Chapter 1

Introduction

Some arguments are good and other arguments are bad. We are all familiar with bad arguments. Discussions of politics, economics, race, and gender, are littered with them. And we generally consider our own arguments to be good. Our arguments for particular views about various hot-button topics are usually successful; or so we think.

But what distinguishes good arguments from bad ones? What must an argument be like—what properties must it have—in order to be good? What, in short, makes some arguments good and other arguments bad?

These questions are not mere academic, ivory-tower speculation. They are extremely important, and their answers really matter. For in arguments, as in life, not just anything goes. Some arguments really do establish their conclusions; other arguments really do not. Some arguments, that is, are good; others, not so much. And in order to change the world for the better—to promote social justice, enact political reform, and so on—we need a theory of the distinction between good arguments and bad ones.

This book is an introduction to two such theories: propositional logic, and first-order logic. Both propositional logic and first-order logic consist of a formal language, an account of what it takes for sentences in that language to be true, and an account of what it takes for arguments in that language to be good. All this can be used to formulate a reasonably decent theory of what makes some English arguments good and other English arguments bad. The purpose of this book is to explain exactly how.

A quick qualification: there are many different ways in which an argument can be good, and this book is only about one of them. Some arguments are good because they are informative. Other arguments are good because they are convincing. Still other arguments are good because they have morally good consequences. And still other arguments are good because they are, in a practical sense, very useful.

In this book, I focus on one way in which arguments can be good: they can be good in the sense of having a good formal structure, that is, a structure of

the sort that propositional logic and first-order logic describe. For it is worth learning about this specific way in which arguments can be good. All the other ways of being good are, unsurprisingly, related to this one. In fact, all good-making features of arguments are extremely intertwined: arguments which have a good formal structure, for instance, are often—not always, but often—more convincing than arguments which do not. So to understand everything that goes into good arguments, we need to improve our understanding of every different way in which arguments can be good. This book contributes to that, by focusing on the particular ways of being good that propositional logic and first-order logic capture.

Think of it like this. Because of how important arguments are, it is worth developing a comprehensive theory of what distinguishes good arguments from bad ones. The theories presented in this book, based on propositional logic and first-order logic, contribute to that comprehensive theory. There are other theories, of the other ways in which arguments can be good—for instance, being good by being informative, or convincing, or moral, or useful—which contribute to that comprehensive theory as well. And those other theories are, of course, worth learning. But the theories based on propositional logic, and first-order logic, are worth learning too. Hence this book.

Here is a rough outline of the theories to come. Roughly put, according to those theories, an argument is good just in case it can be translated into a formal argument—in either the language of propositional logic, or the language of first-order logic—which has the following property: every way of making the premises true is also a way of making the conclusion true. Good arguments, in other words, are arguments in which the truth of the premises *guarantees* the truth of the conclusion. Propositional logic provides one precise account of what it means for a conclusion to be 'guaranteed' by some premises. First-order logic provides another, complementary account.[1]

By the end of this book, you will be able to take an English argument—from a blog post, an online message, a video clip, or wherever—and

  (i) extract a more rigorous English argument from it,
 (ii) translate that more rigorous argument into either the language of propositional logic or the language of first-order logic,
(iii) use propositional logic, or first-order logic, to determine whether that translated argument is good or bad, and then
(iv) determine, on the basis of that, whether or not the original English argument is good or bad.

Taken together, steps (i)–(iv) represent a simple, accessible, yet rigorous theory

---

[1]There are more accounts of how arguments can be good in this formal sort of way. Those accounts are based on logical systems which will not be discussed here. For presentations of some such systems, see (Priest, 2008; Sider, 2010). In this book, I focus on propositional logic and first-order logic because they are perhaps the simplest, strongest, and most influential accounts of one way of distinguishing good arguments from bad arguments. So they are among the most important accounts to study.

of the distinction between good arguments and bad arguments. By learning how to follow steps (i)–(iv), you will learn how to draw that distinction.

As a result, you will be better equipped to evaluate arguments in articles, papers, books, blog posts, social media, advertisements, video clips, television, movies, and more. Steps (i)–(iv), in other words, will help you accurately assess the arguments that you routinely encounter. You will be able to tell when someone gives you an argument that fails to establish its conclusion. You will be able to tell when the assortment of claims that someone makes, passionately and perhaps charismatically, do not actually form an argument at all. And your capacity for constructive self-criticism will improve: you will be better at determining when you have, and when you lack, good arguments for your views. So in general, after learning steps (i)–(iv), you will be a more discerning, clear-headed, and honest reasoner.

And that, in turn, will make you a more effective agent for social change. Of course, to achieve justice, we need more than just good reasoners: we need better politicians, more effective schools, improved health care, and many other things. But we need good reasoners too. We need people who can effectively evaluate—and respond to—the arguments given by politicians, lawyers, lobbyists, CEOs, billionaires, journalists, media personalities, religious leaders, and others in positions of power. And logic can help with that.

A quick note on the scope and limits of this book. Many introductory logic books cover more details than I cover here.[2] Such books are great for those who want to delve into the abstract theory of logic, perhaps because they want to become logicians, or computer scientists, or linguists, or philosophers. But such books are not so great for those who just want to use logic to improve their everyday reasoning about things that matter to them. Similarly, many critical reasoning books cover more social justice issues—and cover them more deeply—than I do here.[3] Such books are great for those who want a big-picture, birds-eye overview of logic and reasoning. But such books are not so great for those who want to learn, in rigorous detail, the core components of logic itself.

This book helps fill the gap between these two literatures. On the one hand, it contains more detailed discussions of the relevance of formal logic to everyday life—the connection, in particular, between logic and social justice—than standard logic textbooks. On the other hand, it contains more detailed discussions of formal logic than standard critical reasoning books.

Here is a summary of the material to come. Part I covers the basic theory of natural language arguments. In particular, in Chapter 2, I explain what natural language arguments are. Then I explain—in a rough and intuitive

---

[2]For other introductions to propositional logic and first-order logic, as well as discussions of those logical systems' more advanced features, see (Chiswell & Hodges, 2007; Enderton, 2001; Shoenfield, 1967).

[3]For examples, see (Burgis, 2019; Cheng, 2018; Linker, 2015; Stebbing, 1939; Zornado et al., 2020).

sense—one way for those arguments to be good or bad.

Part II covers the core theory of propositional logic. In Chapter 3, I introduce the formal language. I also explain how that language can be used to translate English sentences and English arguments. In Chapter 4, I give a fully rigorous account of truth in propositional logic. Then I give a fully rigorous account of what it takes for an argument in propositional logic to be good in the relevant sense. In Chapter 5, I use ideas from Chapter 3 and Chapter 4 to precisify the ideas from Part I. In particular, I use the account of good arguments in propositional logic to sharpen the account of what makes natural language arguments good. Finally, in Chapter 6, I explain how to use these ideas from propositional logic to evaluate the sorts of arguments that appear in social media, news articles, sound bites, radio shows, and so on.

Part III covers some shortcomings of propositional logic. In particular, in Chapter 7, I show that the theory presented in Part II is—though pretty decent—not ideal. That theory classifies some perfectly good English arguments as bad. To correctly classify those arguments as good, propositional logic does not cut it. A better logical system is needed.

Part IV covers the core theory of that better system: first-order logic. In Chapter 8, I introduce the formal language. I also explain how that language can be used to translate English sentences and English arguments. In Chapter 9, I give an account of what it takes for an argument in first-order logic to be good. Fortunately, this account does not require a corresponding account of truth in first-order logic; for that account of truth, it turns out, is too complex to be covered rigorously here. In Chapter 10, I use the ideas from Chapter 8 and Chapter 9 to precisify the ideas from Part I. In particular, I use the account of good arguments in first-order logic to sharpen the account of what makes natural language arguments good. And I explain how these ideas from first-order logic can be used to evaluate the sorts of arguments that you routinely encounter.

Part V contains a review of the entire book. Appendix A summarizes the content of each chapter. Appendix B contains a list of the key definitions.

It is worth making one final remark. In some of the more formal chapters of this book, it can be easy to forget why such dense technical material is worth learning. The main reason is that, as mentioned above, the dense technical material will help you evaluate arguments – and that, in turn, will make you more effective at changing the world for the better. But there are other reasons to learn this material too, reasons which also connect to social justice and political reform. So periodically, near the ends of various chapters whose material gets quite technical, I discuss some other reasons why that technical material is worth your time; those discussions occur in sections called "Applications to Justice."