has been my experience that many mathematicians, even including some accomplished logicians, are unfamiliar with the pluralist approach. I believe that many of them would like it if they gave it a try.

This subchapter is, quite frankly, a sales pitch. I envision an instructor of logic standing in a bookstore or library or at a conference exhibit table, leafing through textbooks and trying to choose one for his/her course. In the next few pages I will attempt to persuade that instructor that pluralism is pedagogically superior to the traditional, classical-only approach — i.e., that it will bring beginning students to a better understanding of logic.

1.10. The traditional and pluralist approaches to introductory logic share the goal of conveying to students not just one or a few separate logics, but also certain deeper ideas that are common to all logics:

> interpretations, derivations, soundness, completeness, independence of axioms, redundancy of connectives, sharing of propositional variable symbols, the finite model property, etc.

Though the two approaches share deeper ideas, they differ greatly on the surface level — i.e., in their choice of examples. Indeed, consider again the diagram in 1.1; the small box in the lower left corner represents the overlap between traditional depth and pluralist breadth. Because the overlap is small, it would not be meaningful to statistically compare the two approaches by giving the same exam to students from the two courses. Instead I will advocate pluralism using arguments based on everyday experience and commonly held principles of teaching.

1.11. *The classical-only approach is unnatural and artificial*, whereas the pluralist approach is motivated by the students' own nonmathematical experiences. Classical logic presented by itself doesn't really make much sense; it embraces non sequiturs such as

if today is Tuesday then the earth is round

— true for a classical logician, but nonsense for anyone else. The student is left wondering *why* implication is defined the way it is defined; it seems rather arbitrary.

Classical logic is computationally the simplest of all the main logics, and most mathematicians are comfortable with it as a method for presenting mathematical proofs. But it does not closely resemble the way that we actually *think* most of the time. Human thought — even that of mathematicians — is a mixture of many logics, not just classical. Human thought may be too complicated to be fully understood, but some of its ingredients are simple enough to analyze. The few logics in *CNL* were selected from the much wider variety in the literature, in part for their computational simplicity, but also for their philosophical and/or psychological significance.

Admittedly, the student's everyday experience is not mathematical, and it is not precise enough to actually be used as a justification in any mathematical *proofs*. But that experience provides intuition and motivation, which are invaluable to beginners.

Some teachers, familiar only with classical logic, may fear that pluralism will open the floodgate of cultural relativism: If all logics are permitted, then no one of them is of any particular interest or value. But just the opposite is true. Accepting an arbitrary-seeming definition as the only correct one deprives us of any meaningful choice, whereas pluralism explores the different advantages enjoyed by different logics.

1.12. *Classical propositional semantics is too easy.* Later in the semester, after students in the traditional course have become accustomed to classical logic and resigned to its artificiality, they learn to use true/false tables. For instance, they can determine that $(P \to \neg Q) \to (Q \to \neg P)$ is true by plugging in the four combinations TT, TF, FT, FF for P and Q.

But then students may not see any need for syntactic derivations. Why bother to prove something, when we already know it is true? (An analogous pedagogical problem arises in the traditional, Euclidean-only course on geometry; see 2.29. Students may use pictures to learn isolated facts about lines and circles without understanding the proofs that connect those facts.)

And even if we persuade them that proofs are somehow worthwhile, students may still have difficulty understanding the criteria that determine whether a proof is *correct*. An omitted step will hardly be noticed, when we already know the conclusion is true. The choice of steps may seem dictated by arbitrary ritual rather than logical necessity.

1.13. *Reasoning requires doubt, which requires plausible contrasting alternatives.* The traditional approach to logic presents classical results — e.g., excluded middle or the Herbrand-Tarski deduction principle — as absolute truths, with no plausible alternatives.

To teach students to reason mathematically, we must encourage them to doubt everything until it is proven. The teacher should be a guide who helps find the right questions, not an authority who dictates the answers. The student should learn to always ask,

what if this is not true?

That question is hypothetical and abstract, and requires more sophistication than is possessed by many beginning students. They need plausible concrete alternatives — i.e., they can ask the easier question

> what happens in one of the other examples
> that we've been studying?

And those examples are useful for more advanced students too, and even for researchers. But to ask about other examples, first we must *have* other examples.

1.14. Even students who have already learned to doubt will gain insights from the additional examples. One of the most widely held principles of teaching is that

> *whenever possible, an abstract idea should be accompanied by several different examples.*

One example — in this case, classical logic — hardly suffices to explain abstract notions such as the law of the excluded middle, the deduction principle, variable sharing, or explosion. See the table in 2.37 for a summary of how those four abstract notions fare in five different logics studied in this book.

Many traditional textbooks follow a route that builds up toward Gödel's Incompleteness Principles. But I think that many of the students studying from such a textbook are lost by the end of the semester, largely because the incompleteness principles are preceded by too few examples of completeness. *CNL* proves completeness for several logics and describes it for several others; see the table in 2.26.

1.15. *The traditional emphasis on predicate logic sometimes obscures basic ideas.* The most fundamental ideas of logic — derivation, interpretation, etc. — are already present, in simpler but still nontrivial form, in propositional logic. Predicate logic complicates the presentation but in many cases does not enrich the ideas.

A good example of this is the classical Herbrand-Tarski deduction principle.[3] The predicate logic version (found in most logic textbooks) is complicated and hard to understand:

[3] Known as the "Deduction Theorem" in most of the literature; see 2.18.

> Let F and G be formulas, and let $\mathcal{H}$ be a set of formulas. Assume $\mathcal{H} \cup \{F\} \vdash G$. Then $\mathcal{H} \vdash F \rightarrow G$ ...
>
> ... *provided that*, in the given derivation $\mathcal{H} \cup \{F\} \vdash G$, none of the steps involves a substitution in which a free individual variable being replaced is one that appears in F.

In effect, the hypothesis "$\mathcal{H} \cup \{F\} \vdash G$" is much more complicated than it at first appears. We need to know not only that G is derivable from $\mathcal{H} \cup \{F\}$, but also *how* it is derivable, all the way down to nitty-gritty details about scopes of quantifiers.

Now contrast that with the propositional logic version. In propositional logic, there are no free or bound individual variables, so we can omit the entire fine print clause. The resulting principle is shorter and simpler, but still retains much of the power of the predicate logic version.

Admittedly, predicate logic is too important to omit altogether. A first course in logic — for most students, the *only* course in logic — should include at least a brief, informal introduction to quantifiers, as *CNL* does in 5.40–5.51. But that topic does not need to skew the entire semester; we don't necessarily have to carry $\forall$ and $\exists$ through our formal development.

1.16. In all fairness, I must also mention one seeming disadvantage of pluralism. Classical logic is computationally the simplest of all the major logics. To compare it with anything else, we must accept some complications.

For instance, teachers who are familiar with classical logic's short formulation (three axioms, two definitions, and one inference rule, given in 25.1.d) may be dismayed to see a dozen axioms in nonclassical logics. Those longer lists are unavoidable. In classical logic we could define two of $\vee, \wedge, \rightarrow, \neg$ in terms of the other two, but for some nonclassical logics we need all four connectives as primitives. (See 10.9.) More primitives require more axioms to govern them.

Still, the dozen axioms need not all be swallowed at one time. They can be digested in several courses, each quite tasty by itself. That is the plan we have described in 1.7.

1.17. *Pluralism is more modern; the prepluralistic view is rather antiquated.* Traditional, classical-only textbooks are still largely built around Gödel's important (but not elementary) discoveries of the 1930s.

I do not advocate teaching to our undergraduates whatever is the latest discovery. Thousands of new mathematical discoveries are published in research journals every year, but most of those discoveries are too advanced and specialized to deserve the attention of beginning students.

Still, when a fundamental (important and/or elementary) development does comes along, we should not overlook it. Such a development has occurred gradually during the late 20th century, in the work of Kripke, Anderson, Belnap, Zadeh, and others. Pluralism (many logics considered simultaneously) has become predominant in the research literature — not as a focus of attention, but as a fact in the background, an accepted and assimilated part of the paradigm. Moreover, enough elementary examples have accumulated in recent decades to make pluralism feasible in the beginners' classroom. I believe that beginning logic students would benefit from pluralism, for reasons indicated in the preceding pages.

1.18. *Pluralism has greater applicability to computer science.* Or at least, so I have been told. I know too little of that subject, so it will not be pursued in the present edition of this book.

Feedback

1.19. Despite the best efforts of myself and the editors at Princeton University Press, I'm sure a few errors remain in this edition. I hope they're small ones, but at any rate I'll list them on a web page when they're reported. I may also post some shorter proofs, if I learn of any. The main page for this book is at

`http://www.math.vanderbilt.edu/~schectex/logics/`

and links from that page will lead to the errata and addenda pages. Also I invite suggestions for alterations for a second edition, which might or might not eventually happen. Email me at `eric.schechter@vanderbilt.edu`.

(Yes, the URL and email are spelled differently.) A second edition, if it happens, probably will have shorter introductory chapters, less repetition, more exercises, and perhaps (if I can find sufficiently elementary ones) proofs of some of the results that are stated without proof in this edition.

Acknowledgments

1.20. This book benefited from discussions with Greg Bush, Elvira Casal, Dave Easley, Mark Ellingham, Jonathan Farley, Isidore Fleischer, Klaus Glashoff, Peter Jipsen, Bjarni Jónsson, Christian Khoury, Ayan Mahalanobis, Ralph McKenzie, Charles Megibben, Peter Nyikos Dave Renfro, Fred Richman, Peter Suber, Constantine Tsinakis, and others. Special thanks go to Norm Megill, who first introduced me to matrix interpretations, and to Robert Meyer, who answered many questions.

This book would not have been written without the help of free Latex software from Donald Knuth, Leslie Lamport, Christian Schenk, Sven Wiegand, and others. Several of the examples in this book were located using MaGIC (MAtrix Generator for Implication Connectives), a computer program made available by John Slaney of the Automated Reasoning Group at Australian National University; my thanks to Norm Megill for porting the program to Windows. I am also grateful to Jonathan Balsam, Nikolaos Galatos, Norm Megill, and the people at Princeton University Press, whose careful reading of previous versions caught many errors. Of course, any errors that remain are my own.

Most of all, I am indebted to the students who permitted me to experiment on them. I hope their pains will be justified by benefits to later students.