CHAPTER 1

The Parts of Public Thinking: Deductive Argument

Here is Shakespeare's Macbeth. After first being told by three witches that he will become both thane of Cawdor (a title equivalent to earl) and king of Scotland, Macbeth discovers that he has in fact been made thane of Cawdor. He muses to himself,

> This supernatural soliciting
> Cannot be ill; cannot be good: if ill
> Why hath it given me earnest of success
> Commencing in a truth? I am thane of Cawdor:
> If good, why do I yield to that suggestion
> Whose horrid image doth unfix my hair
> And make my seated heart knock at my ribs,
> Against the use of nature? (*Macbeth*, I, iii, 130–37)

Or, as one might say less elegantly but more accessibly:

> "Is this prediction by the witches good or evil? I can't decide. If it's evil, why has it foretold successes that have already started to appear, like my becoming thane of Cawdor? If it's good, why does it already have me thinking about murdering the king so I can take his place, even though the thought disturbs me deeply? I just don't know what to make of it all."

Macbeth in this passage is doing what we all do from time to time: taking a tangled collection of thoughts that weigh for and against some idea, and trying to lay them out in a coherent and easily surveyed form.

He's doing it by talking to himself, a practice unjustly frowned upon in some circles as a sign of being either slightly crazy or a professor. But every time you write something down for the purpose of making a problem clearer, in solving an arithmetic problem, for example, or sketching the outline of an essay, you are talking to yourself in more or less the same way. You are figuring out what it's reasonable to say, what it's reasonable to think, by putting your best reasons out where they can be seen and marshalling them into a form that clearly shows what follows from what.

Like many endeavours, this is much easier to do badly than well. For example, those of us lucky enough to be able to walk had to learn how to do it at some point, though we rarely remember the process. Walking is actually a difficult and fragile series of controlled falls from one foot to the other, and for a toddler just getting the hang of it, there's no shortage of ways it can all end in tears. But we manage to make a habit of it, eventually. We even adapt to difficult walking conditions—gravel, ice, cow pastures, stepping stones—by slowing down and moving more gingerly. When we're walking, we automatically recognize circumstances that call for greater care.

Reasoning too is easy to do badly and hard to do well. Unlike walking, we tend not to get much better at it when we are just left to our own devices. When it comes to skating or walking or aerial gymnastics, it's obvious if you're systematically doing it badly: if you wake up in the hospital, something has gone wrong. The same is not always true of critical thinking. There are many, many ways that errors in reasoning can be reinterpreted, overlooked, or rationalized away. So we tend not to be very effective at spontaneously minimizing our ways of reasoning poorly. It takes some special effort to gradually acquire the habit of pulling off those fragile manoeuvres called "good reasoning," or the closely related habit of automatically slowing down and being more careful when you encounter the rational equivalent of a cow pasture or slippery stones.

To acquire these habits is to become a more critical thinker. And a first step in this process is learning how to think clearly about thinking itself. That's where we are starting. This chapter introduces some basic concepts by which reasoning can be categorized and evaluated, and its various roles clarified. We'll begin by reflecting

FIGURE 1.1

Rational Slip-Ups Are Easy in Some Cognitive Terrain

Careful habits in thinking and reasoning are the cognitive equivalent of walking carefully when you know you are crossing a stream with tricky stepping stones. You need to pay attention to avoid the hazards and reach a successful conclusion.

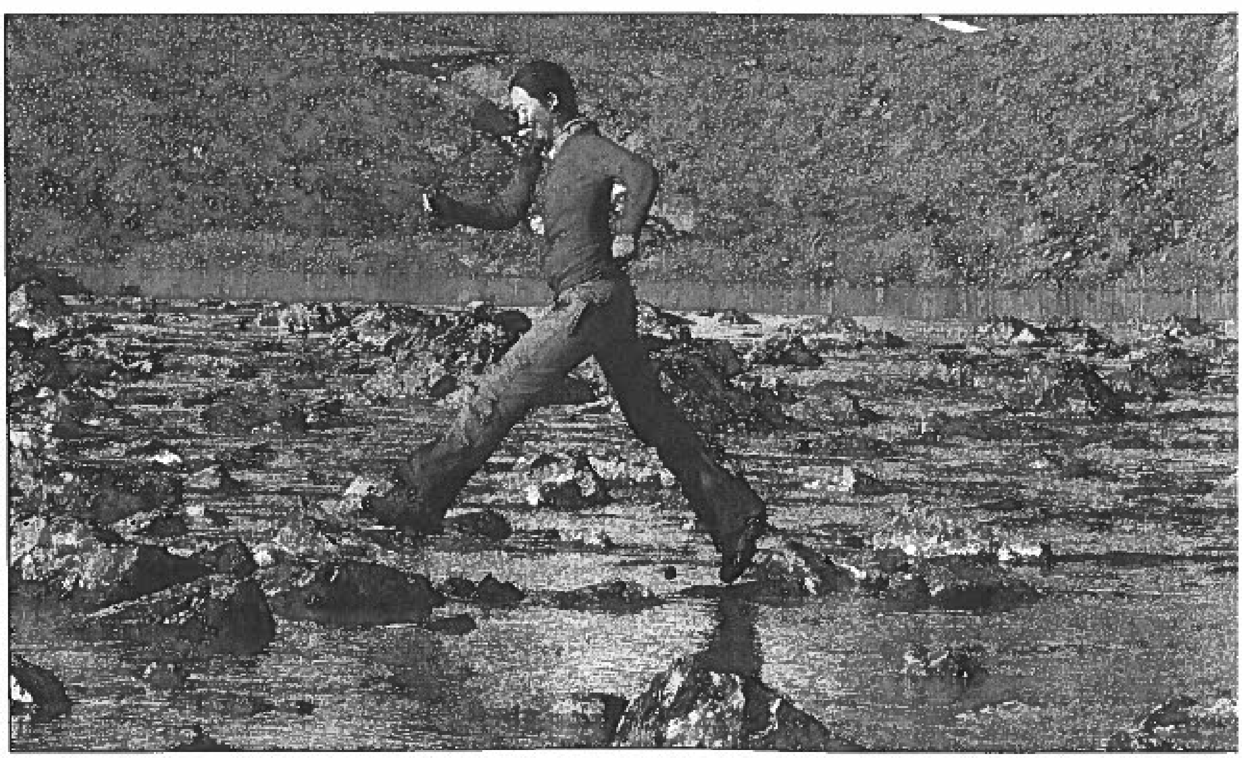

on the *activity* of putting reasoning on public display, as Macbeth did, through the presentation of the fundamental unit of reasoning: the argument.

The rather specialized concepts relating to arguments that we will consider in this chapter involve some fairly powerful idealizations: for example, they treat some fuzzy matters as if they were sharp, in a way that will become clear as we move along. They are also some of the concepts most central to formal logic. Why are we starting this way, you might ask, when the point of this book is not to do formal logic, but to take the blurriness of the world seriously rather than idealizing it away?

There are three reasons. First, the idealized concepts are simpler and easier to get a handle on. We can complicate matters later. Second, lots of reasoning really is quite properly evaluated in accordance with the more idealized concepts. And third, starting off with a more idealized approach to arguments enables us to understand and evaluate non-ideal arguments (the sort much more likely to be encountered) by judging how closely they approximate the ideal. Once we have the basic concepts in place, we can work toward a more general approach to a sometimes clear but often blurry world: learning how to factor in and take seriously the subtleties and nuances of reasoning.

The main focus of this first chapter, then, will be on deductive arguments—the sort of arguments that are logically airtight (or are intended to be so, at least). With the relevant deductive concepts in place, we can then move on to discuss the more common, complex, and subtle arguments in the following chapters.

HAVING REASONS

An overarching goal we all share, presumably, is to hold reasonable beliefs. But what is it to believe something for a good reason? On the one hand, we sometimes say that an event happened for a reason and mean by this only that it had an identifiable cause, that it was not just random or inexplicable. On the other hand, there is a stricter use of the word, meaning that an action or a belief has a **justification** or **warrant**: in other words, it can be rationally defended on the basis of evidence. Not only are there reasons for our beliefs, in the sense that there are explanations for why we come to hold our beliefs, but we moreover *have reasons* for believing what we do. There are reasons why only one side of the moon faces Earth, but the moon does not have reasons for doing what it does. Unlike the moon, people are often in a position to offer a rationale for holding their beliefs or performing their actions.

In general we take ourselves to have reasons for our beliefs. For the most part we must be right, because so many of our beliefs are truisms like "The ground gets wet when it rains" or "Cars have four wheels." Our reasons for believing things like this are so obvious that we virtually never need to articulate them. Nor need we reflect on our justifications consciously. Having reasons for such beliefs is clearly not a matter of continually thinking about either the reasons or the beliefs themselves. It seems to be enough that we could, if necessary, call such reasons to mind. For many such beliefs, no such need will ever arise.

It is important that we could provide them if the need arose, though. From time to time, a belief that seems basic and straightforward to some people will turn out not to be shared by others. Sometimes this may not be worth bothering about. It just seems obvious to me that potatoes with gravy taste better than plain boiled turnips, but if you disagree, that's no problem. (In fact, you can have my turnips.) But under many circumstances, it will be important to find out whether one (or both) of the competing claims is a genuine mistake. The way to pursue an answer is for all sides to express their reasons for their beliefs. For the purposes of settling disputes, comparing explanations, and teaching new information, it becomes essential that we are able to make our reasoning *public*.

SERTIONS AND ARGUMENTS

The most basic kind of communication for each of these purposes is the practice of presenting statements as true. Typically the point of saying something we think is true is to have other people also accept that it's true. Or, even if they don't agree, at least it is useful and enlightening to have their disagreement explained, so we can see where they part ways with what's been said. To present some claim as if it were true is to **assert** it. Whether in a classroom, in the workplace, in social environments, or through the media, a great deal of our communicative exchanges consist of **assertions**, as we go about telling one another the facts as we see them. "It was a bad idea for Athens to invade Sicily during the Peloponnesian War." I think that's true, and I'd be fairly pleased for you to think so as well.

If issuing assertions were the full extent of such communication, though, there wouldn't be much point to it. If all we did in such contexts was to declare this or that to be true, then you'd say your piece, and I'd say mine. And if we disagreed? "No, no," you'd reply, "it was a perfectly good idea for Athens to invade Sicily." Then it would be my turn to repeat my assertion, and we'd both just keep repeating ourselves until we got bored or hungry.

Of course there is much more to it. To assert is, among other things, to undertake a kind of obligation: the obligation to *defend* or *retract* the assertion in the face of questioning or when confronted with evidence to the contrary. For this reason the fundamental units of rational exchange are not assertions, but **arguments**.

The word "argument" is often used to mean a disagreement or even a fight. "We had an argument," one might say, without meaning to suggest that anything particularly involving reasonableness or rationality took place. But there is also a common and natural sense of the word that means *the presentation of reasons*. That's the sense that we're refining for the purposes of explaining and analyzing critical thinking. An argument, we will say, is a set of statements that are presented as true and that have a very important internal relation: some of the statements are **premises**, which are intended to provide rational support for (and ideally to establish the truth of) a further statement, the **conclusion**. An argument is premises given in support of a conclusion.

The property of an argument that succeeds in supporting its conclusion is *soundness*. Strictly speaking, a sound argument proves that its conclusion is true. This property can be broken down into two sub-properties of the argument: first, that it is *valid*, and second, that it has all true premises. Like the term "argument" itself, "sound" and "valid" have ordinary uses that can vary widely. Below we will define them in more specialized ways that let us speak precisely about the virtues and failings of argument.

The overall picture, then, is that assertions rest on arguments that support them, usually implicitly, but that could be made explicit if necessary. It's somewhat analogous to a simple currency system, in which people exchange paper money instead of gold or goods because they're confident they can swap the paper for the goods when they need to. The idea of the gold or the goods, available if requested, underwrites the exchange of paper money and coins; the idea of reasoning and arguments, available if requested, underwrites the exchange of assertions. Speakers are not normally taken to be making wild guesses when they assert. The idea is that a speaker has an argument tucked in her pocket that can be pulled out if necessary; she could produce at least the bare bones of reasons that supported her assertion.

Moreover, just as we saw with Macbeth and his dilemma, it's important to understand that one of the people who might be informed and affected by some argument you produce in support of an assertion is *you*. One of the central themes of this book is that critical thinking begins at home, that its immediate application is in understanding how your own reasoning succeeds and fails. Quite apart from the effects of your argument on others, you may change your own mind in various ways by going through the process of mounting an argument. You might convince yourself of something altogether new, or you might strengthen or weaken some opinion you'd already held. You might even convince yourself that one of your beliefs should be rejected. The process of making reasons public is not just about persuading other people. It's also about evaluating one's own reasoning by putting it in the clearest and most explicit form possible. The image of the argument tucked in your pocket is doubly appropriate, given how hard it can be to precisely identify the jumbled contents of your own pockets: the best way of deciding what's really in there is to pull it all out into the open and get a clear look at it.

WHAT MAKES AN ARGUMENT (GOOD)?

There are at least two ways of approaching the study of arguments and arguing. Each approach carries with it a distinct conception of how arguments succeed and fail. We will make use of both notions of argument, which are closely related and normally travel together.

Definition 1: Argumentation is a rational practice.

This approach to arguments stresses the fact that arguing is a process, one that occurs in a communicative context. Argumentation is a practice by which we aim to show the reasonableness of an assertion, up to whatever standard of reasonableness is called for in that context. (Or, in failing to do so, to learn that the assertion was not as well founded as it had first appeared.) Presenting an argument is a way of making good on the obligation to support an assertion, and there are as many different practical aims of argument, in this sense, as there can be reasons for making an assertion. An argument can be a means of education or of explanation.

It can be a way of rationalizing or ordering reasons, and it can be a way of explicitly weighing competing reasons against each other. In each of these roles, though, a good argument is the presentation of a collection of premises that jointly are rationally persuasive of a conclusion. Taken together, the premises make it reasonable to believe the conclusion.

Definition 2: Arguments are linguistic or logical objects.

It can also be very illuminating to think of an argument just as a set of sentences or propositions, factoring out considerations about the speaker, the audience, the context, and the broader goal of giving the argument. Are these sentences true, and do they jointly imply this conclusion? This way of asking whether the argument is a good one just sets aside worries like whether an audience will find the premises *reasonable*, worries that can be viewed as merely practical complications that arise when speakers use the argument. In other words, a good argument in this sense can be defined in terms of the truth of the premises and their logical relation to the conclusion, without any allusion to speakers or hearers or context. It is in this sense that a good argument is defined as sound. This means, first, that it is valid: its premises are relevant to its conclusion in such a fashion that if the premises are true, then the conclusion must be true. And, second, a sound argument has all true premises. More precisely (since an argument might contain some irrelevant premises) we should say that all its **essential premises**—the ones that would remain if all irrelevancies were removed—must be true.

Obviously the first definition is the less precise and less technical of the two. Indeed, the second definition is really a refinement of the first. The idea behind using this more technical definition is that the somewhat vague notion of reasonableness can be clarified in terms of the more exact and idealized notion of soundness. In particular, when some argument fails to show its conclusion to be reasonable, this can be diagnosed in terms of its unsoundness: either it has an invalid structure, or it has one or more premises that are not true. Still, it is important to understand how these two conceptions of argument can come apart.

For example, Definition 1 reflects the idea that a good argument is supposed to be an *effective* one. The reasonableness of an argument has to be measured, at least in part, by the effects the argument has, or would have, on a basically reasonable person. We need the qualification about basically reasonable people, because we don't want to end up saying that an audience can make a speaker's argument a bad one just by being sufficiently closed-minded or, for that matter, that an audience can make a lousy argument good just by being sufficiently gullible. As to how we precisely define a reasonable person, it's a hard question. The short and imperfect answer is, by extrapolating from how people reason in uncontroversial cases.

FURTHER REFLECTION

"I am bound to furnish my antagonists with arguments, but not with comprehension."
(Attributed to Benjamin Disraeli)

NEL

Definition 2, on the other hand, makes no mention of effectiveness. An argument's being sound does not mean that anyone believes, or even ought to believe, that it is sound. We all know that, in some contexts, something can be true even though the currently available evidence suggests otherwise. So suppose we had a sound argument containing such premises. Being sound, it would be a good argument in the sense of Definition 2, but it would not give anyone (in that context) a reason to believe its conclusion. It would fail by the standards of Definition 1, because, while all its premises were true, they would not themselves be reasonably acceptable as true.

Suppose, for example, that there were a species of lizard not yet discovered by anyone, hiding in the badlands of south-central Alberta. Suppose also that I was putting together a critical thinking lecture and made up a sample valid argument using the (entirely fictional, so far as I knew) premise that there is a new species of lizard hiding in the badlands of south-central Alberta. In my lecture I say, "Either there is a new species of lizard hiding in the badlands of south-central Alberta or five plus five equals eleven. But it's false that five plus five equals eleven. So there is a new species of lizard hiding in the badlands of south-central Alberta."

Now, that argument is valid (it's an instance of *Disjunctive Syllogism*, discussed below), and in the situation described, it would also be sound because both its premises are true. But in giving this sound argument by accident, without actually having reason to believe the first premise (in particular because I have no reason to believe the first half of the first premise), I don't *do* the kind of thing that arguments are useful for doing. That is, I don't give anyone—not even myself!—rational grounds to believe the conclusion. When engaged in the practical business of giving an argument, in other words, it is not enough to produce premises that happen to be true, independent of whether you possess or have presented reasons to think that they are true. It is important that your premises are mutually accepted between you and your audience if you aim not only to give a sound argument, but also to give one that will convince a reasonable person. These considerations frequently drop away when we are analyzing the argument merely as a set of statements in accordance with Definition 2. And they drop away entirely when we do formal logic, which is the fundamentally algebraic study of valid argument-forms.

In the context of communicating an argument, the issue of acceptable premises frequently leads to at least some degree of regress of justification: the need to give some new argument to support one of your original premises, when it turns out not to be acceptable to your audience. As arguments get spelled out in actual discussions, then, they tend to acquire complex structures (although sometimes they have complex structures as initially presented, too). An argument can have not just an overall conclusion, but various sub-conclusions as well—propositions that are supported by a sub-argument and that function as premises in the larger argument. In principle, this process of embedding one argument within another can go on indefinitely, since a sub-argument too might contain a premise that has to be defended. In practice, it can go on until everyone is bored *and* hungry. Certainly it isn't rare for the discussion of some initially tidy argument to expand in surprising and mentally taxing ways, as the premises of the arguer's supporting arguments at each level of regress are called into question.

Finally, the idealization represented in Definition 2 only captures the virtues of deductive arguments, a special class of arguments whose properties will be discussed below. Many other arguments fail to meet these standards, while still counting as reasonable by the lights of Definition 1. To evaluate these arguments we eventually need to refine further concepts besides validity and soundness, especially for arguments falling under the broad heading of **evidential reasoning**. To begin, however, we will think of arguments as aiming at soundness, while bearing in mind the more general idea of argumentation as a rational practice.

SOME BASIC VOCABULARY OF COMMUNICATION AND ARGUMENTATION

Assertion: The act of stating something as if it were true.

Proposition, statement, sentence, claim: What you say in order to make an assertion. I will use these terms interchangeably for the most part, though for some purposes there are important technical differences among them.

Premise: A statement intended to provide rational support for some other statement (a *conclusion*), often in conjunction with other premises.

Conclusion: A statement intended to be rationally supported by a set of premises.

Argument: A collection of premises that justify, or are supposed to justify, a conclusion.

Validity: A structural property of arguments. An argument is valid just in case there is no way for the conclusion to be false if all the premises are true.

Soundness: A two-fold property of arguments; an argument is sound if it (i) is valid and (ii) has all true premises. By definition, a sound argument proves its conclusion.

Inference: The act of reaching a conclusion on the basis of some premises.

Premises	{ **entail** / **logically imply** / **establish** }	conclusions
Conclusions	{ **are a logical consequence of** / **follow from** }	premises
People	{ **draw** / **infer** }	conclusions

IS GOOD ARGUMENTATION A MATTER OF BEING LOGICAL?

The term "logic" and its variant forms are often used in non-technical ways, and it may seem natural to think of the goal of critical thinking as *being logical*. On both the technical and non-technical understandings, this is probably a mistake.

First, the technical side, where "logic" has at least two distinct meanings. There is a broad sense in which it means something like the *study* of rules of correct inference, and there is a singular sense in which it means a particular system of

DETAILS MATTER

A common error is to confuse *infer* with *imply*. Implication is the abstract relation between the premises and conclusion of a valid argument. People infer conclusions; premises imply conclusions.

Sometimes *imply* is used to mean something like indirect or sly suggestion. Both senses of implication are distinct from *infer*. To keep the logical entailment sense of implication clearly separate from the suggestion or insinuation sense, it's useful to call the former *logical implication*.

Wrong: This evidence infers that we should support public libraries.

Right: I inferred from this evidence that we should support public libraries.

Right: This evidence implies that we should support public libraries.

inferential rules. In the latter sense we can speak of *a* logic or of *different* ones: logic is not a single universally received system of rules. In fact, there is considerable debate over what ought to count as logical laws, with each different set of proposed logical laws giving rise to a different system of logic. Inferences that are valid in one system may not be valid in another.

Historically, going back at least to Aristotle, it was thought that logic was **univocal**, meaning that the term had only a single meaning or interpretation. It was believed that true logic flowed from three basic axioms. Later, some philosophers even described these axioms as the "**Laws of Thought**":

The **Law of Identity** (For any proposition P: P if and only if P)
The **Law of Non-Contradiction** (Not both P and not-P)
The **Law of Excluded Middle** (P or not-P)

This collection of axioms gives us what's called **classical logic**. Today some logicians continue to see these axioms as definitive of genuine logic, but all logicians are aware that there are important challenges to the universal application of classical logic. I will explain a bit about these challenges, just as a means of acknowledging them and to provide the keywords for further investigation, should you be interested.

Intuitionistic logic, for example, is a well-developed formal system that does not include the Law of Excluded Middle as an axiom. This approach may tolerate vagueness and fuzzy boundaries better than classical logic, which tends to idealize and sharpen distinctions conveniently but (potentially) inaccurately. Intuitionism also sets the bar higher for certain kinds of proof, since, without Excluded Middle, you can't just take a disproof of not-P as a proof of P. You need a *direct* proof in intuitionistic logic.

Dialetheic logic, by contrast, keeps the Law of Excluded Middle, but gives up (or restricts) the Law of Non-Contradiction. The latter tells you that if a collection of propositions contains a contradiction, the collection is incoherent. But in some situations, dialetheists argue, this produces such unwelcome consequences that we should limit the application of Non-Contradiction. For instance, in a system of legal judgments, different precedents may deliver contradictory results for a single case without it being clear that one precedent has the priority. In that sort of case we can either look around for arbitrary grounds on which to claim priority for one precedent, or we can discard that subset of the legal system as incoherent, or we can accept that there is at least a small class of acceptable contradictions. Dialetheists take the third option to be the most reasonable, all things considered.

Other systems of logic are distinguished by their giving up other theorems besides the so-called Laws of Thought; examples are *Relevance logic* and *Quantum logic*, if you feel inclined to investigate this more yourself. The point is not really to explain any of these alternatives to classical logic, but just to flag the fact that what counts as logic is not a settled question. Still other logical systems are extensions of classical logic, with certain complications factored in. Such complications can include the formalization of more notions like *belief*, *knowledge*, *obligation*, *possibility*, and *temporality*, among others. These systems are called **modal logics**. Again, the overall lesson is just that logic is not monolithic. There is a great deal of variety and many live philosophical issues involving logic and logics. While you might still read or hear about the "Laws of Thought," it is probably safer to think of them as weaker and less absolute than that name suggests. Perhaps the "Municipal Bylaws of Thought" would be a better label, indicating that the sharp and convenient categories of classical reasoning apply within some city limits, but are not necessarily suited to the wider and wilder regions beyond. While neither the details of logical formalism nor the higher-order philosophical debates about logic are our focus here, they are important and conceptually rich topics of study.

On the non-technical side, there are uses of the word "logic" that are likely to cause or embody confusions if we run them together with the technical senses. In many cases the term is used casually as a loose synonym for "sensible" or "plausible" and is applied not just to inferences and statements but to people and their choices.

> "You'd rather have cold hands than wear gloves? That's just not logical."
> "If he's athletic, then it's logical to think that he's tried skiing at some point."
> "She's very logical, so she probably doesn't like art, poetry, and music."

Each of these looser uses of the word can be fine on its own terms, but none of them is what the study of logic is about. Science Officer Spock of *Star Trek*, for example, while forever expounding on what is logical and illogical, almost never uses the terms in any precise sense. In the strict sense, to call an inference logical is simply to say that it is valid, while a "logical truth" is not just (perhaps not even) a plausible statement, but one that must be true given the logical laws one is working with. Many of the non-technical claims that something or someone is logical or illogical are best understood merely as statements of agreement or disagreement, which invoke the term "logic" in an attempt to gain an aura of authority. Rather than using sloppy terms as code words for agreement and disagreement, our aim is to learn the most precise terminology for categorizing and diagnosing arguments. So we won't just call arguments logical or illogical, since that is a sloppy and unrevealing way of using the terms.

WHAT ISN'T AN ARGUMENT?

I once owned a wheeled office chair that had the unnerving habit of throwing me across the room if I leaned back in it. It had only three wheels, so it balanced poorly. Yet while it was not a particularly good chair, it still counted as a chair. In fact we might reason this way: *it was a lousy chair; hence it was a chair.* With one armrest removed, it would still have been a chair. But with both armrests, two wheels, the backrest, and part of the seat removed (and destroyed, if that affects your intuitions), the remaining bit might have been at most a *broken chair*, or maybe we would then have said that it was no longer a chair at all. So it is with most kinds of thing: the lines are blurry between good instances, bad instances, and non-instances.

We have made a start on explaining conceptions of a *good argument*. But as we will see below, some of the ways arguments can fail to be sound—by failing to be valid, for instance—are, when taken to extremes, ways that a set of statements can intuitively fail to be an argument altogether. It would be silly and awkward to say that only successful arguments count as arguments. On the other hand, it seems strange to suppose that any set of randomly chosen sentences counts as at least a bad argument, or a **fallacious argument**, as bad arguments are known. Fortunately not much hangs on this distinction in practice. For our purposes, we will treat something as an argument if the context in which it is presented indicates that it is *intended* as an argument. So whether something counts as a random outburst of unrelated statements or as a dreadful fallacy might depend, in context, on what we glean of the speaker's intentions.

Still, there are systematic forms of reasoning that are easily mistaken for arguments, even though they are better categorized in other ways. Insisting, repeating, declaiming, vowing, defining, and stipulating should not be confused with arguing (though of course such acts may be *part* of giving an argument). Nor should mere assertion itself. The common usage of the word "argue" might extend to saying that someone has argued when she has only given unsupported assertions. But in the sense of argument we are using—as the basic unit of rational exchange—mere assertions are not arguments.

EXPLANATION VERSUS ARGUMENT

More subtle is the distinction between arguments, in which evidence is given in the form of premises to defend a conclusion, and **explanations**, in which we appeal to some facts in order to rationalize, or make sense of, some other facts. Explanation is a form of reasoning that is broadly distinct from argument while often overlapping with it. A working definition of the difference is that arguments aim at showing some statement to be worth believing, while explanations aim to make better sense of something already believed. Take, for example, the following statement:

> "The ice is thicker at the far end of the rink because the ground slopes in that direction."

The context of utterance will determine whether a statement like this is functioning as an argument or as an explanation. If the speaker is proffering the slope of the ground as a reason to believe that the ice is thicker at the far end of the rink (especially if the latter is not yet known to be the case), then the utterance is intended as an argument. But if the audience is fully aware that the ice is thicker at the far end and simply wants to know *why* this is the case, then the utterance is better taken as an explanation. In other words, the speaker need not be trying to persuade us that the ice is thicker at the far end. That may well be known already, so the explanation is not, or need not be, an argument itself. Instead the aim may just be to make sense of some known fact by fitting it into a pattern of understanding: laws, regularities, and familiar experiences.

As we will see in the case of arguments, explanations can have implicit elements—parts that are supposed to be understood from the context. The explanation above is just such an example. The speaker presumably expects us to understand that ice can be made by flooding an area with large amounts of water, that water runs downhill, and that water doesn't freeze instantly, so that the water will have time to run downhill before freezing. There are different kinds of explanation, moreover. Sometimes an explanation is **causal**, describing (in part) the prior conditions that caused some event. The explanation for a bicycle chain's getting rusty might be that it got wet, for example. Other times an explanation aims to rationalize, to order reasons and definitions, or to sort priorities according to principles of reasoning. An explanation of why there are as many even numbers as there are even and odd numbers together makes no appeal to causes, for instance, but is no less an explanation.

The argument-explanation distinction cuts across other distinctions of reasoning. There can be causal arguments and causal explanations, moral arguments and moral explanations, mathematical arguments and mathematical explanations. But aside from general virtues like clarity, a specific key virtue of good arguments is shared with explanations: roughly speaking, both arguments and explanations are supposed to teach us something. Circular or question-begging arguments, as discussed in Chapter 4, fail this condition inherently; these are fallacious arguments of the form "P, therefore P." Explanations too are open to an analogous problem of **pseudo-explanation**, primarily a matter of providing a triviality or a mere label when an explanation is called for.

The French author and playwright Molière famously parodied the pseudo-explanations of the medical doctors of his era by having one of his characters explain the fact that opium causes drowsiness in terms of a "*virtus dormativa*" of the opium. That is, opiates have the power to cause sleep. To say that opiates cause

LINK

The interrelations of argument and explanation are a complex topic, the surface of which we will barely scratch. Chapter 2 discusses a form of reasoning called **abduction,** in which the fact that a hypothesis would *explain* known facts is cited as *justification* for believing that hypothesis.

sleep because they have a *virtus dormativa,* however, is merely to label the thing requiring explanation, rather than explaining it. Explanations are not necessarily arguments, but explaining in a circle is as unhelpful as arguing in a circle.

UNDERSTANDING VALID ARGUMENT FORMS

So far we have a technical definition of a deductive argument, telling us that arguments aim to be sound, and that a sound argument is valid with all true premises. Why should this two-part definition of soundness be so heavily emphasized? There are plenty of good reasons that could be spelled out in terms of the noble goals of conceptual inquiry, but for now my answer is just this: because laziness is wonderful when you can get away with it.

The fact that there are two conditions that a sound argument must meet—being valid and having all its essential premises be true—means that there are broadly two ways an argument can fail to be sound: by having one or more false essential premises, or by having an invalid structure. In general it's harder to discover whether premises are true than to determine whether the argument is valid. To learn whether some alleged statement of fact is true, you might have to go out and actually *look at the world.* It can be a messy and time-consuming business. So it makes sense to evaluate an argument's structure before focusing on the truth of its premises. After all, an argument that would not support its conclusion even if its premises were true can be rejected without checking out the truth of its premises. If it's invalid, then it's unsound irrespective of whether the premises are true. This can be a very efficient way of evaluating an argument.

Let's clarify the ideas of validity and soundness by looking at a few examples. Some of these examples are deliberately strange, in order to make the truth or falsity of the premises and their relevance to the conclusion perfectly obvious.

First, here is an argument that is valid but unsound:

1. If Pierre Trudeau was the prime minister, then he was a professional dance instructor.
2. Pierre Trudeau was the prime minister.

Therefore,

3. Pierre Trudeau was a professional dance instructor.

Since Premise 1 is false, the argument is not just goofy but unsound. A sound argument has all true premises. Still, for all its silliness this argument has an important virtue: it is valid. Its premises relate to its conclusion in such a way that, *if* the premises were both true, then the conclusion would be true as well. The same can be said of the following argument, even though *both* its premises are false.

1. If John A. Macdonald, first prime minister of Canada, was a dentist, then he was a professional dance instructor.
2. John A. Macdonald, first prime minister of Canada, was a dentist.

Therefore,

3. John A. Macdonald, first prime minister of Canada, was a professional dance instructor.

The premises are not true, but if they *were* true, then the conclusion would be as well. That is, if you think of a far-fetched world in which these premises are true—one in which dentists all teach dance, and in which John A. Macdonald was a dentist—then you're thinking of a world in which John A. Macdonald taught dance. So this argument, however silly, satisfies one condition for soundness (it's valid) but fails the other condition (it has false premises). The argument is valid but unsound.

It should be clearer now what it means to say that validity is a *structural* property of arguments. We are not concerned with whether the particular premises are actually true when we evaluate the validity of some argument, but only with whether the conclusion would be true if the premises were true. This enables us to factor out the specific content of the premises. For instance, both of the arguments we just examined have the same structure:

1. If P then Q
2. P

Therefore,

3. Q

The arguments are valid because this is a valid *form* of argument; its technical name is **Modus Ponens.**

The letters P and Q in this argument form are just placeholders; they each stand for any declarative sentence that we might plug into that structure. So we can substitute any two (declarative) sentences for P and Q in this argument framework and produce actual arguments. To say that the structure is valid is to say that any choice of sentences for P and Q that makes both premises true also makes the conclusion true. Notice that having true *premises* does not necessarily mean having true Ps and Qs! If one premise in an argument is not-P (i.e., "It is not the case that P"), then only a *false* sentence P will make the premise not-P true. In the modus ponens case, as long as we pick P and Q in such a way that it's true both that P, and if P then Q, then Q will also be true. That's a structural property of the argument—a property that remains even if P and Q don't happen to make the premises true. Looking again at the first two sample arguments, then, we can see that even though they have false premises, they still possess at least the virtue of validity. If their premises were true, then their conclusions would be true as well.

Another basic and important valid argument form is called **Modus Tollens:**

1. If P then Q
2. Not Q

Therefore,

3. Not P

FURTHER REFLECTION

There are a few Latin terms that are almost unavoidable in critical thinking, since the study of reasoning goes back to antiquity and inherits many terms from that history. Where possible I will use English terms. In addition, for purposes of clarity I will capitalize the names of fallacies. In general, though, this is unnecessary.

Again, to say that this is a valid structure is just to say that any choices of P and Q that make both (1) and (2) come out true are guaranteed to make (3) come out true as well; there is no way for the conclusion to be false if the premises are true. For example, if P = "Canada is a republic" and Q = "Canada has no monarch," then both

1. If P then Q

and

2. Not Q

are true. (If Canada is a republic, then Canada has no monarch; and it is not the case that Canada has no monarch.) Since Modus Tollens is a valid argument form,

3. Not P

should be true as well . . . and so it is. Canada is not a republic.

So far we have seen two examples of arguments that are valid but unsound. Next we will see an argument having the opposite problem. This argument has all true premises—indeed, it has a true conclusion—but it is invalid, and hence unsound.

1. 42 is a number.
2. Continents drift.
3. Pigs grunt.

Therefore,

4. Socrates is Greek.

Here each of the premises is true, but. . . yuck. It's only because the parts of this argument are numbered, and connected with "therefore," that we'd even imagine it to be intended as an argument. The premises do not relate to the conclusion in any way; in particular, the premises could all be true even were Socrates from Lithuania. So this argument is invalid and hence is unsound. This is the case even though the argument has not just all true premises, but a true conclusion as well. It's unsound anyhow. A lousy argument with a true conclusion is a lousy argument: after all, the idea is not simply to have the alleged conclusion happen to turn out true. The idea is for the conclusion to be *supported* by the premises. (By the same token, a lousy argument for a conclusion does not mean that the conclusion is *false*, either. It just means that the argument gives us no reason to believe it.)

This argument, however, is sound:

1. Vixens are foxes.
2. All foxes are mammals.

Therefore,

3. Vixens are mammals.

As an aside, notice that for the sake of convenience the above argument is framed in what are called **categorical** terms: it's a matter of how the categories *Vixen*, *Fox*, and *Mammal* are related. This way of framing logical inferences was largely the extent of logic for many centuries and is still frequently taught as informal logic owing to its ease of use in natural language. Since the rise of modern logic in the late nineteenth and early twentieth centuries, a more rigorous and general treatment is given to inferences of this sort, fitting them into the form of

conditional reasoning (like "If P then Q"). In the above case, the same argument could be given using a **quantifier expression** (like "some" and "every") and a **variable** like *x* or *y*—a kind of placeholder for objects that could be plugged into the statement.

1. For every object *x*, if *x* is a vixen then *x* is a fox.
2. For every object *x*, if *x* is a fox then *x* is a mammal.

Therefore,

3. For every object *x*, if *x* is a vixen then *x* is a mammal.

I will not focus on the distinction between categorical and quantificational reasoning, nor on the axioms or deduction rules that characterize either approach to formal logic, beyond the cases useful for explaining validity and soundness. For example, we now have an argument with the valid form of **Hypothetical Syllogism**. The premises seem quite clearly true, given the definitions of "vixen," "fox," and "mammal." So the argument is sound; it proves its conclusion true.

Now let's consider an example of a sound argument that falls outside the conditional mould:

1. Either foxes are mammals or rabbits are birds.
2. Rabbits are not birds.

Therefore

3. Foxes are mammals.

This argument also has a valid form for which there is a name: **Disjunctive Syllogism.**

1. P or Q
2. Not Q

Therefore,

3. P

Again, the premises of the sample valid argument are clearly true, so the argument is sound. Finally, we can see how more complex arguments can chain together such basic argument forms, proving intermediate conclusions along the way. This argument makes use of both Disjunctive Syllogism and Modus Ponens.

1. Either foxes are mammals or rabbits are birds.
2. If foxes are mammals, then vixens lactate.
3. Rabbits are not birds.

Therefore

4. Foxes are mammals.

Therefore,

5. Vixens lactate.

Lines 1 and 3 imply line 4 by Disjunctive Syllogism, then 2 and 4 imply 5 by Modus Ponens.

While people rarely present their reasoning in such a regimented format, this is the refined conception of what arguments are. Even when they don't come packaged in this way, arguments *can* be packaged in this way for the purposes of clarifying and assessing them. When we assess them—when we are interested in their soundness, or at least in their approximation to soundness—it is useful to put them in the format of numbered premises and conclusions. Throughout the book we will see examples of how to do this.

CHECKING AN ARGUMENT'S VALIDITY: THE METHOD OF COUNTER-EXAMPLE

Now we are in a better position to appreciate the utility and convenience of testing an argument's soundness by first testing its validity. What is needed, then, is a method for discovering whether an argument is invalid. There are several such methods, but a particularly quick and useful way of testing for invalidity is called the **Method of Counter-example**.

Our definition of validity tells us that there is no way for the conclusion of a valid argument to be false if all the premises are true. This means we can tell that an argument is invalid if we can think of ways for the premises all to be true while the conclusion is false. That's the Method of Counter-example in a nutshell: think of a situation in which the premises would be true but the conclusion would be false. If you can think of such a situation (or "construct a model," as some say), then the argument is invalid. The truth of its premises doesn't guarantee the truth of its conclusion, because you've thought of at least one situation in which the premises would be true but the conclusion false.

To take a simple example: every natural number has an immediate successor (i.e., the number you get by adding 1 to it); therefore, every natural number is the successor of some other natural number. A quick and easy way to see that this is an invalid argument is to look for a counter-example: some number that is consistent with the premise but inconsistent with the conclusion. Since 0 is clearly such a number, it shows the invalidity of the argument. Consider too this slightly more complex example:

1. The club president appoints the treasurer.
2. The chair of the club's board of governors appoints the vice-president.

Therefore,

3. The treasurer and the vice-president are appointed by different people.

Here again we can test for invalidity by thinking of scenarios that are consistent with (1) and (2), but inconsistent with (3). What about the scenario in which a single person serves both as president and as chair of the board of governors? Nothing in the premises rules this out, yet the conclusion would be false in that situation. So the argument is invalid; its premises would be true and its conclusion false if the president was also the chair.

Again, if all we're concerned with is whether a deductive argument provides grounds to accept its conclusion, it makes a great deal of sense to first consider its validity. Assessing the truth of premises can be hard work; we circumvent that problem altogether if we can see that *even were the premises true* they would not support the conclusion.

Here is an argument that a reader could reject without even knowing what the key terms *mean*, still less whether the premises are true:

1. If something is a dioxin then it's a polycyclic compound.
2. Some polycyclic compounds are carcinogens.

Therefore,
3. If something is a dioxin, then it's a carcinogen.

Because validity concerns an argument's form alone, a valid argument should remain valid (though not necessarily sound!) if we uniformly substitute one predicate or name for another throughout the argument. We can test the above argument for validity by substituting some more familiar terms for the potentially unfamiliar ones and seeing whether we can thereby make the premises true and the conclusion false.

1. If something is a fox, then it's a mammal.
2. Some mammals are herbivores.

Therefore,
3. If something is a fox, then it's a herbivore.

Clearly this argument is invalid—both premises are true and yet its conclusion is false. But structurally it's the same as the argument above it. Using the same argument structure (i.e., just uniformly substituting "fox" for "dioxin," "mammal" for "polycyclic compound," and "herbivore" for "carcinogen" everywhere those terms occur) we've found an interpretation that makes the premises true and the conclusion false. Hence the original argument—*whatever* its key terms mean—is invalid. That's not to say its conclusion is false; it's just to say that that argument doesn't give us reason to *believe* the conclusion. So the argument fails.

Now you might think, "But what if I'd chosen 'carnivore' instead of 'herbivore'? Then the premises and conclusion would all have been true." That's correct, but it's irrelevant. An invalid argument can happen to have true premises and a true conclusion; even in that case it still fails the condition that the premises are what *establish* the truth of the conclusion. The existence of even one counter-example—either a scenario in which the actual premises are true and the conclusion false, or a scenario on which a structurally identical argument has true premises and a false conclusion—shows that an argument is invalid. So when we test for invalidity, we have to actively search for counter-examples.

Overall, then, testing for validity is a very useful first tool in the evaluation of arguments. It can cut short a potentially long and difficult investigation into the argument's soundness. Of course, one might nevertheless want to investigate the truth of the argument's premises—just in the interests of learning and thoroughness. A complete evaluation of an argument will still consider both validity and the truth of the premises.

VALID ARGUMENT FORMS

Here is a selection of further argument forms recognized as (classically) valid.

Simplification

P and Q
Therefore,
P

"Eric and Ellen are both doctors. Therefore, Ellen is a doctor."

It might seem that such a simple and obvious inference rule could never need to be invoked in everyday communication. But this sort of inference can be usefully invoked in setting aside irrelevancies, as in "If she's a convicted criminal but truthful, then she's truthful, and her word can be trusted." It may also apply to complex claims of predication, sometimes flagged with the Latin term **a fortiori,** meaning "with stronger reason."

A: "The rules say that prisoners aren't allowed to vote."
B: "Well, that probably doesn't apply to *repentant* prisoners."
A: "A fortiori, a repentant prisoner is a prisoner. So the rule applies to them too."

A's reply here implicitly invokes a simplification rule, moving from "repentant and prisoner" to "prisoner." (Not all predications will fit this logical mould, however. For example, a *young head of state* might simply be young *for* a head of state, rather than both young—period—and a head of state.) We've already seen another example of a fortiori reasoning, in the case of my office chair: it was a lousy chair; a fortiori, it was a chair.

Conjunction

1. P
2. Q

Therefore,

3. P and Q

"Eric is a doctor. Ellen is a doctor. Therefore, Eric and Ellen are both doctors."

Notice that putting the conclusion in this way, rather than just writing "and" between the two premises, requires that the term "doctor" is being used univocally—that is, with one meaning, non-metaphorically in both cases.

Addition

1. P

Therefore,

2. P or Q

"Foxes are mammals. Therefore, either foxes are mammals or cows are mammals."
"Foxes are mammals. Therefore, either foxes are mammals or lizards are mammals."

Here it does not matter whether the statement added using "or" is true or false. For reasons discussed below, the resulting statement is true as long as *at least* one of the sub-statements is true. If we already know that P is true, this guarantees that "P or Q" is true—no matter what Q is.

This shows how truth conditions can come apart from the conditions for a cooperative or useful assertion. We might be guaranteed to say something *true* with "Foxes are mammals or lizards are mammals," if we know that foxes are mammals. But normally it would be uncooperative or misleading to assert such a truth if we know that lizards are not mammals, or even if we have no idea either way.

Hypothetical Syllogism

1. If P then Q
2. If Q then R

Therefore,

3. If P then R

> "If the dollar is devalued, exports will rise. If exports rise, then unemployment will fall. Therefore, if the dollar is devalued, unemployment will fall."

Constructive Dilemma

1. P or Q
2. If P then R
3. If Q then S

Therefore,

4. R or S

> "Either it will snow tomorrow or there will be a quiz in class. If it snows tomorrow, classes will be cancelled. If there's a quiz in class tomorrow, I'll fail it. So either classes will be cancelled tomorrow or I'll fail a quiz tomorrow."

Destructive Dilemma

1. If P then R
2. If Q then S
3. Not R or not S

Therefore,

4. Not P or not Q

> "If Zainab called her mother, the answering machine took a message. And if her brother called her mother, the line was busy. But the machine didn't take a message, or the line wasn't busy. So Zainab didn't call her mother, or her brother didn't."

You should be able to see how some of these valid forms are constituted out of simpler forms. Of course, one can have fabulously long and complicated argument forms; here I list just some simpler ones to aid in the recognition of valid arguments when you encounter them in natural language.

OTHER STRUCTURAL PROPERTIES OF ARGUMENTS

Besides being valid or invalid, an argument can be organized in a range of ways important to assessing the support it gives to its conclusion. Most of the arguments we have examined so far can be classified as **linked arguments**: their premises essentially tie together to support a single overall conclusion. Modus Ponens is a clear and simple example of this: by themselves, neither *P* nor *If P then Q* are grounds to conclude *Q*. But together they link up to imply the conclusion. Linked arguments are really the most fundamental sort of argument. But arguments can be structured in more complex ways that build on linked arguments.

One example is a **convergent argument**, in which a range of independent grounds for a conclusion are assembled together as premises. No premise in a convergent argument requires the other premises in order to support the conclusion; rather, each premise directly supports the conclusion. Suppose, for example, that I said, "Cars are bad for the environment. They burn petroleum, which pollutes the atmosphere, and the methods of their construction involve mining and industrial practices that pollute the soil and water as well." Here a collection of reasons are given for thinking that cars are bad for the environment, none of which rests on any of the others. In the case of deductively valid arguments, this kind of argument can amount to overkill, since just one sound argument for a conclusion is sufficient to establish its truth. But it's one thing to present an argument as sound, and another thing for one's audience to accept it as sound. Often it makes sense to provide convergent reasoning in the hopes that an audience who denies the truth of one premise or the validity of one of the premise-conclusion inferences will still find another of the converging lines of argument persuasive. That is, maybe someone is unconvinced that automobile exhaust is particularly bad for the environment, but accepts that mining and industrial practices are. Even if both lines of argument are sound, in this case casting the net more widely ensures that the argument is effective.

In the case of evidential reasoning, discussed in Chapter 2, convergent arguments are especially important. A collection of convergent premises, each independently supporting the conclusion *to a degree*, can add up to a stronger argument than any of the premises on its own. We can think of convergent arguments as collections of distinct simpler arguments. This is important for the purposes of analyzing arguments, because finding a false premise in a convergent set of reasons doesn't entail that the overall argument isn't sound. The other lines of reasoning in the convergent argument might nevertheless be sufficient to entail the conclusion.

Another way of putting together basic arguments into complex forms is through **sequential arguments**. These are cases in which premises establish intermediate conclusions, which then serve as premises for some further conclusion.

1. Literacy is an important life skill.
2. Anyone who gets accepted to a major university is literate.
3. Jill got accepted to a major university.

Therefore,

4. Jill is literate.

Therefore,
5. At least one person is literate.
Therefore,
6. At least one person has an important life skill.

In this argument, lines (4) and (5) are both intermediate conclusions that follow from earlier premises and premises from which subsequent conclusions follow. Many arguments are sequential in this fashion, proving intermediate conclusions for the purpose of making the move to the main conclusion especially clear and self-evident.

Finally, arguments can incorporate elements of all three structures, with complex hierarchical orderings. Someone (in particular my mother-in-law) might say:

> "Oat bran is healthy, and it's tasty. Besides, it's affordable and can be added to many different dishes. This makes it good value for the money. So it's a good ingredient to have in the pantry."

This reasoning really consists of four convergent lines of argument for an intermediate conclusion, with a main conclusion that follows sequentially. But the convergent arguments aren't entirely independent of one another; plausibly, it's the fact that oat bran is healthy that makes its (alleged) tastiness and affordability relevant. After all, automotive antifreeze might be tasty and affordable, but it's poisonous—so these properties on their own wouldn't really count in its favour as a food choice. The argument is partly convergent, then, but with some implicit linkage or interdependence between some of the convergent lines of reasoning.

Constructing an effective argument is partly a matter of understanding these sorts of dependence relations among the premises and being able to make them clear, just as analyzing an argument effectively requires perceiving these relations. In the case just given, for example, it would be a serious mistake to think that the other convergent lines of argument remained in force if it turned out that oat bran was in fact unhealthy.

TRUTH CONDITIONS

The first element of soundness is validity; now let's consider the second element. What does it mean to say that a premise is true? When we evaluate the truth of a premise, we are in effect checking to see whether there is an appropriate fit between the state of the world and the how the claim represents the world as being. We use the term **truth conditions** to mean how things would have to be in order for the statement to be true.

In many cases it may seem that explaining the truth conditions of a sentence is as simple as removing the quotation marks from it. The sentence "Felix the cat is on the mat" has truth conditions that are simply explained as a matter of Felix the cat's being on the mat. But this is not much of an explanation in a wide class of cases. Even this apparently simple case is complicated by there being more than one cat named Felix in the world, so particular *uses* of the sentence will have different truth conditions. My utterance of that sentence in a room with a salient cat named "Felix" that is on a mat will normally be judged true, even if there are thousands of cats named Felix around the world that are not on mats at that moment. But even in that sort of immediate context, it still might be clear to everyone involved that I am actually making a claim about some other cat named Felix, perhaps far away (in an attempt to display my psychic abilities, say). Under *these* conditions, the presence of the local Felix on the local mat would not suffice to make my statement true. When we are considering what makes some statement true, a range of such interpretive issues can arise. Some of these are discussed below; others are examined in detail in Chapter 3.

TRUTH AND REASONABLENESS

It is arguments, and not single statements, that are valid or invalid, sound or unsound. On the other hand, it is statements and not arguments that are true or false. The term "reasonable" is broad enough to apply at various levels, but we will use it, like truth, to apply to individual statements rather than arguments. A reasonable statement on our definition is one with sufficient evidence, all things considered, to render it acceptable in a given state of information.

We often call statements true even though we are rarely, if ever, in a position to know that they are perfectly accurate, immune to falsification by any subsequent discovery of evidence. This is a significant difference between truth and reasonableness, though. If our assertion later turns out to have been mistaken, we say that it was false, but this is consistent with continuing to hold the assertion was reasonable at the time it was made. While we *aim* at the truth in making and defending our assertions, the means by which we aim, and typically the only virtue we can be certain our assertions have, is their reasonableness according to our information at any given time.

Truth is often formally understood to be a discrete concept: either a statement is true or it is false, with no intermediate cases. Another way of putting this is to say that statements have one of only two **truth values**. Strictly speaking, statements are not "sort of true" or "pretty much true." On this view, the word that means "approximately true" is "false."

An important philosophical question is how to characterize truth in general. For example, we may wish to admit the possibility of statements that are neither true nor false. But admitting this possibility is still consistent with the idea that statements are either true or not true, with no middling cases; it's just that there could be more ways of being not-true than of simply being false. However, even if we accept truth as **bivalent** (having only two possible truth values), this does not mean that we talk nonsense if we say that some statement is only approximately true. We might just interpret such talk as not quite literal. Such claims could be understood as meaning that there is some *other* statement, quite similar to the one in question, that really is true. When we say, "It's approximately true that Ted lied," we could be understood to mean "It's true that Ted did something approximately like lying." In other words, we need not suppose that an approximately true statement has some property of *semi-truth;* certainly it's simpler to suppose that such a statement is strictly false, but has a similar meaning to a (genuinely, fully) true statement.

There are philosophically deep waters here. Whatever we end up saying about truth, though, *reasonableness* clearly comes in degrees. The evidence available in support of some statement can range anywhere between overwhelming, considerable, marginal, and none. Understanding how we can reliably measure the degree of evidence for a statement is a major aspect of critical thinking. In general it is a mistake to suppose that we can call a valid argument unsound only if one or more premises are absolutely known to be false. An argument still fails by the lights of Definition 1 in the event that there's no positive reason to accept the premises; we don't actually need positive reasons to *reject* the premises in order to say that the argument fails. For in the case where there's simply no reason to accept the premise(s), there is correspondingly no reason to accept the conclusion. But that's what an argument was supposed to give us.

On the other hand, just as a matter of the meaning of the word *true*, the following seems correct: enough reason to accept some claim is enough reason to accept it as *true*—however provisionally. To the extent that we regard a statement as worth asserting, and to the extent that we are even interested in the *possible* truth of some statements, sound arguments will be crucial for revealing what entails and what is entailed by those statements. So even though the definition of soundness implicates truth rather than degrees of reasonableness, the standards of deductive reasoning remain relevant to our assessment of the consequences of sets of beliefs or claims, virtually without regard to what we say specifically about the standards for regarding claims as true—that is, without regard to how much evidence makes it reasonable to assert a claim.

NECESSARY TRUTHS AND DEFINITIONAL TRUTHS

An important class of true statements is those that *must* be true, either because they are truths of logic and mathematics, or because they are true by definition in a broader sense. Many philosophers of logic and language would say, for example, that when you realize the following sentence is true—It is not the case both that all bears hibernate and that not all bears hibernate—you are realizing something not about bears, but about logic. That sentence is an instance of a logical law, the Law of Non-Contradiction: not both P and not-P. Because it holds for all declarative sentences, irrespective of their subject matter, the truth of any instance of the law is really a matter of logic and not about the workings of the world. Even if an utterance of the sentence "Felix is on the mat" is true in some situations, this is merely a **contingent truth**; things might have turned out differently. But to the extent that Non-Contradiction is a logical law, the sentence above and any other instance of the law are **necessary truths**: they would be true no matter how things might have turned out (even if there were no bears at all).

Propositions of mathematics are standard examples of necessary truths, since they too, if true, could not have been false. Their truth follows from the definitions of the basic objects and operations of mathematics. Necessarily, then, 15 is greater than 6; this is simply a matter of the definitions of "15," "greater than," and "6." But other definitions abound in the way we communicate, including in argumentation, without their status always being so obvious. Someone learning about football might hear it said that the ground can't cause a fumble if the quarterback slides feet-first. This might sound like an assertion about the proper way to avoid having the ground cause a fumble—i.e., by sliding with the proper technique. In fact it's true by definition; someone who said this would not be making a claim about the causal powers of the ground, but rather laying down part of the definition of a fumble in football. According to the rules, if the quarterback slides feet-first and the ground knocks the ball loose, it doesn't count as a fumble.

Because necessary or definitional truths are often said to really be about logic or language, many people find it tempting to say that they are only *trivially* true, or that they convey no genuine information. There is little point to taking such a view, from our perspective. What matters is recognizing what sort of truth we are dealing with. Naturally if a premise in an argument is recognized to be true by definition, this removes any concern about confirming its truth by investigating facts or demanding further proof. But it's also important to be sensitive to at least two possibilities when dealing with such premises. First, there is the prospect that a claim equivocates between definitional and contingent status—that it gets used in both ways within the same argument. This is the problem with the No True Scot fallacy, for instance, discussed in detail in Chapter 4. And there is also the serious question about the power of definitions and who recognizes them. The fact that one side in a discussion thinks that a premise is true by definition does not settle the question; there are many, many examples of definitions that properly change over time to accommodate new discoveries or new perspectives.

When nineteenth-century scientist John Dalton re-introduced the ancient Greek notion of the *atom* to physics and chemistry, he defined atoms as indivisible (which is even what the word "atom" means, in Greek). But clearly it would not have made sense, upon the later discovery that atoms could be split, to say, "No, this discovery must be mistaken; after all, it's true by definition that atoms are indivisible!" Similarly, there was a time in most nations and cultures, including Canada, at which the legal definition of "voter" excluded women. The fact that this was true by definition, in some sense, was not a good argument against extending the right to vote to women, since the very question at issue was whether this was a *good* definition of voter. Premises can be treated as true by definition only if the acceptability of the definition itself is not contested.

Finally, questions of definition give rise to the important distinction between **necessary and sufficient conditions**. Objects or types of object are often defined in terms of a set of conditions that they satisfy. For example, the definition of a mammal typically includes such properties as being warm-blooded, having body hair, and nursing offspring by lactation. When we reason about how such properties or concepts fit together, we need to see the logical relations that such a definition imposes.

We can say that being warm-blooded is a necessary condition for being a mammal; if it's not warm-blooded, it's not a mammal. But the opposite is not true. Being a mammal is not a necessary condition for being warm-blooded. (Birds are warm-blooded too, so something can be warm-blooded without being a mammal.) Rather, being a mammal is a sufficient condition for being warm-blooded; if it's a mammal, it's guaranteed to be warm-blooded. Nor is being warm-blooded sufficient for being a mammal, since the other conditions must also be satisfied.

It is often important to note these distinctions. "Terry must be a good teacher," someone might assert. "After all, she's very enthusiastic about her subject area." Here we can reply that enthusiasm is at most a necessary condition for being a good teacher, but not a sufficient condition. Other things besides enthusiasm are needed. For the same reason, the inference from "Terry is not a good teacher" to "Terry is not enthusiastic about her subject area" also does not follow. Terry might fail to be a good teacher by failing to satisfy other necessary conditions instead. As a final example that we have already explained, notice that an argument's validity is a necessary condition for its soundness, while its soundness is a sufficient condition for its validity.

Of course some relations between conditions are neither necessary nor sufficient. Having excellent diction might be *correlated* with being a good teacher, in the sense that speaking clearly makes it *likelier* that one is a good teacher, without being necessary for it.

TRUTH CONDITIONS OF COMPOUND SENTENCES

Premises and conclusions are rarely all simple statements. Usually the claims made in an argument are complex in some respect—for instance, involving two or more sentences joined together. For this reason the assessment of an argument requires that we pay attention to the truth conditions of such complex sentences, since the way they are joined makes all the difference to what they entail and what entails them.

Simple (or atomic) statement: A sentence that does not contain another sentence as one of its parts. For example, "My dog has fleas" and "Continents drift."

Conjunctive statement, or conjunction: A compound statement containing two or more sub-statements (called **conjuncts**), usually joined with the words "and" or "but." A conjunction is true if and only if both of its conjuncts are true. To put it slightly more technically, a statement of the form "P and Q" is true just in case P is true and Q is true. Near-synonyms like "as well as" might be used instead. A term like "but" indicates a different speaker's attitude than "and" but also creates a conjunction. (This point is explored in the section on *rhetoric* in Chapter 3.)

Disjunctive statement, or disjunction: A statement of the form "P or Q" is true just in case at least one of P and Q is true.

A compound statement containing two sub-statements (called **disjuncts**), joined with the word "or" or near-equivalents like "alternatively." A disjunction is true if and only if one of its disjuncts is true.

Famously the word "or" can be understood *inclusively* or *exclusively*. To say that a disjunctive statement uses the inclusive "or" is to say that *at least* one of the listed disjuncts is true; in other words, the disjunction is also true if *both* its disjuncts are true. The exclusive "or," by contrast, applies when one and only one of the disjuncts is true. "Fish or cut bait!" is an idiom meaning "Do one thing or the other, but not both"—that is, it exhibits the exclusive "or." But for most purposes, it is best to treat "or" inclusively, as far as the meaning of the word itself, and to regard the exclusive "or" as an artifact of **implicature** in some contexts: that is, a further interpretation that goes beyond the strict and literal meaning of the words. (Implicature is discussed in more detail in Chapter 3.)

The intuitive difference between disjunctive and conjunctive statements, for the purposes of logic and evidential analysis, is that it's particularly easy for a disjunctive statement to be true. A disjunctive statement is true provided *any* of its disjuncts are true, while a conjunctive statement is true just in case *all* of its conjuncts are true.

Conditional statements: A statement of the form "If P then Q" is true unless P is true but Q is false.

Conditionals are sentences with an if-then form. In the above example, we would call P the **antecedent** (the "if" part), and Q the **consequent** (the "then" part).

The question of the truth-conditions of conditional statements is surprisingly complicated, so the definition given above should be taken as quite tentative. This is the definition of the if-then symbol employed in classical formal logic, so it has considerable currency, but it does have some unhappy consequences. For example, by this definition the sentence "If continents drift, then grass is green" is true, as is "If continents don't drift, then grass is purple." (You can check these examples against the above definition to confirm that they count as true.) Why should we adopt such a definition?

The short answer is this: because it captures many of our intuitions when understood properly. Often when we *use* a conditional statement, we intend to convey some sort of explanatory relation between the antecedent and the consequent—for instance, that P is what *made it the case* that Q. But that too can often be explained as something like an implicature. What matters for the sake of argument analysis, in any event, is whether an argument containing conditional statements *requires* them to be interpreted in the explanatory sense or not. As long as we're clear on what the argument requires, the task is just to determine whether it succeeds on that interpretation.

It can be an easier matter to identify a conditional statement in English than to recognize which *direction* the if-then relation runs, since the antecedent need not come before the consequent on various constructions. For example, each of the following makes the same conditional claim:

Abigail swims if Bill bikes.
If Bill bikes, Abigail swims.
Bill bikes only if Abigail swims.
Only if Abigail swims does Bill bike.

The latter two of these sentences may seem surprising at first. Why should two such strong-sounding claims have the same truth-conditions as the first two? The answer emerges quickly upon reflection. If it really is true that *if Bill bikes, then Abigail swims,* then there is no case in which Bill bikes but Abigail doesn't swim. But that's just to say that Bill bikes only if Abigail swims. The sentence need not be telling you anything about the explanation for Bill's biking or Abigail's swimming; it need not be proposing a causal connection between the two events. It just specifies a correlation between the truth-values of the two sub-sentences: that whenever "Bill bikes" is true, then so is "Abigail swims." In general, then, the following two formulations are logically equivalent.

If P then Q
Q only if P

On the other hand, their different word orders and different emphases can mean that the two logically equivalent formulations are suited to convey different messages in practice. This, again, we can treat as a matter of **rhetoric** and implicature, discussed in Chapter 3.

There is a genuine distinction between two kinds of conditional statement in natural language, however. They are **indicative conditionals** and **subjunctive conditionals**.

Basic indicative conditional: If P then Q.

Subjunctive conditionals: If it were to be the case that P, then it would be the case that Q. There are some important differences in how to analyze reasoning that employs these two kinds of sentence. However plausible the definition of truth conditions for conditionals given above (a refined definition that's usually called the **material conditional**), it cannot be correct for conditionals in the subjunctive mood. For example, the argument form Hypothetical Syllogism is valid for indicative conditionals, interpreted as material conditionals, but fails for subjunctive conditionals since there are counter-examples to it. The following argument looks like Hypothetical Syllogism, yet even though each premise seems plausible, and the conclusion formally appears to follow from them, the conclusion is intuitively false.

1. If baseball player Nolan Ryan had thrown a cricket ball, he would have been a cheater.
2. If baseball player Nolan Ryan had played cricket, he would have thrown a cricket ball.

Therefore,

3. If baseball player Nolan Ryan had played cricket, he would have been a cheater.

With true premises and a false conclusion in this example, Hypothetical Syllogism is clearly not a valid form for subjunctive conditionals.

Now, you might notice that in this argument I've switched the order of the premises, in comparison to the Hypothetical Syllogism argument form defined earlier. The order of premises normally should not affect the validity of an argument. Why do you think I made this switch, then?

The reason is that the order of premises does affect the interpretation here. Subjunctive conditionals are often used to express **counterfactuals**, or statements about the way things might have been but are not. This idea is sometimes expressed in terms of various **possible worlds**, a convenient shorthand for talking about how things might have been. The reason the above argument goes wrong is that both premises are counterfactual conditionals, but they are about *different* possible worlds—different ways that things might have been. The possible world in which Nolan Ryan is a baseball player but uses a cricket ball anyhow is different from the one in which he uses a cricket ball because he is a cricket player. If you look at the above argument and take its premises in reverse order, chances are that your interpretation of which possibilities are under consideration would have been initially set by your interpretation of Premise 2. That is, if Premise 1 followed Premise 2, it would ring false—reading them in that order, you would begin by thinking of the situation in which Premise 2 is true, and *in that possible world* Premise 1 is false. To make the counter-example more effective, then, I reversed the order of the premises. As written, the argument makes it easier to interpret each premise as true, relative to different possible worlds. But the rules of inference we've sketched so far don't allow for that particular subtlety.

Formally spelling out valid inferences for counterfactual conditionals requires the resources of modal logic, mentioned earlier as a family of different logics that allow for talk of possibility and necessity (among many other things). We won't delve into this topic, but informally we can draw an important lesson: whenever we are confronted with complex reasoning about possibilities, including assertions about what would have happened under different circumstances, we have to bear in mind how the different possibilities fit together. For similar reasons, the conjunction rule also fails in modal contexts: from "Possibly P" and "Possibly Q," we can't conclude "Possibly P and Q." Consider the case in which Q is *not-P*, for example. It's possible the Canucks will win the Cup in the year 2020, and it's possible the Canucks will not win the Cup in 2020, but it's not possible that the Canucks will both win and not win the Cup in 2020. The details of the possibilities in question are often essential to making reasonable arguments about would have been the case, what might occur, or what must have happened.

Negation: A statement of the form "Not-P" (that is, "It is not the case that P") is true if and only if P is false.

It is only three letters in English, but the word "not" is tricky. Like a range of related terms and prefixes like "no," "non-," "anti-," and "false," it performs the logical role of negation. In one sense, negation is simple: we use it to say what isn't the case. But there are many different ways of doing so, while the meaning of "not" itself is less precise and less univocal than it may seem. From logic to linguistics to computer science, finding a complete and correct account of negation is a much-sought goal.

If we were going to do formal logic, we would have to elaborate considerably on the *meaning* of negation, especially by taking a stand on whether not-not-P is equivalent to P. (An inference rule known as **Double-Negation Elimination**, its plausibility as an axiom essentially reduces to that of the Law of Excluded Middle.) While this enormously simplifies a logical formalism, it doesn't always square with our use of negation in natural language. For instance, when asked a question like "Is Ted pleasant?" one might answer, "Well, he's not unpleasant." For predications that can be of this neither-nor sort, Double-Negation Elimination seems to force sharp distinctions where none can be found. Bearing in mind that it can be an over-simplification, though, for the most part we can use the simplest and most obvious story: not-P is true when P is false. For our purposes, then, the more important subtleties involve the interpretation of ambiguities that arise from the negation of compound and complex statements. When we encounter an utterance like "Seabiscuit was not the fastest, prettiest, gutsiest racehorse," it's important to recognize that we don't know exactly to what the negation applies, at any level more precise than the whole predication: it could be any subset of *fastest, prettiest, gutsiest, racing,* and *horse* that is being denied.

LINK

This is a kind of **scope ambiguity,** discussed in more detail in Chapter 4.

This is an especially common ambiguity when negations are uttered as "prosentences"—words that stand for sentences the way that pronouns stand for nouns. If someone says "Continents drift" and I reply "Yes" or "True," I am using a prosentence to endorse their utterance, as a handy way of agreeing without uttering the sentence myself. Essentially I too am saying, "Continents drift." But if someone says, "Three members of Grant Devine's provincial cabinet were criminally prosecuted," and I reply, "No" or "False," it is far from clear what I'm saying *except* that I deny the whole utterance exactly as worded. Is it because I deny that it was Devine's cabinet members that were prosecuted? Or do I deny that those prosecuted were cabinet members? Or that they were criminally prosecuted, as opposed to civilly sued or not prosecuted at all? Do I deny that there were as many as three, or might I be taking exception to some understatement, on the grounds, say, that there were in fact *more* than three or that they were not simply prosecuted but convicted and sent to prison? Since the mere negation of a complex statement does not reveal such distinctions, the assessment of a negated premise or conclusion must be based on the least specific interpretation available, unless the surrounding context really does make clear what is intended. And by the same token, when we are constructing arguments ourselves, we should place our *not*s carefully to communicate our views as clearly as possible.

COMPLEX STATEMENTS

A crucial aspect of analyzing reasoning is understanding how much information a quite short statement can express—how many distinct claims are encoded by a grammatically complex sentence.

> Ted was cranky and annoying despite getting his way, as usual.

The truth-conditions of this sentence include all of the following:

1. Ted was cranky.
2. Ted was annoying.
3. Ted got his way.
4. Either Ted is usually cranky and annoying, or Ted usually gets his way, or both. (This is disjunctive because it's unclear what the phrase "as usual" modifies in the initial sentence.)

Hence if any of 1 to 4 is false, so too is the original sentence. A compound and complex sentence like this can end up embodying many distinct factual claims, each of which must be evaluated if the statement is offered as a premise or a conclusion.

Overall the lesson is that being really precise in one's claims requires great care in how one frames them. On the flip side, someone else's claims may require very careful evaluation to separate what's reasonable and clear in them from what is vague and unreasonable. An imprecise or ambiguous statement can seem plausible and significant on a quick reading, even when one interpretation of it is plausible but trivial, while another interpretation is exciting but implausible. Sensitivity to the precise effects of particular word choices and to the effects of context is a crucial ingredient in assessing claims that we encounter and in making reasonable statements ourselves. This is harder work than one might expect, given that we've used language most of our lives. But little of our everyday use of language provides us with the skills to write or speak with the great clarity that some situations require.

FACTUAL AND NON-FACTUAL STATEMENTS

Many thinkers in various fields have thought it plausible that certain kinds of statements do not really have truth conditions at all, or have very different truth conditions than they seem to have. The most-discussed class of statements for which these views have been widely held are value-theoretic statements: that is, statements involving moral concepts like right and wrong, good and evil, and statements involving aesthetic notions like beauty and ugliness.

Both within and without professional philosophy, many people have thought that moral statements either have no real truth conditions—they are neither true nor false—or they have truth conditions involving something other than what the statement seems to be about. Some people have thought that a statement like "Killing innocent people is morally wrong" is really just an expression of emotion, on a par with "*Booooo* to killing innocent people!" or "Killing innocent people? Yuck!" Since utterances of the latter two sorts arguably are neither true nor false, the same can be said of moral statements in general on this view. Others have taken ethical statements to contain an implicit reference to some behavioural code, so that our example would be more clearly rendered as "By the standards of my culture, killing innocent people is unacceptable." On this view our sample moral statement does have truth conditions, but they implicate facts about our culture more than any facts about killing *per se*. And on the aesthetic front, it has reached the status of folk wisdom that "beauty is in the eye of the beholder" or, as an older saying had it, *De gustibus non est disputandum*: on matters of taste there can be no argument.

None of these views is entirely accepted by philosophers today. Partly this is because it's difficult to spell out just what kinds of talk are and are not beyond the categories of truth and falsity. If we think of a statement's truth as consisting in some arrangement of physical things, for instance, on the model of the cat's being on the mat, then certainly it's difficult to think of some straightforward truth condition for the statement that something immoral has occurred. But it is equally hard to visualize just what arrangement of physical things would suffice to make true various other kinds of statement that may seem less contentious than the moral or aesthetic cases. What are the truth conditions for the claim that monetary inflation has occurred, for example? That there are now more ways than ever of proving some theorem? That celebrity is a fetish within the mainstream media? In each case we might answer the question by appeal to other concepts inherent to the relevant field—economics, logic, sociology—but so too in the moral case might we appeal to justice, fairness, rights, and so forth. Tracing the important differences between

ANOTHER EXAMPLE

The Dog Walking Ordinance[1]

The following reputed transcript of a borough council meeting in England illustrates the potential difficulty of finding precise and unambiguous language to express something that seems like a simple idea.

Councilor Trafford took exception to the proposed notice at the entrance of South Park: "No dogs must be brought to this Park except on a lead." He pointed out that this order would not prevent an owner from releasing his pets, or pet, from a lead when once safely inside the Park.

The Chairman (Colonel Vine): What alternative wording would you propose, Councilor?

Councilor Trafford: "Dogs are not allowed in this Park without leads."

Councilor Hogg: Mr. Chairman, I object. The order should be addressed to the owners, not to the dogs.

Councilor Trafford: That is a nice point. Very well then: "Owners of dogs are not allowed in this Park unless they keep them on leads."

Councilor Hogg: Mr. Chairman, I object. Strictly speaking, this would prevent me as a dog-owner from leaving my dog in the back-garden at home and walking with Mrs. Hogg across the Park.

Councilor Trafford: Mr. Chairman, I suggest that our legalistic friend be asked to redraft the notice himself.

Councilor Hogg: Mr. Chairman, since Councilor Trafford finds it so difficult to improve on my original wording, I accept. "Nobody without his dog on a lead is allowed in this Park."

Councilor Trafford: Mr. Chairman, I object. Strictly speaking, this notice would prevent me, as a citizen, who owns no dog, from walking in the Park without first acquiring one.

Councilor Hogg (with some warmth): Very simply, then: "Dogs must be led in this Park."

Councilor Trafford: Mr. Chairman, I object: this reads as if it were a general injunction to the Borough to lead their dogs into the Park.

Councilor Hogg interposed a remark for which he was called to order; upon his withdrawing it, it was directed to be expunged from the Minutes.

The Chairman: Councilor Trafford, Councilor Hogg has had three tries; you have had only two...

Councilor Trafford: "All dogs must be kept on leads in this Park."

The Chairman: I see Councilor Hogg rising quite rightly to raise another objection. May I anticipate him with another amendment: "All dogs in this Park must be kept on the lead."

This draft was put to the vote and carried unanimously, with two abstentions.

value-theoretic discourse like aesthetics or morality and other discourses is a subtle matter; in the absence of such subtlety it is unwarranted to dismiss moral and ethical claims as meaningless, or as lacking truth values.

For our purposes it is safest to treat moral discourse as differing from other discourses, not in being non-factual where they are factual, but in being less apt to general agreement on fundamental assumptions. Statements about morality, humour, etiquette, aesthetics, and other matters with some significant degree of subjectivity can be perfectly acceptable as premises in an argument, provided the interlocutors share the relevant assumptions. If at first they do not, then discussion may bring them around to agreement; otherwise, however, it may be difficult for them to have a meaningful exchange. But then, this much is true whenever fundamental assumptions are not shared, no matter what their topic.

This is not to say there are no lasting philosophical issues regarding the alleged contrast between **subjective** and **objective** statements or disciplines. I am, however, counselling that we treat the distinction (whatever it amounts to) as irrelevant, focusing instead on the more general communicative problem of settling on premises acceptable to those involved in a discussion. Labels like "objective" and "subjective" can easily do more harm than good, leading us to polarize the categories under discussion, so that we take the terms deemed objective to be immune to conceptual criticism while lumping the subjective issues in with matters of whim or fleeting conventions. While many writers over the past few decades have claimed that objectivity isn't all it's cracked up to be, an equally important observation is that strictly subjective ways of speaking—ones that implicate human responses or judgments in their truth-conditions—are also much more robust than a simplistic conception would have it. Judgments about colour, for example, are in some sense subjective. But if "Hurting the innocent is wrong" were no worse off truth-wise than "The sky is blue," there would be no obvious general reason to worry about arguments involving allegedly subjective moral statements!

People considering the following argument, for example, might agree on the first (non-moral) premise but disagree over the second (morally loaded) premise.

Mixed moral and non-moral premises in argument:

1. If we allow inflation to rise, then the shoe manufacturers will lose their jobs.
2. It would be unjust for the shoe manufacturers to lose their jobs.

Therefore,

3. If we allow inflation to rise, then something unjust will occur.

But then, they might just as easily disagree over the truth of the first premise while all accepting the second. Either way, the argument appears valid, and it seems entirely reasonable for someone defending the argument to offer reasoning in support of *either* disputed premise. So we end up treating value-theoretic statements much the same way we treat statements more generally: as open to disagreement, but effective in argument provided they are acceptable to all discussants. Reasoning about values is discussed in more detail in Chapter 3.

CONCLUSION

For the most part we want to believe reasonable things and to communicate our beliefs in a way that makes their reasonableness clear. Arguments are the basic units of these processes. The clearest and most idealized conception of argumentation is that of deductive argument, according to which a good argument is a sound one—an argument that is valid and has all true essential premises. An argument is valid just in case there is no way for its premises to be true and its conclusion false; the idea is that if the premises are true, they prove the conclusion. To say that the premises are true is to say that they describe the world accurately, or that they are laid down as commonly accepted definitions. Premises can be treated as true for argumentative purposes when they are jointly acceptable to both the speaker and audience. If they are not accepted by the audience then, whether or not the premises are true, further argument in support of them may be required as a practical measure for the argument to be rationally persuasive. The truth values of compound premises and conclusions are determined by the truth values of their subsentences and vary according to the way those subsentences are connected. Validity and soundness are technically all-or-nothing threshold concepts, though we will see how this idealization is usually only approximated to some degree in actual argumentation.

REVIEW QUESTIONS

1. For every valid argument form listed in this chapter, construct your own example.
2. Explain the concepts of linked, convergent, and sequential arguments. Which is the most fundamental kind, and why?
3. Read through recent national or local newspapers to find examples of ambiguous negation. For each example, explain the different possible interpretations of what is being denied.

NOTES

1. From R. Graves and A. Hodge, *The Reader Over Your Shoulder* (New York: Macmillan, 1943). Reprinted in E. Nagel, "Symbolic Notation, Haddocks' Eyes and the Dog-walking Ordinance," in *The World of Mathematics*, Vol. III, ed. J. Newman, 1890–91 (New York: Simon and Schuster, 1956).