

Arthur, John
Studying Philosophy
2nd Ed. 2004
Pearson

5

THE TOOLS OF THE TRADE: LOGIC AND CRITICAL THINKING

WHAT IS REASONING AND WHY DOES IT MATTER?

Much of what goes on in universities and colleges involves reasoning, whether it's developing and defending our own ideas or understanding and assessing somebody else's. We reason with one another, whether about how nature works, what society is like, or how we ought to act. And it's usually good that we do reason, considering the alternatives. Logic and critical thinking is the field that studies reasoning, in particular, how to distinguish good reasoning from bad. It is a complex, fascinating subject, and I will only scratch the surface here. But in my experience, even a brief introduction to critical thinking and logic is useful to students. It enables them, and indeed anybody who studies it, to become more aware of reasoning processes and more adept at making good arguments as well as avoiding being fooled by bad ones. It is also helpful, as you will see, in understanding precisely what is being claimed—what something means—and also in noticing when meaning is unclear. So my goal in this chapter is to give you some tools that can make you a better (more logical) reasoner, whether as a reader, discussant, or writer.

First, I want to say something about reasoning in general, and its relationship to logic and critical thinking. Universities are wonderful places: even the United States Supreme Court has held that they are special. University faculty, and students, are protected by the Constitution to speak their minds without fear of reprisal. But before you tee off against the administration's policies, you might want to keep in mind that such freedoms are not absolute. Private institutions have more authority to restrict speech than state-run ones, and there are always limits on the time, place, and manner of speech, even if what you say is otherwise constitutionally protected. Though you are entitled to express your opinions about most everything, you can't do it by spray-painting a building or disrupting a classroom. Nevertheless, both students and faculty do enjoy broad First Amendment protection, which allows us to say

unpopular and even offensive things to one another without fear of reprisal. (That same First Amendment also protects those who disagree with you, however, so don't be surprised if you get back as good as—or even better than—you give!) Universities are therefore inherently controversial and divisive places where ideas are tried out, criticized, defended, and sometimes ridiculed.

Why encourage all this ferment, debate, and arguing? The answer, of course, is that we learn from it. Testing theories about how nature works, developing new ways of interpreting the world, and offering new philosophical, historical, or other ideas about human life and societies all rely on the availability of criticisms and commentaries. We can proceed no other way, at least not very effectively. Reasoning with each other, in the broad sense of proposing ideas and theories for discussion and criticism, is the heart of a university: we can't ultimately know if what we are saying is true or good or valuable unless we are lucky enough to attract critical attention. But not all attention is helpful, of course, and that's where reason comes in. Somebody who burns a book or screams epithets at a speaker is probably enjoying the attention, but at the end of the day that's not nearly as helpful as offering thoughtful criticism based on reasons. We learn from such outbursts that the person is offended or angry, not whether the person is entitled to be.

If I am right, and free speech and debate exercised in the service of reason are at the heart of a university, then what is reason? Begin with the idea of an argument, not in the sense of things people *have* with one another, but in the sense of things that we *give* to one another. To give an argument is to give reasons for a conclusion. Reasons are given to change people's beliefs, though there are other ways to change people's beliefs as well. You may not believe it is a good idea to say you were wrong until I produce a pistol, for example. That threat, however, is not an argument—at least not in the sense that we use the term. Besides threats of force, reasons should also be distinguished from (mere) emotional appeals and rhetoric. Suppose two people give word-for-word the same talk defending affirmative action or global warming, say, but one does so with great emotional appeal and the other simply reads the speech coldly, without feeling. Has the first given any better reason to believe the conclusion? No. Will the first be more persuasive? Probably. My point is that the sense of "reasoning" I have in mind is normative, i.e., it is concerned not with what actually works to persuade people, but what should work. Reasoning is about what people *should believe, based on evidence*. That's why, although it may be easier, and even kinder, to lie and praise a friend's argument, in another sense it is less friendly. The paper has mistakes in it (let's suppose) and so you are not giving your friend advice that will help get to the truth (or perhaps a good grade, for that matter).

If reasoning is what we want to do, and inherent in reasoning is a distinction between good and bad ways to persuade people, then how do we know when reasoning is "good" or "bad"? The answer is logic. All disciplines, whether economics, philosophy, history, or whatever, share one thing. It is not

what they study, or even the method they use. What they do all share is the goal of giving arguments, which means they depend on the norms or rules that we call "rationality." Indeed, every time people try to understand or explain something, they assume that there is an important difference between good reasons or explanations and bad ones. Otherwise, why would anybody bother to study a subject? We might as well just speculate about the answer to a question, or just make up an answer. Logic is the name we give to the standards and principles of rationality and rational inquiry. So all rational inquiry depends on logic.

Why then bother to study logic? There are many reasons, but I will mention just two. First, it enables us to learn a great deal about our own language. Like all other spoken languages, English words and sentences are full of ambiguity. Different words can mean different things, as can different sentences. Translating English sentences into precise logical statements helps people appreciate the complexity of their language and to become more adept at saying just what they mean. Second, it is worthwhile to study logic because it helps people reason better. Learning and applying the rules that mark the difference between good and bad reasoning, as well as the many forms of "fallacious" reasoning people sometimes employ, is helpful in analyzing what we hear or read as well as in formulating our own arguments and reasoning.

I will begin by saying something more about the nature of an argument, including the difference between deductive and inductive arguments. Then in the next section I will discuss critical thinking by looking at the most common types of bad reasoning, usually termed "fallacies," that can often lead people to believe something when they shouldn't. I conclude the chapter with a brief introduction to deductive logic.

THE NATURE OF ARGUMENTS

Making and challenging arguments is what we do when we are reasoning. Arguments are offered or disputed in order to persuade people, and good arguments are ones that *should* persuade while bad ones should not. But what then is an argument? The answer is *statements* in the form of *premises* and *conclusions*.

Let's begin, then, with the nature of language. It helps to think of language as a box of tools enabling us to accomplish a wide variety of goals. But notice how many different things we do with language besides trying to say what is true. Sometimes we may use language to give a warning, for instance, when a friend yells "Look out!" as a car speeds toward you. Other times we use language to express our feelings, for example by saying "Ouch," or even to perform an action, as when a person says "I do" during a marriage ceremony. On other occasions we may use language to pray, or to sing songs,

or to make promises. The list goes on and on. But much of the time, when we use language, we are trying to say something that is either true or false. We are making a statement of fact. By "fact" here I don't mean just scientific fact, but anything that we want another person to believe. It could be anything from "Grass is green" and "I like coffee" to "Abortion is murder." Let's call that use of language, where we are saying something that is true or false in the broadest sense, a statement.

Arguments are statements, but not just any old group will do. They have a certain form or structure because they comprise at least one premise and one conclusion. A *premise* of an argument, then, is a statement or claim of a particular sort, namely one that is supposed to provide good reason to believe another statement or claim, the *conclusion*. So an *argument* therefore is at least two statements: a premise and a conclusion.

But how do we know which is a premise and which is the conclusion? The answer depends on the meaning and the context. Premises are often shown by *premise indicators* such as "because," "for the reason," and "since." Conclusions can often be identified by *conclusion indicators* such as "hence," "therefore," "it follows that," or even "in conclusion." But those indicators are not always there, and so in identifying arguments and their structure there is no substitute for reading, listening, and thinking carefully about what is meant. There is therefore lots of room for argument about the nature of arguments themselves.

Once we notice arguments, we see that they surround us. We think about, analyze, are given, and give arguments all the time. So, for instance, I might say that:

It rained just recently. I know because the streets are wet.

The conclusion of the argument is that it rained recently; the premise is that the streets are wet. Sometimes the premise or premises are stated first, but sometimes later on, as in this example. You can't begin to understand an argument until you know the meaning of what is said, including which statements are premises and which is the conclusion. That's the first step.

But notice, also, that there is another, unstated premise in that argument I just gave, namely that the streets being wet makes it certain, or at least more likely, that it has rained recently. That distinction, between arguments claiming to justify a conclusion beyond doubt and ones that only make the conclusion more likely to be true, is the difference between deduction and induction. We will take up those issues shortly.

So we often give an argument without explicitly stating all the premises. We even sometimes don't state the conclusion, as in:

The movie will be boring, and anyway I am tired.

Here the missing statement is the conclusion that we should not go to the movie.

Arguments are also sometimes presented in a linked series, like a chain. This is almost always the case in longer passages and whole essays. In that case, the link between premises and conclusions is not direct, like the one about wet streets, but instead involves intermediate steps. Consider, for instance, this argument:

Global warming is a problem. We can see from our own experience this summer that the planet is getting warmer. Because of the heat there is drought and, ultimately, starvation.

This argument is more interesting than the rain one. For one thing, it shows clearly what I said a moment ago—that the first step in understanding an argument is to understand what is meant by its different parts. We often have to read and think carefully about language in order to even recognize an argument, let alone assess its merit. The ultimate conclusion in the global warming argument—what we are asked to believe—is that global warming is a problem. Two premises, according to the argument, justify the conclusion: that the planet is getting warmer and that increased temperatures cause drought and starvation. But notice also that embedded in the argument is still another, subsidiary argument: that we have reason to believe the planet is getting warmer (intermediate conclusion) because we have had a hot summer. In other words, the statement that we have reason to believe the planet is getting warmer does double duty: it is both a premise, supporting the ultimate conclusion that global warming is a problem, and a conclusion that is based on the premise that the summer has been hotter than usual. As with the first argument, here too we can see that there are unstated premises. Making those explicit is again essential in understanding and ultimately in assessing an argument. So here is an interpretation of the arguments embedded in those three sentences. (There could be others, and you might want to think about whether or not you agree with my interpretation.)

First argument:

Premise: This has been a hot summer.

Intermediate conclusion: The planet is getting warmer.

Concluding argument:

Premise 1: The planet is getting warmer. (This is also the conclusion of the first argument.)

Premise 2: Warming of the planet is causing drought and, ultimately, starvation.

Final conclusion: Global warming is a problem.

So what have we said so far? That arguments are defined as at least two statements, one of which is a premise and the other a conclusion; that statements are forms of language in which we make a claim that something is true, or false; and that premises and conclusions are often identified by indicator words, such as “because,” “since,” and “therefore.”

Keep in mind that the logical test of an argument depends on the relationship between the premises and the conclusion: Do those premises, if true, give good reason for accepting the truth of the conclusion? There are two types of arguments because there are two forms that this “good” relationship between premises and conclusion can take. Sometimes the premises, if true, give good but less than conclusive evidence for the conclusion. These arguments are termed inductive ones. They rely on patterns and other similarities to give us reasons to believe what we might otherwise not be able to know. So, for example:

Premise: Most lakes in this region are polluted.

Conclusion: Lake James is probably polluted.

(This argument has a hidden premise: that James is a lake in this region.)

But notice that though this is a good argument another type of argument would establish the same conclusion with greater certainty, assuming its premises were true.

Premise: All lakes in this region are polluted.

Conclusion: Lake James is polluted.

The first type of valid argument, where the conclusion does not follow conclusively and with certainty, is termed “*inductive*” while the second type, where the premises guarantee the conclusion, is “*deductive*.” If the premises of a deductive argument are (or were) true, then the conclusion must (or would) absolutely have to be true. With inductive arguments, however, the premises only make it likely or probable that the conclusion is true. Inductive arguments can therefore be *stronger* or *weaker*, while deductive ones are more like a light switch or a math problem: They are either right (*valid* is the term logicians use) or they are wrong (*invalid*). Unlike inductive arguments, then, in deductive ones the conclusion does not “go beyond” the premises.

Unlike deductive ones, inductive arguments, even if they are very strong, leave open at least the theoretical possibility that the conclusion is false, while the premises are true. My belief that you will not win the 100 million dollar lottery with your one ticket is very well justified inductively. But the trick with inductive arguments is to know how strong they are, that is, the degree to which the premises justify the conclusion. The fact that I have won five con-

secutive bets at the race track by itself gives little reason to think I will win the next bet. You'd want to know, for instance, whether I am a novice who just got lucky or an expert who has found a good way to predict a horse's success.

As you may have guessed, inductive arguments are common, and come in many forms. Here are four of the most typical types of inductive arguments.

Generalizations, in which, for instance, we infer from a sample of people what the attitudes of all the people in the group are.

Predictions, of things like weather.

Causal inferences, in which, for example, we conclude that cancer is caused by smoking.

Analogies, in which we say, for instance, that abortion is acceptable because it is like getting something back that you own, your body, from somebody who has gotten to use it without your permission.⁴

Those examples can be generalized, so here are more abstract, formal examples of the four types of inductive arguments:

Generalization:

Premise 1: Very few As are also B.

Premise 2: x is an A.

Conclusion: So x is probably not B.

Prediction:

Premise: In the past, when I have done X then Y has followed.

Conclusion: If I do X now, Y will likely follow.

Causal inference:

Premise 1: More people in this area have characteristic C than is normal.

Premise 2: People in this area have been exposed to D at higher than normal levels.

Conclusion: Exposure to D causes C.

Analogy:

Premise 1: Everything I have examined with characteristics A, B, C, D, and E is also a Q.

Premise 2: t has characteristics A, B, C, D, and E.

⁴Judith Jarvis Thomson uses that analogy in Judith Jarvis Thomson, "A Defense of Abortion," *Philosophy and Public Affairs*, 1, no. 1 (1971), pp. 47-66.

Conclusion: So t is also probably a Q.

These arguments will be either strong or weak, depending on the specific details of argument that are given. Probability theory and statistics are the branches of mathematics that are central to inductive reasoning. In standard cases of inductive generalization and other uses of statistical sampling, for instance, the larger the sample (and the more randomly it is chosen) the better. Returning to the race track example, the more races I have predicted, the better the evidence you have about my next bet being right. Similarly, asking a small number of people who they voted for is not as good a predictor of an election as asking a large group; asking people only in a heavily Democratic district is not as good as asking a cross-section of the population.

Another important issue in inductive logic involves causal inferences. We need to know, sometimes desperately, whether something in the environment is causing illness, or whether it is just a "statistical accident." Statistical accidents do occur: Imagine taking a bag of pennies and dumping them on the floor. It's virtually certain you would find some of them grouped together. So when residents of a town have a higher-than-average incidence of cancer, there are always two possible explanations: Either something we want to know about is causing the higher incidence, such as water pollution, or it's just a statistical anomaly, like the coins bunched together. Similar issues arise in trying to determine when discrimination is present. Is the higher death rate among babies of one particular ethnic group due to prejudice and therefore poorer care given by doctors and nurses, or to other factors, such as mothers' eating habits or drinking and drug taking?

Finally, there is also a distinction between higher and lower "levels" of induction. The fact that this computer has worked well for some time would normally be a good reason to think it will continue to work. But at some point, because we know from many other cases that computers eventually wear out, the fact that this one has worked for a long time begins to work *against* the belief it will continue to work. We know from looking at more cases over a longer term that this one, too, will eventually wear out. That higher level of inductive generalization allows a more refined judgment about an individual case. So, again, we see that probability theory and statistics are important to anyone who wants to do inductive reasoning well.

Some speculate that human beings use induction instinctively. Induction may well have even an evolutionary basis, since our ancestors who relied on it presumably lived to reproduce others who, like themselves, were good at induction while those who were not so good were also not successful at reproduction. Learning to recognize and run from dangerous animals meant they lived to produce offspring with the same talent. But whatever induction's origins, it's hard to imagine humans surviving long, let alone prospering, without the ability to learn from past experience and to generalize from that experience about what will happen in the future.

INFORMAL FALLACIES

There are two ways to look at arguments from a logical point of view. You can study them formally, to see if they are strong (if inductive) or valid (if deductive). I've said a little about induction and will have more to say about deductive logic shortly. Another way to study arguments, besides looking to see if they follow good statistical or other formal rules, is to look at some of the ways reasoning typically goes off the tracks. The name given by logicians to the mistakes in reasoning we typically make is *fallacies*. So although there are probably about as many fallacies as there are people who reason, there are some forms of bad argument to which human beings are especially vulnerable. These are the subject of this section.

The point of studying fallacies is similar to the reason it is useful to think about the nature of arguments and probability theory. By paying close attention to language and good reasoning patterns, we can become better reasoners.

Fallacies can be divided into groups. Sometimes the premises, even if true, do not support the conclusion. The premise may be either completely irrelevant to the truth of the conclusion, or though relevant there still may be a crucial "gap" between the premises and the conclusion that makes believing the conclusion unreasonable given the premises. Other times the argument is unacceptable for other reasons. As you read through the different types of fallacies, think about why they are fallacies and also about when an argument similar to that one might not be bad or fallacious.

Ad Hominem. Ad hominem means "to the person" or, literally, "to the man." The fallacy of ad hominem takes place when, instead of attacking a person's position or argument, the arguer instead attacks the *person* on irrelevant grounds. This fallacy often involves questioning a person's past activities or associations, or needlessly referring to the person's background or interests. Sometimes such appeals carry great emotional feeling, but unless the personal trait is somehow relevant to the argument being offered, an ad hominem attack is not relevant to the conclusion. An obvious and therefore not very persuasive example is this: "Clinton's personal moral behavior was despicable, so his welfare reform proposals, which punish people who don't work, couldn't possibly make sense."

Appeal to Ignorance. This fallacy occurs because we often assume that since somebody has not proven something is true, it must be false (or vice versa). This is a common mistake. The fact that you cannot prove God exists does not mean He doesn't exist, just as the fact we have not found out all the details of how evolution works does not show the theory itself is false. But that said, it is often interesting and controversial to determine exactly when it is a fallacy to appeal to ignorance. Sometimes the fact that something has not been proven true *should* be regarded as evidence against it. The fact that a

good mechanic has gone over the car with a fine-toothed comb and found nothing is, presumably, good evidence the car is safe to buy. But the fact that we don't know exactly how cancer is caused by some chemical is not good evidence that the chemical is not the cause. It might be the cause, and we just haven't figured out how it works yet.

Appeal to the People. This fallacy can take many forms, but the basic idea is that the arguer fallaciously appeals to the popularity of an idea or of an attitude as a justification. "All my friends are doing it" is a child's version, but adults are also often tempted to believe what others believe when that popularity is insufficient to justify the belief. The bandwagon effect is what we sometimes call it, and it is common. It is interesting to note in this connection that voting and majority rule are not so much methods to establish a truth, as a decision procedure that is accepted because it is fair. The truth about the world is not usually decided by a vote.

Appeal to Force. This fallacy occurs because the premise, that somebody is powerful, does not warrant the conclusion that the person is right or the belief is correct. "Your money or your life" may give good practical reason to turn over your wallet, but it is not a good reason for concluding you should do it. "If you don't go to class, you will fail" is again a good reason in the prudential sense, but it does not justify the professor's reasonableness in making the student attend.

Straw Man. Here the argument is fallacious because the arguer is attacking a position that is not, in fact, the one that the opponent is defending. So the attack is irrelevant. The author has in effect created a weaker, straw figure and merely attacked *it*, leaving untouched the original argument that was supposedly being criticized. But sometimes people think they have found a mistake in another's argument when, in fact, they have not interpreted the argument correctly, i.e., they have not given the author the benefit of the doubt or have overlooked a more reasonable interpretation that does not involve the writer in a mistake. So the straw man fallacy and the principle of charity are, in that sense, related: people who commit the fallacy are often not being charitable in interpreting what others mean. This is a pretty common fallacy, though identifying actual cases is not always easy, and it may be controversial whether there really is a straw man. Sometimes people do give bad arguments, after all.

Over-Generalization and Susceptibility to Counter-Examples. Everybody generalizes when they think and write, and sometimes it's perfectly acceptable. But careful reasoners avoid reaching conclusions based on small samples, stereotypes and the like. Some generalizations are "empirical," i.e., they are about the world and how it works. Others are about language, ideas, rules and other standards. It's probably safe to say, based on available

information, that "The British people supported the Allied effort in World War II." However, if an author says that "U.S. youth hated the Viet Nam war" there should be some pretty strong reasons to justify such a generalization. The same is true of statements like "Philosophers are skeptical about the objectivity of morality" and "Welfare recipients don't want to work." While such statements are undoubtedly sometimes true in individual cases, it's dangerous to make such sweeping claims without good evidence.

There are different ways to question a generalization. If the generalization is thought to apply without any exceptions, one *counter-example* can serve as a refutation. (There is a famous case in which a philosopher was giving a formal talk at the meetings of the American Philosophical Association. He told his audience that while it is common in English to use a double-negative to mean a positive—it is "not unlikely," for instance—you never, he said, hear a double positive used as a negative. Another philosopher in the audience then yelled out, "Yeah, yeah.") We often use counter-examples to attack claims that are conceptual and moral. If it is suggested, for instance, that only rational human beings can have rights, you might suggest the counter-example of infants, or perhaps even Mr. Spock. (If a philosophical generalization is conceptual, as this one about rights seems to be, then it can be useful to think up imaginary cases as well as real ones.) Or suppose a theorist claimed that all laws are rules, or that all rules are general commands. So thinking up a counterexample to a generalization is one of the most familiar ways to dispute it. If the generalization is only that, a generalization, and is not meant to be universally true, then of course merely finding one counter-example does not refute it.

Gamblers' Fallacies. These fallacies are similar to over-generalizations. While common among gamblers, they are also found in business and everyday life. One is the familiar idea that having lost the bet the last five times in a row, "My number has got to come up next time." The problem, of course, is that whether the number comes up next time is causally and statistically independent of what happened in the past. A flipped coin still has a 50/50 chance of being heads, whatever it's been in the last few (or few hundred) tosses. Whether your car will make it to California may not, however, be independent of how it has been doing the past year. So whether or not the future is able to be predicted from past experience is itself a question that needs further thought; you can't just assume things will be the same, or will change. Another gamblers' fallacy involves the belief that if we increase our bet each time, then eventually we must win. This is not true, as you can see if you calculate what will happen if somebody starts out betting one dollar, then after losing the first time, bets two, then three, then four, etc. On the other hand, *doubling* your bet after each loss will ensure that you eventually win (as long as you don't run out of money!).

Throwing Good Money (or Time) After Bad. Here's another favorite of mine, sometimes referred to as the fallacy of sunk costs. Suppose you have already spent five dollars on a ticket, but have since learned the movie is no good and you will probably not enjoy it. Many people will say, "I already spent the money, so I might as well just go." But that's a mistake. If you would rather do something else, like go to a different movie, you should not compound your loss and pass up something you prefer just because you have made a mistake. It's better to admit the mistake and do something you would actually enjoy. This same idea, that we have something invested and so we should stick with it, can also occur in other contexts—when we have spent time on a project, for instance. It applies to students, too. If you have written a draft of a paper that has not panned out, don't think you must keep it. Sometimes it's best to start again. On the other hand, sometimes you can save even a bad paper by showing that what seemed a plausible idea didn't work out. Blind alleys are sometimes worth writing about, after all.

Questionable Analogy. Arguing by analogy involves asserting that because two objects, situations, or events are similar in certain respects, the two must therefore share a further characteristic. Bad analogies are uncovered by pointing out that the two things are not in fact alike in the relevant ways, or that even if they were alike that doesn't prove they share the further property. One famous example of analogical arguing (though it is much debated whether it is a fallacious analogy) used in an article that I mentioned earlier involves a person who is kidnapped and hooked up to a violinist who needs the kidnapped person's kidneys for nine months. The author, Judith Jarvis Thomson, used the analogy to show that because the violinist has no right to use the other person's kidneys, an unborn baby has no right to use a mother's body either, and so the mother can get an abortion.

Inconsistency. People are often inconsistent in some way, holding onto a premise or conclusion in one context that is incompatible with ones accepted in a different context. (This is a fallacy involving mistaken premises because inconsistent premises cannot both be true. Something has to give.) Rarely are people clearly and obviously inconsistent, though in a more subtle way, it's always a danger. We may say something in one context, writing about one issue, that we would not be willing to say in another, not having thought about the latter context. Or we may use a term in a subtly different sense in an argument, making the argument appear stronger than it is. Which leads to the next category.

Equivocation and Ambiguity. Sometimes an argument depends on the author's equivocation between two different senses or meanings of a word or phrase, but you should reject the argument, because a key term or concept is used in one way in one context but another way elsewhere, thereby making a

key premise false. So, to use a simple example, you should not accept the following argument: "People who own banks are probably quite wealthy, and Sue owns the bank next to her river, so Sue is probably wealthy." Other examples are not so silly, though; people often confuse themselves and others with imprecise words used in different ways in different contexts. You can imagine, for example, that a person might accept the idea that lung cancer kills people and that smoking causes lung cancer yet not know or appreciate sufficiently how smoking kills.

Begging the Question. This fallacy occurs when, instead of giving a reason that establishes the conclusion, the conclusion is simply assumed. It is a fallacy because having assumed the (doubtful) conclusion as a premise, the arguer has not given a good reason to accept the conclusion. Often this occurs when a writer uses words that, on careful inspection, merely assert in different terms what's supposed to be proven. Here's an example: "God exists, because the Bible says so. And the Bible must be true, since it's the word of God." You can sometimes see that a paper includes this mistake if a critic who reads it can reasonably say, "Well, of course, if I accepted *that* as a premise, then I'd have to agree with you. But that premise is just another way of stating what you want to prove. And so you haven't proved anything." Politicians and political writers are especially prone to commit this fallacy when they evade or don't answer a question. In doing so, they assume what is at issue, and so are guilty of begging the question. The fallacy also can occur in a longer argument chain, because of the distance between the premises and the conclusion, or where people ignore crucial but controversial or even very doubtful premises and simply assume they are true.

The phrase "begging the question" has come into popular usage recently, but it is often used to mean something different. Sometimes people will say that someone "begs the question of . . ." and mean only that they have "raised the question of. . ." I have heard logicians get upset with that usage, claiming it is not the right way to think about begging the question. But since words' meanings can change, there may be little point in complaining. Still, I do agree that the logician's use of the term, the fallacy of begging the question, is also an important idea to keep in mind.

Slippery Slope. This is just what it sounds like. Slippery slope arguments assume that if we take a first step, or two, then inevitably we will be led down a slippery slope to someplace we don't want to go. Because we reject the end point, it is claimed, we must therefore reject the beginning. This "slope" can be either logical or practical. It could be logical if the idea is that anybody accepting some argument would then logically have to accept a further, obviously unacceptable conclusion as well. Or else you might think that were we to follow this advice we'd be led down a path that would have disastrous practical or policy consequences. Sometimes slippery slope arguments are good,

but often they are fallacies because the author simply assumes, without justification, that the first step will lead down the slope. To assess such arguments, it is important to think carefully about whether, and why, the first step will supposedly lead to the next and the next. If you unreasonably assume it does, you are guilty of the fallacy.

False Dilemma. In this case the fallacious arguer unreasonably reduces the alternatives by arguing, for instance, that there are only two possibilities when, in fact, there are others. You can criticize a false dilemma most easily by pointing out the missed alternative, i.e. by showing that it really is not a dilemma at all. Sometimes it is assumed, for example, either that God exists and is the author of the moral law, or else there is no morality and everything is permitted. But there are alternative justifications of morality, besides divine commands, which shows that (at least without further evidence) this is a false dilemma. Another example might be those who assume human free will is a myth or else our choices and decisions are uncaused. It could be argued that this is a false dilemma, since we could be free and at the same time all of our actions could be caused by prior events.

Red Herring. This fallacy gets its name from the fact that red herring fish were used to throw dogs off the scent when they were trying to track a person. The fallacy occurs when the arguer diverts attention away from the real issues onto an irrelevant side issue. It differs from the straw man fallacy because there the opponent's argument is misinterpreted while in red herring the arguer changes the subject, ignoring the opponent's argument altogether.

False Accusation of a Fallacy. It's frequently controversial whether or not a fallacy is really present, as I have suggested, so sometimes people even add "False accusation of a fallacy" to the list of fallacies. Be sure, therefore, to think about the possibility that you may have committed the mistake of falsely accusing the writer of committing a fallacy. Maybe the slippery slope fallacy is not a fallacy in this case, and the slope really is slippery. Or maybe your claim to have identified a doubtful analogy or straw man can itself be questioned.

So it's critical, both in assessing arguments and in giving them, to pay close attention to the various fallacies that might be committed. There are many types, and the ones I mentioned are only the most common. Finally, here's a very helpful exercise. Take the list of fallacies I have given, and go through a newspaper or magazine to see if you can find examples of the fallacies. Alternatively, listen to some advertisements or to a political debate. Write down the argument (the premises and conclusion) being presented, and then ask yourself which, if any, of the fallacies I have described is being committed here. Doing that will make you a better reasoner and therefore, I am con-

fident, a better student.

DEDUCTIVE LOGIC

Fallacies of the sort we have been discussing are sometimes thought of as occurring in “informal” logic, as opposed to “formal.” Formal logic gets its name because the correctness or incorrectness of the argument can be identified by the form of the argument alone. Informal fallacies don’t work that way. There the mistake is not formal, but depends instead on the specific content or meaning of the premises and conclusion, whether or not language is used ambiguously, whether there are irrelevant appeals to force or ignorance, and so on.

In this last section, I want to discuss briefly some of the central themes in formal, deductive logic. As I said, it is concerned with argument forms alone, not their particular content. To see what that means, notice that different statements, while they mean different things, can have the same logical form. “Some doctors are males” and “Some horses are brown,” for instance, have the same logical form. So do “Professors are teachers” and “Cats are mammals.”

Validity and Soundness

A *valid* deductive argument is one whose premises, if true, establish the conclusion with certainty. It is impossible, in other words, for the premises to be true and the conclusion false. Sometimes it is obvious if an argument form is valid or not. For example, it’s pretty clear that the following is formally valid. We don’t need to know what A and B mean to know that if the premises are true then the conclusion must also be true.

Premise 1: If A then B. (For example: If Sam is a cat then Sam has a lung.)

Premise 2: A. (Sam is a cat.)

Therefore: B. (Sam has a lung.)

It’s also pretty obvious that this similar argument form is *not* valid, because there are plenty of instances where the premises are true but the conclusion is false.

Premise 1: If A then B.

Premise 2: B.

Therefore: A.

If A stands for “Sam is a cat” and B for “Sam has lungs,” then both the premise and conclusion are true but, because Sam is in fact a dog, the conclusion is false. That single example shows the argument is not formally valid, even

though there are many other examples where the premises are true and the conclusion is also true.

But some arguments are less obviously either valid or invalid, and showing which is a complex problem. As I discussed earlier in the chapter, there are also argument “chains” in which premises establish a conclusion, but then that conclusion in turn provides premises for more arguments that eventually lead to the ultimate conclusion. Formal logic is concerned with seeing if there is a valid route to a conclusion from a given set of premises.

It is also sometimes tricky to distinguish deductive from inductive arguments. Recall that inductive arguments do not establish their conclusions with certainty; they only show the conclusion is more likely true, given that the premises are true. Even the fact that an argument uses “most” or “some” is not necessarily an indication that it is inductive. Consider the following argument:

Premise 1: Most doctors are wealthy.

Premise 2: Most doctors are conservative.

Conclusion: At least one doctor is wealthy and conservative.

This argument has the form:

Premise 1: Most As are Bs.

Premise 2: Most As are C.

Conclusion: At least one A is both B and C.

Despite using “most,” a term usually found in inductive arguments, this is a deductively valid argument because the premises, if true, establish the conclusion with certainty, i.e., it is impossible for the premises to be true and the conclusion false. An invalid deductive argument would therefore be one in which it is claimed that the premises establish the conclusion with certainty, but they do not.⁵

But it is also the case that any of the valid forms can have false premises and a false conclusion and still remain valid. So, for example, the following argument is valid even though its conclusion is not, in fact, true:

Premise 1: If a candidate won the popular vote, he won the presidency.

⁵So it should be clear from what’s been said that the common idea that induction involves going from the particular to the general and deduction from the general to the particular is mistaken. Though some forms of inductive inferences do that, others don’t. And though it is sometimes true that deductive arguments go from the general to the specific, it is not always so. “All men are mortal and no mortals are supernatural, therefore no men are supernatural” is a valid deductive argument.

Premise 2: Al Gore won the popular vote.

Conclusion: Al Gore won the presidency.

What's wrong with the argument is not that it is invalid, but that at least one of its premises is not true. *If* it were the case that candidates who win the popular vote win the presidency, then Gore would, indeed, be president. The problem is therefore not with the logic; it's not that there is a "gap" between the premises and conclusion. Rather the problem is that the one of premises is false. Logicians refer to it as deductively *valid, but unsound*, to indicate that distinction.

The reverse is also true: you can have an invalid deductive argument form with true premises and a true conclusion. Here is an example:

Premise 1: If it had rained, then streets would be wet.

Premise 2: It didn't rain.

Conclusion: The streets are not wet.

This argument has this form:

Premise 1: If R then W.

Premise 2: Not R.

Conclusion: Not W.

But this is not a valid argument form, even if—as in the interpretation of it involving streets and rain—the conclusion and the two premises are all true. (For example, a fire hydrant could have been on.) Or consider the following argument with the same form:

Premise 1: If Felix is a dog, then Felix has lungs.

Premise 2: Felix is not a dog.

Conclusion: Felix does not have lungs.

If Felix is a cat, then though each of the premises is true (If Felix were a dog, he'd have lungs and Felix is not a dog) we cannot therefore conclude Felix does not have lungs since Felix is a cat.

Categorical Statements and the Square of Opposition

It is useful to begin studying deductive logic with the four classic forms of categorical statements. These express the logical relations among classes or groups of things, which is why they are unsurprisingly termed "categorical" statements. So imagine that As are any class or category of things whatso-

ever, from horses, people, or red herrings to sunny days, electrons, or things on your desk. Bs, let's also assume, are another class. Also assume that "Some" means at least one, though it could include more than one. There are four categorical statements, with the traditional letter-names of A, E, I, and O.

All As are Bs (called form A).

No As are Bs (form E).

Some A is a B (form I).


Some A are not a B (form O).

These can each be put into diagram form by imagining circles or Venn diagrams where shading indicates there is nothing there, a capital A, B, etc. refers to the particular class of things, and lower case x refers to an individual member of the class that exists. Here is how the four are diagrammed.


(A) All As are Bs. (For example: All apples are fruits.)

(E) No As are Bs. (No apples are pears.)

(I) Some A is a B. (Some apples are red.)



(O) Some A is not a B. (Some apples are not red.)



The differences between these four can be expressed in terms of their *quality*

