CRITICAL THINKING

A CONCISE GUIDE

TRACY BOWELL

GARY KEMP

■ Critical Thinking

This is the best single text I have seen for addressing the level, presumptions, and interests of the non-specialist. The authors have a fine knack for articulating simply and clearly the most elementary – but also the most important – aspects of critical thinking in a way that should be clear to the novice.

Charles Ess, *Drury University*

Attempts to persuade us – to believe something, to do something, to buy something – are everywhere. What is less clear is how to think critically about such attempts and how to distinguish those that are sound arguments. *Critical Thinking: A Concise Guide* is a much-needed guide to argument analysis and a clear introduction to thinking clearly and rationally for oneself. Through clear and accessible discussion, this book equips students with the essential skills required to tell a good argument from a bad one.

Key features of the book include:

◆ Clear, jargon-free discussion of key concepts in argumentation.
◆ How to avoid common confusions surrounding words such as 'truth', 'knowledge' and 'opinion'.
◆ How to identify and evaluate the most common types of argument.
◆ How to spot fallacies in arguments and tell good reasoning from bad.
◆ Topical examples from politics, sport, medicine, music; chapter summaries; glossary and exercises throughout.

Critical Thinking: A Concise Guide is essential reading for anyone, student or professional, at work or in the classroom, seeking to improve their reasoning and arguing skills.

For further exercises and information, please go to www.routledge.com/textbooks/criticalthinking

Tracy Bowell is lecturer in Philosophy at the University of Waikato, New Zealand.

Dr Gary Kemp is lecturer in Philosophy at the University of Glasgow, UK.

Critical Thinking

A Concise Guide

■ **Tracy Bowell**

and

■ **Gary Kemp**

◼ Contents

Contents

■ Preface

Like all authors of texts on critical thinking or critical reasoning, we have tried to write a book that is genuinely useful. But our conception of what is useful differs somewhat from that of most of those authors.

On the one hand, we have avoided formal logical methods. Whereas the application of formal methods is justified primarily by its value in coping with complex logical structure, the logical structure of everyday argumentation is very seldom so complex that an argument's validity, or lack of it, cannot be revealed to ordinary intuition by a clear statement of the argument in English. Yet no formal means short of the first-order predicate calculus is sufficient to represent the logic of the majority of everyday arguments. Rather than compromise by presenting less comprehensive formal methods that are useful only in a narrow range of cases, we have avoided them entirely.

On the other hand, we have discussed and employed the *concepts* of logic more thoroughly than is customary in texts that avoid formal methods. We have defined them as accurately and in as much detail as we could, without superfluous refinement or inappropriate theoretical elaboration. We have done this for three reasons. First, it is only by grasping those concepts clearly that the student can achieve a stable and explicit understanding of the purposes of presenting and analysing arguments. Second, facility with those concepts enables the student to think and to talk about arguments in a systematically precise way; it provides a common currency in terms of which to generalise about arguments and to compare them. Third, experience, including our teaching experience, suggests that the concepts of logic themselves, when they explicitly appear in argumentative contexts, are amongst the most persistent sources of confusion. A symptom of this is the relativism that is so often encountered and so often lamented. At the root of this, we assume, are certain equivocations over the word 'truth'. We have tried to clear these up in

Preface

a commonsense and non-dogmatic way, and thereby to clarify further concepts that depend on the concept of truth, such as validity, probability, inductive force, soundness, justification and knowledge. We hope that clarity about these concepts, and the ability to use them with confidence in analysing arguments, will be amongst the most valuable accomplishments to be acquired by studying this book.

We do not entirely accept the view that examples in a book on critical thinking should be real, or even that they should be realistic. Of course, the aim is that students should be able to deal with real arguments. But whereas real examples typically call for the exercise of several strategies and the application of various concepts at once, those strategies and concepts have to be learnt one at a time. Unrealistic, trumped-up examples are often much more useful for illustrating isolated concepts and points of strategy. We have tried to vary the realistic with the unrealistic as the situation recommends.

Chapter summaries are provided along with exercises, but we have not provided answers to the exercises. One reason is that in many cases there is no demonstrably correct answer. Another is that in many cases there is too much scope for different sorts of answer, different ways of putting a point or representing an argument. And more generally, we think it is just too tempting for students to consult the answers too soon. We anticipate that teachers will want to assign further coursework, especially where this involves the procurement and analysis of material from newspapers and the like.

Every effort has been made to trace all copyright holders of extracts used in this book. If any have been inadvertently overlooked, the publishers will be pleased to make the necessary arrangements at the first opportunity.

Tracy Bowell, University of Waikato
Gary Kemp, University of Glasgow
January 2001

■ Introduction and preview

We are frequently confronted with people's attempts to get us to do and believe things, and often these attempts are arguments. This book aims to familiarise you with techniques of argument analysis and assessment so that you become able to assess the strength of any argument you may encounter and thus determine whether you ought to be persuaded by it.

Chapter 1 introduces the concept of argument as it should be understood for the purposes of critical thinking. Argument is distinguished from other linguistic means of getting people to do and to believe things. We introduce a method for laying out arguments so as to understand them more clearly, and we discuss various ways in which language can obscure an arguer's intended meaning.

Chapter 2 introduces two of the concepts required for the analysis and assessment of deductive arguments: validity and soundness. We discuss the assessment of validity and soundness, and explain the meaning and use of the Principle of Charity.

Chapter 3 continues our coverage of the concepts central to this book, this time for analysis and assessment of inductive arguments: inductive force and inductive soundness. We also discuss inductive inferences and degrees of probability.

Chapter 4 is a detailed discussion of rhetorical ploys and fallacies, two species of what we call 'sham-reasoning'. Common species of each are considered, and using the concepts and techniques covered in previous chapters, we provide a method for exposing and assessing fallacious reasoning.

Chapter 5 covers in more detail the techniques required for reconstructing arguments so that their strength can be accurately assessed. We demonstrate techniques for deciding which material is relevant to an argument; for dealing with ambiguous and vague language; for uncovering an argument's hidden premises; for adding connecting premises; for dealing with practical reasoning, and for dealing with causal arguments.

Introduction and preview

Chapter 6 is concerned with further concepts and techniques for argument assessment. We introduce the concept of rational persuasiveness, and introduce further techniques for assessing arguments and for refuting them. We also present argument trees as a graphical method of representing an argument's form.

Finally, in Chapter 7 we consider some of the philosophical issues underlying the concepts and techniques used here. We discuss truth and the relationship between truth, belief and knowledge, and relate these issues to the concept of rational persuasiveness.

Each chapter concludes with a chapter summary and some exercises. Where appropriate, the reader is encouraged to look outside the book for further examples to serve as exercises.

Chapter 1

Why should we become critical thinkers?

The focus of this book is written and spoken ways of persuading us to do things and to believe things. Every day we are bombarded with messages apparently telling us what to do or not to do, what to believe or not to believe: buy this soft drink; eat that breakfast cereal; vote for Mrs Bloggs; practise safe sex; don't drink and drive; don't use drugs; boycott goods from a particular country; abortion is murder; meat is murder; aliens have visited the earth; the economy is sound; capitalism is just; genetically modified crops are safe; etc. Some messages we just ignore, some we unreflectively obey, and some we unreflectively reject. Others we might think about and question, asking 'why should I do, or refrain from doing that?', or 'why should I believe that, or not believe it?'.

When we ask the question 'why?' we're asking for a **reason** for doing what we are being enjoined to do, or believe what we are being enjoined to believe: Why should I vote for Mrs Bloggs, or eat this particular breakfast cereal? Why should I believe that meat is murder, or that the economy is sound? When we ask for a reason in this way we are asking for a **justification** for taking the action recommended or accepting the belief – not just a reason, but a good reason that ought to motivate us to act or believe as we are recommended to do. We might be told, for example, that Wheetybites are a nutritious, sugar-free, low-fat breakfast cereal; if so, then if we want to eat a healthy breakfast, we've been given a good reason to eat Wheetybites. If, on the other hand, we are given only state-of-the-art marketing techniques – for example, images of good-looking people happily eating Wheetybites with bright red strawberries out of fashionable crockery – then, although an attempt has been made

to persuade us to buy Wheetybites, it would not appear that any attempt has been made to provide good reasons for doing so.

To attempt to persuade by giving good reasons is to give an **argument**. We encounter many different types of attempts to persuade. Not all of these are arguments, and one of the things we will concentrate on early in this book is how to distinguish attempts to persuade in which the speaker or writer intends to put forward an argument from those in which their intention is to persuade us by some means other than argument. Critical thinkers should primarily be interested in arguments and whether they succeed in providing us with good reasons for acting or believing. But we also need to consider non-argumentative attempts to persuade, as we must be able to distinguish these from arguments. This is not always straightforward, particularly as many attempts to persuade involve a mixture of various argumentative and non-argumentative techniques to get readers and listeners to accept a point of view or take a certain course of action.

You may find it surprising to think of an 'argument' as a term for giving someone a reason to do or believe something – telling them why they should boycott certain products or disapprove of fox-hunting for instance. Perhaps in your experience the word 'argument' means a disagreement – shouting the odds, slamming doors, insults, sulking. In fact in some of those situations the participants might actually be advancing what we mean by an argument, putting forward a well argued case for washing up one's dishes for example, but in many cases, they will not be arguing in the sense we have in mind here.

The sort of argument we have in mind does occur frequently in ordinary, everyday situations. It is by no means restricted to the works of Plato, Descartes and other scholars famous for the arguments they put forward. You and your acquaintances give each other reasons for believing something or doing something all the time – why we should expect our friend to be late for dinner, why we should walk rather than wait for the bus, and so on. Open a newspaper, and you'll find arguments in the letters section, editorials and various other discussion pieces. On television and radio broadcasts (especially current affairs shows) and in internet discussions you'll find people arguing their case (though they may well also resort to other persuasive techniques as well). The same thing occurs in a more elevated form at university and college. Throughout your time as a student you will hear lecturers and other students arguing for a point of view, and in set readings you will encounter attempts to persuade you of various claims about all manner of issues.

If you develop your ability to analyse people's attempts to persuade so that you can accurately interpret what they are saying or writing and

evaluate whether or not they are giving a good argument – whether, for example, they are providing you with a good reason to believe that fox-hunting should be banned – then you can begin to liberate yourself from unquestioningly accepting what others try to persuade you of without knowing whether you have a good reason to be persuaded.[1]

But then, you may ask, why is it liberating to demand reasons before you are persuaded to adopt new beliefs? Isn't it less trouble to go through life unreflectively doing more or less as you please and not worrying too much about whether you have good reason to do or believe something beyond whether or not you want to? Well, it may often be easier in the short run, but it might lead to a life dominated by bad decisions and discontentment. Socrates, the Ancient Athenian philosopher famously argued that 'the unexamined life is not worth living'.[2] While this may or may not be true, the only way to find out is to approach the issue in a critical and rational manner. Paying attention to arguments gets you (eventually) to the truth of a matter, thereby making the world and the people in it easier to comprehend and to deal with.

Even if a desire to discover the truth does not seem a sufficiently strong reason for being concerned about having good reasons to justify your actions and beliefs, there are various life situations in which the ability to interpret and evaluate someone's case properly may be crucial to someone's well-being, or even to their remaining alive. For example in a court trial the jury is instructed to convict an alleged murderer if the prosecution has proved their guilt beyond reasonable doubt. The jury is being asked to consider the prosecution's case (which is ideally an argumentative attempt to persuade them of the guilt of the accused), and the evidence they offer at each step of making that case. It has to consider whether there is good reason to accept the argument or whether some holes in it mean that there must be some doubt about its truth. The skills of evaluation and interpretation involved in argument analysis are what we use (or ought to use) in determining the strength of the prosecution's case in such situations. In fact in any situation in which we have to make decisions, be they about our lives or the lives of others, there is no substitute for the ability to think logically and to detect errors in the thinking of others.

---

1   Although this book emphasises the value of reason and the benefits of using techniques of persuasion that are rational, we should also bear in mind that what is claimed to be rational is not always rational, and certainly does not always have positive consequences. Historically, for example, those who wield power have often granted themselves authority over what counts as 'rational', condemning as 'irrational' what threatens the status quo. The correct response to that sort of rhetorical manoeuvre, however, is not to say 'so much the worse for rationality, then!'; the correct response is to question whether the charge of irrationality is justified, or whether the term is merely being abused or manipulated. Rationality in itself is a neutral force, independent of anyone's particular interests or beliefs.

2   Plato, *Apology*, 38a, Harmondsworth: Penguin, 1969, p. 72.

Why should we become critical thinkers?

Plato, the student of Socrates, provides us with many good examples of the power and value of thinking critically. Plato's dialogues are dramatisations of imaginary philosophical exchanges between the character Socrates and various other men of Athens and other Greek city-states. The character Socrates as he appears in Plato's dialogues was based on the actual man Socrates, who was something of a civic institution in Athens. He was not a professional teacher and wrote very little. Rather he preferred to express his ideas in discussion. He was a familiar figure in the public areas of the city – in the *agora* (the public meeting place), at the *gymnasium*, or simply walking the streets immersed in discussion with friends and followers. These discussions did not typically involve Socrates lecturing others about his own beliefs; rather he challenged the beliefs of others, leading them to re-examine their beliefs, assumptions and prejudices – getting them to consider whether they had good reasons for holding those beliefs. For Socrates, the first step towards developing knowledge and wisdom is coming to see how much less we know than we think we do.

However, the esteem in which Socrates was held by the young Plato was not universally felt. In challenging widely-held beliefs and received opinions, and daring to ask questions like, 'what makes life worth living?', Socrates was considered a threat by the Athenian political establishment. He was seen as undermining the state's authority through his questioning of accepted beliefs about such subjects as courage, justice and the good life. This is why he was eventually tried on trumped-up charges of 'corrupting the young', found guilty, and put to death. Thankfully these days, in the United Kingdom at least, society is sufficiently tolerant that critical thinking will not usually get you into trouble. But it is a good reflection of the importance of the skills you are developing that those who hold power sometimes fear the effects of those who can think critically about moral, social, economic and political issues.

Beginning to think critically: recognising arguments

We do many things with language – state a fact, ask a question, tell someone to do something, insult someone, praise someone, promise to do something, swear an oath, make a threat, tell a story, recite a poem, sing a song, say a character's lines in a play, cheer on a football team. Throughout this book we write about 'attempts to persuade' – by argument and by other means. As we've mentioned, not all attempts to persuade using language are attempts to persuade by argument. Others are attempts to persuade by means of rhetorical devices. In Chapter 4 we discuss the most common of these devices in detail. For the time being

we'll just make some remarks about rhetoric in general. For our purposes rhetoric is defined as follows:

> *Rhetoric*: Any verbal or written attempt to persuade someone to believe, desire or do something that does not attempt to give good reasons for the belief, desire or action, but attempts to motivate that belief, desire or action solely through the power of the words used.

The crucial thing to understand here is that an attempt to persuade by argument is an attempt to provide you with reasons for believing a claim, desiring something or doing something. Arguments appeal to your critical faculties, your reason. Rhetoric, on the other hand, tends to rely on the persuasive power of certain words and verbal techniques to influence your beliefs, desires and actions by appeal to your desires, fears and other feelings.

Threats and bribes are special cases that may appear to count as rhetoric according to our definition. In fact they are closer to argument; for they work by announcing to the recipient that they have a good reason to act as suggested. For example, if Smith attempts to persuade Jones to lend him her car by threatening to inform the police that she uses a fake driver's licence, then he is implicitly giving her a reason to lend him her car – if she doesn't do so, the police will find out about the driver's licence; since she doesn't want that to happen, she has a reason to lend him the car. Although threats and bribes may be immoral, and may motivate partly by appeal to our fears and greed (among other feelings), they do motivate through force of reason and for that reason do not count as rhetoric.

Rhetorical techniques can be manipulative and coercive and their use should generally be avoided by those who aspire to think critically and persuade by reason. That is not to say that rhetoric is always undesirable. Often it is used to great effect for good causes. Consider this excerpt from Sir Winston Churchill's famous speech to Parliament during the Second World War in which he attempts to rein in a sense of celebration at the success of the evacuations of British troops from Dunkirk, and to remind parliamentarians and the public generally that there was still a long way to go in defeating the Nazis and their allies. Churchill uses some remarkably effective rhetoric in a good cause and might well be admired as a talented rhetorician. But his speech does not amount to an attempt to persuade by argument.

Why should we become critical thinkers?

> The British Empire and the French Republic, linked together in their cause and in their need, will defend to the death their native soil, aiding each other like good comrades to the utmost of their strength. Even though large tracts of Europe and many old and famous States have fallen or may fall into the grip of the Gestapo and all the odious apparatus of Nazi rule, we shall not flag or fail. We shall go on to the end, we shall fight in France, we shall fight on the seas and oceans, we shall fight with growing confidence and growing strength in the air, we shall defend our Island, whatever the cost may be, we shall fight on the beaches, we shall fight on the landing grounds, we shall fight in the fields and in the streets, we shall fight in the hills; we shall never surrender, and even if, which I do not for a moment believe, this Island or a large part of it were subjugated and starving, then our Empire beyond the seas, armed and guarded by the British Fleet, would carry on the struggle, until in God's good time, the New World, with all its power and might, steps forth to the rescue and the liberation of the old.

On the other hand, those who try to persuade you of not such good causes might also be effective, persuasive rhetoricians. European fascists of the 1930s – Hitler, Mussolini, Franco – provide good examples of this.

Of attempts to persuade that are arguments, not all are good arguments. So when analysing attempts to persuade we have to perform three tasks:

- ◆ The crucial first stage involves distinguishing whether an argument is being presented. We need to **identify** the issue being discussed, and determine whether or not the writer or speaker is attempting to persuade by means of argument.
- ◆ Once we have established that the writer/speaker is presenting an argument, we can move to the task of **reconstructing** the argument so as to express it clearly, and so as to demonstrate clearly the steps and form of the argument's reasoning.
- ◆ A clear reconstruction makes our third and final stage – **evaluating** the argument, asking what's good about it and what's bad about it – much easier to perform and to justify.

In subsequent chapters we explain in detail what we mean by reconstruction, and explain what makes an argument a good one. Our aim is not to help you acquire the basic comprehension skills that you need to work out what a passage or speech is about. We assume that you already have that skill, though working through this book might help you to hone it more finely. So we will begin with the first step, by considering how to distinguish arguments from other ways of putting forward opinions and persuading people to act.

When we put forward an argument we are either advancing an opinion (a claim which we think is true) or recommending an action. In either case we give a number of claims intended to support the claim or the recommendation. However, these two types of arguments can be collapsed into one. For we can think of an argument that recommends an action as advancing a claim to the effect that the hearer or reader should, or ought to, do such and such. For example an argument whose aim is to get you to buy Wheetybites can be understood as advancing the claim: You ought to buy Wheetybites.

Thus all arguments can be understood as attempting to provide reasons for thinking that some claim is **true**. The nature of truth is a deep and controversial philosophical issue that we do not need to contemplate here. We are working with an ordinary, non-theoretical concept of truth such that to say that someone's claim is true is to say that what it says is how things really are. For example, if someone makes the true claim 'Moscow is larger than Paris', then according to our intuitive conception of truth, it is true just because Moscow *is* larger than Paris. Our working definition of truth then, is as follows:

> **To say that a claim is *true* is to say that what is claimed is how things actually are.**

A claim, however, does not constitute an argument. An argument needs more than one claim: it needs the claim of which the arguer hopes to convince his or her audience, plus at least one claim offered in support of that claim. To illustrate the difference between arguments and claims, consider these **unsupported claims**:

► It's going to rain later.

► The Labour Party is making a better job of running the country than the Conservative Party did.

► Philosophers are odd, unworldly people.

► The world is facing environmental catastrophe.

The following examples, by contrast, attempt to give some **support** for these claims. Whether they provide adequate support is something we will look at later. The important point is to see the difference between this set and the first set:

Why should we become critical thinkers?

▶ It's going to rain later, I know because I heard the weather forecast on the radio and it's usually reliable.

▶ The Labour Party is making a better job of running the country than the Conservatives did. Unemployment is down, prosperity is up and the pound remains strong. These are the crucial signs that the country is doing well.

▶ I've met a few philosophers in my time and they've always been strange people, heads in the clouds, not really in touch with the real world. Philosophers are odd, unworldly people.

▶ Climate scientists predict that the world is facing environmental catastrophe, and they are the experts on these issues.

There are special terms for the two parts of arguments: the primary claim, the one we are trying to get others to accept, is the **conclusion**. The supporting claims, the ones that are intended to give us reasons for accepting the conclusion are the **premises**. As with the word 'argument', we are using the word 'premise' here in a restricted way that does not correspond to all the ways the word is ordinarily used. People sometimes respond to someone's expression of opinion by saying, 'that's just your premise, but no one knows that for *sure*'; they do so to cast doubt on the truth of the claim being made. That is not the sense of the word 'premise' used in the discussion and analysis of arguments: for this purpose, a premise is simply any claim put forward as support for the conclusion of an argument, however certain or uncertain that claim may be.

We can now give a working definition of **argument**:

> An *argument*: A set of propositions of which one is a conclusion and the remainder are premises, intended as support for the conclusion.

And what's a proposition?

> A *proposition*: the factual content expressed by a declarative sentence on a particular occasion. The same proposition may be expressed by different sentences. For example, on a given occasion, 'The Government has

decided to hold a public enquiry into the affair' would
express the same proposition as 'It was decided that the
Government would hold a public enquiry into the affair.'

One upshot of this is that different sets of sentences could express the same argument. Another is that when a sentence is used rhetorically, its rhetorical aspect, which we will call its **rhetorical force**, is not part of the propositional content that it expresses; rather, it is the emotive or otherwise suggestive window-dressing which surrounds the proposition and is used to persuade us to believe or do something. The point is best grasped when we consider sentences that express the same proposition but have different rhetorical force. The sentence 'She is bringing up her children on her own' expresses the same proposition as the rhetorically charged 'She's a single mum'. But while the former merely expresses a fact about someone's family arrangements, the second, by its use of the emotive and politically significant term 'single mum' might function not only to inform us of a fact, but also to manipulate our sympathies concerning the person in question (depending upon our beliefs and feelings about parenthood).

An argument may be about any subject and have any number of premises, but it will always have only one final conclusion. This argument has just one premise:

Bart has two sisters.

Therefore, Bart is not an only child.

This has two:

Helping someone to commit suicide is the same as murder.

Murder is wrong.

Therefore, helping someone to commit suicide is wrong.

And this one three:

Car use is seriously damaging the environment.

Reducing car journeys would reduce damage to the environment.

We should do what we can to protect the environment.

Therefore, we should use cars less.

Why should we become critical thinkers?

As you can see, arguments for analysis are set out in a particular style with the premises listed in the order that they occur in the reasoning process and the conclusion appearing at the bottom. We can refine this style and further clarify the argument by numbering the premises P1, P2, and so on, and drawing a line between the last premise and the conclusion, which we mark with a 'C'. The line between premises and conclusion is called an **inference bar**, and its purpose is to distinguish steps in reasoning. The bar should be read as standing for 'therefore'. This style of setting out arguments is called **standard form**. The purpose of setting out arguments in this manner is to maximise clarity. Using this method helps us to see the stages of reasoning clearly and to make comparisons between arguments of similar form. When dealing with arguments as they are ordinarily presented, distinguishing the exact conclusion from the premises, the premises from each other, and the premises and conclusion from other, irrelevant material can be difficult. Writing the argument in standard form provides us with the most comprehensive and clearest possible view of it, ensuring that while discussing the argument and attempting to evaluate it, we do not lose track of exactly what the argument is.

A number of the exercises included in this book require you to set out arguments in standard form. To do this is to **reconstruct** the argument, and the end product – the argument set out in standard form – is called a **reconstruction** of the argument, or an **argument-reconstruction**. In reconstructing arguments you should follow the example below by taking these steps:

- ▶ Identify the conclusion.
- ▶ Identify the premises.
- ▶ Number the premises and write them out in order.
- ▶ Draw in the inference bar.
- ▶ Write out the conclusion, placing 'C' in front of it.

Thus the previous example looks thus in standard form:

P1) Car use is seriously damaging the environment.
P2) Reducing car journeys would reduce damage to the environment.
P3) We should do what we can to protect the environment.

C) We should use cars less.

Identifying conclusions and premises

The question of whether a passage or speech contains an argument is the question of whether the speaker or writer is attempting, by means of that passage or speech, to persuade his or her audience of some conclusion by offering premises in support of it. This is a question about the intentions of the writer or speaker – 'What does this person intend to do with these words here?' – that cannot always be answered unless we know something of the **context** – the circumstances in which the passage or speech appeared or took place. But even when we've determined that an argument is being advanced, its premises and conclusion are often buried deep amongst the other elements of a speech or text, and there are no hard and fast rules for distinguishing the propositions that form an argument from those that perform some other function in a text or speech. Identifying arguments is largely a matter of determining what the author or speaker intends by interpreting her words (spoken or written), and this comes with practice. Often writers and speakers leave some of their premises unstated because they think that readers or listeners will know what they have in mind. So in interpreting arguments we may have to add premises to make their structure and content complete. Further, people do not always express their arguments in very clear language, so we have to clarify each proposition before we can command a clear view of the argument as a whole (we look at difficulties with linguistic meaning later in this chapter).

Identifying conclusions

Once you have determined that a text or speech contains an attempt to persuade by argument, it is easiest to proceed first by identifying its conclusion. Determining whether a passage contains an attempt to persuade by argument and identifying the conclusion of that argument do not always occur independently however. Sometimes you will identify the conclusion in the process of working out that a passage does indeed contain an argument. On other occasions you may have already worked out that a passage contains an argument by paying careful attention to the writing style and the context without yet having identified the conclusion. We will, in any case, treat these processes as independent steps in argument analysis.

The conclusions of the following examples are probably clear from the first reading:

> Since Jo Bloggs is a politician and politicians are always corrupt, I guess Jo Bloggs is corrupt.

Why should we become critical thinkers?

> I'm anti-hunting because I believe that hunting foxes is wrong. After all, it's wrong to kill simply for pleasure and fox-hunting involves the killing of innocent animals for pleasure.

Before moving on, make sure that you can identify the conclusions in each of these examples.

Several points make the identification of conclusions an easier task

1   Once you have decided that a passage or speech contains an attempt to persuade by argument, try to see what the main point of the passage or speech is. Ask what point the speaker or author is trying to establish; that point will be the conclusion. Once you come to reconstruct an argument for analysis, paraphrasing the main point as one simple proposition will make the argument easier to handle. Bear in mind that a writer or speaker may make the same point in a number of different ways, so you may have to settle upon one particular way of expressing it.

2   Any proposition on any topic can be a conclusion. It is possible to attempt to argue for any claim, from the highly theoretical to the most mundane. So the type of subject matter of a proposition – religion, morality, science, the weather, politics, sport – is not in itself a guide to identifying whether or not that proposition is intended as the conclusion of a passage's argument. The premises and conclusions of arguments should ideally be expressed in declarative sentences, but in real-life arguments they may be expressed otherwise. When reconstructing arguments, we may need to rewrite premises and conclusions as declarative sentences in order to clarify the propositions expressed. For example, the apparent question, 'Aren't all socialists idealists?' might be used to express a premise that all socialists are idealists. The types of linguistic phenomena that need to be rewritten for clarity's sake are discussed in detail later in this chapter.

3   A single text or speech may contain several arguments for several different but connected conclusions. Sometimes we argue for one point, then a second, and then use those conclusions as premises in an argument for a third and final conclusion. These chains of arguments are known as **extended arguments** and we look at them in more detail later.

4   A helpful guide to recognising arguments are words which usually indicate that a writer or speaker is putting forward an argument. For

example if someone says, 'Given the facts that A, B, and C, **it follows that** D', you can be sure that D is the conclusion of the intended argument (and that A, B and C are the premises). Other common **conclusion indicators** are:

- Therefore . . .
- Hence . . .
- Thus . . .
- It can be concluded that . . .
- So . . .

Usually (though not in all cases) these words or phrases will follow the sentences that express an argument's premises. Another way of expressing an argument is to contain the premises and conclusion in one sentence with an indicator word separating them. For example, in the sentence 'The fact that Matthew Bordillo is a politician proves that he has a very big ego.', the conclusion that Mr Bordillo has a very big ego is separated from the premise that states that he is a politician by the indicator '**proves**'. Other words that serve the same function are:

- . . . implies . . .
- . . . establishes . . .
- . . . shows . . .

Commonly, a writer or speaker will state the conclusion of their argument before stating the premises. There are indicator words that are typically placed after the conclusion in these cases. For example in the sentence 'Gordon Brown must be a very important man since he is Chancellor of the Exchequer.', the conclusion that Mr Brown must be a very important man is separated from the premise stating that he is Chancellor of the Exchequer by the indicator word '**since**'. Other words and phrases that serve the same function are:

- . . . because . . .
- . . . for . . .
- . . . follows from the fact that . . .
- . . . is established by . . .
- . . . is implied by . . .

These indicators are not foolproof and should not be treated as a substitute for careful identification and interpretation of attempts to persuade by argument. Not all arguers will help the critical thinker out by making use of indicator words. The fact that a sentence does not include an

indicator word is not a reliable reason for thinking that it is not part of the expression of an argument contained in the text or speech under consideration. If a passage does not appear to have any **conclusion indicators** then an alternative way of trying to identify the conclusion is to try inserting conclusion indicators at appropriate places in sentences that appear to be good candidates for the conclusion. Then see if the passage or speech still reads or is heard smoothly and its meaning is unchanged. There are no conclusion indicators in the following speech, but it is still an attempt to argue:

> I think that Dinnah should sue the local council. They have admitted that they were negligent in not mending the cracked pavement that she tripped over when she broke her ankle and that's sufficient grounds for compensation.

Here if we try placing the conclusion indicator 'therefore' at the beginning of the second sentence ('They have admitted that they were negligent . . .'), it becomes clear that it is not the conclusion of the intended argument. Inserting 'because' between the first and second sentence (and thereby joining them to make one sentence), on the other hand, leaves the meaning intact and makes it clear that the conclusion – the claim that the speaker wants us to accept – appears at the beginning of the speech. Of course, when we write out the argument in standard form we change the order of the sentences and place the conclusion at the end preceded by the inference bar. Notice that the second sentence contains two premises so that in standard form the argument would be written thus:

P1) The local council has admitted negligence.
P2) An admission of negligence is sufficient grounds for compensating an injured party.

_____

C) The local council should compensate Dinnah.

5 So far we have only discussed **explicit conclusions** in which a writer or speaker expresses her conclusion directly and more or less clearly. However, there are occasions when conclusions remain unexpressed. These are **implicit conclusions**. They are only implied or suggested by the actual text or speech content, not explicitly expressed by it. This usually happens when the speaker or writer thinks that the context is sufficient to make the conclusion obvious so that it literally 'goes without

saying'. This is often a bad idea as the conclusion is not always as obvious to those whom one is trying to persuade as it is to the persuader. Try to avoid implicit conclusions in your own writing and speech. It isn't clear, for example, what, if any, conclusion is implicit in the following:

> There's so much pornography available on satellite television these days and young people are so easily influenced, it's bound to result in a social collapse into an orgy of rape, abuse and indecency.

Identifying premises

As you go through the process of identifying an argument's conclusion, it is likely that you will also spot some or all of its premises. Thus the stages of identification are not entirely separate. The identification of an argument's premises is a search for reasons given by the writer or speaker to think that their conclusion is true. Like the identification of conclusions, much of the process of identifying premises amounts to close and charitable reading of what a writer or speaker says; but again there are some helpful guides:

1   Ask yourself what the writer's or speaker's reasons for believing their conclusion are. What evidence does the writer or speaker give to think that the conclusion is true? The propositions that you come up with in response to these questions are likely to be the premises of the intended argument.

2   Like conclusions, premises can have any subject matter whatsoever. It does not matter whether a proposition is controversial or unanimously agreed, it can still be a premise.

3   Most real examples of writing and speech will embed their arguments within other language that is not intended as part of the argument itself, although some of this language may be used in what in Chapter 4 we call **sham-reasoning**. Again it helps to work out the overall structure of the passage when trying to identify the premises. Consider the following:

> I really think the Government should reconsider its policies on higher education. Education is such a complicated topic, and their policies are just more poll-driven nonsense; Blair and his cronies are so image-oriented with their expensive suits

> and so on, they invite pop stars to their parties and comport themselves as if they too were pop stars, just out to sell themselves really.

In this example the speaker gets side-tracked into commenting upon the prime minister's suits and party guest-lists, and fails, beyond the vague charge that the government's policies are 'poll-driven nonsense', to offer a substantive criticism. Most of what is said is at best only obliquely relevant to the issue.

4 As with conclusions, there are certain words that are usually (but not always) reliable indicators of the presence of premises – **premise indicators**. We have already seen some of these because they mark the speaker's or writer's move from premises to conclusion or from conclusion to premises ('since', 'because', 'is implied by', and so on.) There are other words and phrases that introduce sentences stating a premise or premises. A speaker or writer might state their conclusion and then begin the next proposition with such phrases as:

- My reason is . . .
- My evidence for this is . . .
- This is so because . . .

For example:

> I put it to you that Ms White killed Colonel Mustard in the ballroom with the candlestick. The reason I make this claim is that on the night of Colonel Mustard's death Lady Scarlet saw Ms White in the ballroom beating Colonel Mustard over the head with the very candlestick that was later found to have Ms White's fingerprints and Colonel Mustard's blood on it.

Other premise indicators may occur at the beginning of a sentence containing both the premise and the conclusion. For example:

> On the basis of the fact that they have promised big tax cuts, I conclude that the Conservative Party will probably win the next general election.

We should be aware, however, that **indicator words** can be used with other meanings. The sentence, 'Since 2000 I have been a student at the University of Anytown' does not express the conclusion of an argument because 'since' is used here to talk about a period of time. 'Because' can

be particularly tricky as it is frequently used in **explanations** that are not intended as attempts to persuade. These are cases where we have to think hard about the context of the word's use and what is the most likely intention of the speaker or writer. We need to work out whether they are telling us that such-and-such event occurred as a result of some other event – that is, whether they intend to assert a causal relationship. For in that case, 'because' is being used to introduce an explanation, not an argument. For example:

> The tap is leaking because it needs a new washer.

This is an explanation of why the tap is leaking; 'because' is used to indicate a causal relationship rather than a connection between premise and conclusion.

This dual role of 'because' – its role in both arguments and explanations – can be confusing where the explanation of actions is concerned. Suppose you are driving fast and your passenger asks, 'Why are you driving so fast?' You assume your passenger is not in any way suggesting that you shouldn't drive so fast. You think they don't mind in the slightest. You assume they are merely curious as to why you're driving fast – whether it's because you're late, being chased by the police, or perhaps testing the limits of your new car. Your reply, however, is simply 'Because I enjoy it'. This would be an explanation: you are telling your passenger why you're driving fast, not trying to persuade them of some further proposition.

But suppose when your passenger asked 'Why are you driving so fast?', you didn't assume that they didn't mind your driving fast. You think perhaps they do mind. So you take the question as demanding a **justification** for your driving so fast. If you now say 'Because I enjoy it.', then you would be arguing, roughly, that it is all right to drive at such a speed on the grounds that you have a right to do what you like. In that case 'Because I enjoy it.' would be a premise of an argument, which might initially be expressed thus:

> It's OK for me to drive as fast as I like, because I like driving fast. I think we should be free to do anything that we enjoy.

It might be re-written thus in standard form:

P1) I enjoy driving fast.
P2) It is acceptable for me to do anything I enjoy.
_____

C) It is acceptable for me to drive fast.

Why should we become critical thinkers?

As demonstrated by the following examples, 'Since', 'thus' 'so' and 'therefore' may also be used in explanations that are not intended to provide reasons for acting or believing something:[3]

▶ Since we forgot to add yeast, the bread didn't rise.

▶ We forgot to add yeast, therefore the bread didn't rise.

▶ We forgot to add yeast, thus the bread didn't rise.

▶ We forgot to add yeast, so the bread didn't rise.

5 Again, as with conclusions, a text or speech may not include specific premise indicators. Context is the best means of identifying premises in such cases. It may also help to try adding premise indicators to propositions to see if the passage or speech still runs smoothly.

6 Ordinary language can make identifying arguments more difficult than it might otherwise be because people do not always express all of their premises explicitly. Thus many attempts to persuade by argument rely on **implicit premises**. In Chapter 5 we will discuss the interpretation of hidden premises and the reconstruction of arguments to include them.

Intermediate conclusions

The conclusion of one argument may serve as a premise of a subsequent argument. The conclusion of that argument may itself serve as a premise for another argument, and so on. A simple illustration:

Fido is a dog. All dogs are mammals, so Fido is a mammal. And since all mammals are warm-blooded, it follows that Fido is warm-blooded.

In this argument, an **intermediate** conclusion – that Fido is a mammal – is used as a premise for a further argument, whose conclusion is that Fido is warm-blooded. We represent **extended arguments** of this kind like this:

---

3   While reading this book you may also have noticed a further use of 'thus'. 'Thus' can be used to mean 'in this way' and often precedes an example or a quotation.

P1) Fido is a dog.
P2) All dogs are mammals.

C1) Fido is a mammal.

P3) All mammals are warm-blooded.

C2) Fido is warm-blooded.

We give the two conclusions numbers: C1 is the conclusion of an argument whose premises are P1 and P2; C2 is the conclusion of an argument whose premises are C1 and P3. So C1 is both the conclusion of one argument and the premise of another.

Normally, in such cases, the last conclusion reached (the one with the highest number) is the proposition which the arguer is most concerned to establish. It is the ultimate target. So we call this simply *the* conclusion of the argument, whereas any other conclusions, reached as steps along the way, are called **intermediate conclusions**.

We sometimes want to concentrate for a moment on a particular part of an extended argument. In the above case, for example, we might be particularly interested either in the first part of the argument, or in the second. We will sometimes speak of the argument *from* P1 and P2 to C1, or of the argument *from* C1 and P3 to C2. We can also speak of the **inference** from P1 and P2 to C1, and the **inference** from C1 and P3 to C2.

The use of the word 'inference' in logic and critical thinking is another case where a word is used in a somewhat restricted sense in comparison with ordinary language. All reasoning consists of inferences, in the logician's sense of the word – each step of reasoning, each jump from premise or premises to conclusion, is an inference. Contrary to the way the word is often employed ordinarily, there need be nothing doubtful about an inference. We sometimes say, 'but that's just an *inference*', meaning to cast doubt upon whether a given proposition should really be accepted on the basis of others. But in our sense of the word, an inference may be completely certain, not subject to doubt. For example it is an inference, in our sense, to go from 'John is a classical musician' to 'John is a musician' – despite the fact that there can be no doubt that if the first proposition is true, then so is the second (in the terminology to be introduced in Chapter 2, it is a **valid** inference).

Why should we become critical thinkers?

Linguistic phenomena

As we've seen, once we've determined that a text or a speech contains an attempt to persuade by argument, the remainder of argument reconstruction is largely a matter of interpreting the speech or text as accurately as possible. Here we are trying to work out what the speaker or writer intends readers or listeners to understand and consequently do or believe on hearing or reading their words. Phenomena in ordinary language sometimes make this task more difficult because they obscure speakers' and writers' intended meanings and therefore make it difficult to tell which proposition their sentences are supposed to convey. So aspirant critical thinkers need to be aware of the ways in which language can work to hide writers' and speakers' meanings and must practise spotting potentially problematic sentences. At this stage you should aim to be able to recognise these sentences and to be able to give the possible interpretations of them; that is, the propositions that they could be used to convey.

Ambiguity

A sentence is **ambiguous** in a given context when there is more than one possible way of interpreting it in that context – that is, if there is more than one proposition it could plausibly be taken to express in that context. There are two types of ambiguity.

Lexical ambiguity

This is a property of individual words and phrases that occurs when the word or phrase has more than one meaning. The set or group of things to which an expression applies is called its **extension** (it helps to think of an extension as all the things over which the word or phase extends or spreads itself). Thus the extension of the word 'student' is the set of all students. An ambiguous word or phrase, then, has two or more separate and different extensions – it picks out two or more different sets of things. Ambiguous words and phrases can bring their ambiguity into sentences, making those sentences capable of having more than one possible interpretation. The word 'match' is one such word. The sentence 'He is looking for a match' could be intended to mean any of the following propositions:

> ▶ He is looking for a small stick of wood with an inflammable tip.

▶ He is looking for another one the same [as this one].

▶ He is looking for [wants] a game of tennis (or some such).

Notice that it is not only nouns that can be lexically ambiguous. Suppose you are going to meet someone for the first time and all you've been told about them is what a friend has told you – 'She's a hard woman'. This could mean: She's a difficult person; She's an aggressive person; She has a very well-toned, muscular body. Whichever interpretation you adopt will have an important effect on your expectations of the woman in question. When interpreting sentences that are lexically ambiguous, we have to focus on the context in which they are written or said and the consequent probability of each of the possible interpretations being the correct one. For instance the sentence 'A visitor to the zoo was attacked by the penguins' is lexically ambiguous because the preposition 'by' has two possible meanings in this context. The sentence could express either of the following propositions:

▶ The penguins attacked a visitor.

▶ A visitor was attacked beside the penguins' enclosure.

However, in the absence of any information about a vicious penguin, and given what we know about the usually non-aggressive behaviour of penguins towards zoo visitors, it would probably be reasonable to interpret the sentence as intended to express the second proposition.

There are a few words that are not really ambiguous but may seem so when we hear them, though not when we see them written. This is because the words, though spelt differently, sound the same. For example when heard, as opposed to read, the question 'Are you a mussel (muscle) man?' could be either an enquiry as to someone's taste in seafood or as to his physique. Of course, once we see the question written, we are in no doubt as to its meaning.

The examples considered so far are relatively simple to understand because the alternative meanings of words such as 'match' and 'hard' are very different. However, instances of lexical ambiguity also occur when a word has alternative meanings that are much closer together. Such cases are much harder to interpret and we need to pay a lot of attention to the context in which the word is being used and to the probability of the speaker or writer intending one interpretation rather than the other.

Suppose someone argues, 'Fewer women have the ability to do complex mathematics than men'. The speaker or writer might mean to say that men have a greater **natural** or **innate capacity** for complex mathematics than women; that is they might want to express the idea

that if the same number of women as men were to study complex mathematics under the same conditions, more men than women would succeed. If that were the intended interpretation, the claim would be sexist if it unjustifiably assumed that gender is a biological condition for success in mathematics. On the other hand, the speaker or writer might have intended to claim that as a matter of actual fact, there are fewer women than men successfully working at complex mathematics. This is true, but it is much less likely to be intended or interpreted as sexist: the relevant statistics are easily obtainable, and indeed might be cited by opponents of sexism as evidence that social factors encourage men, but discourage women, from specialising in mathematics.

The ambiguity here is due to the use of the word 'ability'. 'Ability' can refer to one's natural potential to do something – a potential with which one is born, and which may require training before engendering the actual capacity to do the thing – or it can refer to an actual capacity immediately to do the thing, a capacity which one may have acquired partly or wholly by training or practice. Often, it is not clear which meaning the speaker or writer intends. As you can see, the two ways of interpreting the word 'ability' are not entirely unrelated, as they are in the case of 'match'. Nonetheless, correct interpretation of which meaning a speaker or writer wishes to convey is crucial in determining how one should evaluate and respond to their remark.

Syntactic ambiguity

This occurs when the arrangement of words in a sentence is such that the sentence could be understood in more than one way (as expressing more than one proposition.) You will probably be familiar with examples of syntactic ambiguity as it is often the basis of jokes and newspaper headlines that appear odd. For example '33-year old Mrs Jones admitted to dangerous driving in Leeds Crown Court yesterday.' could mean either of the following:

- ▶ In Leeds Crown Court yesterday Mrs Jones admitted to dangerous driving.
- ▶ Yesterday Mrs Jones admitted to driving dangerously inside Leeds Crown Court itself.

The sentence is syntactically ambiguous because it could, consistently with English grammar, be used to express either proposition. But since the second interpretation is extremely unlikely, it is unlikely that an actual use of this sentence would be ambiguous. But consider this case:

> US President Bush has cancelled a trip to Scotland to play golf.

We can easily imagine a real context in which this sentence is ambiguous as to whether the purpose of the cancelled trip was to play golf or whether the trip was cancelled so that the president could play golf.

Once we decide the most likely interpretation, we should always rewrite the ambiguous sentence so as to eliminate the ambiguity. For example we might rewrite the above sentence thus:

> In order for him to play golf, US President Bush trip to Scotland has been cancelled.

Notice that in cases such as this we have to change the sentence quite radically to rid it of the syntactic ambiguity and clarify its meaning. Consider a further example:

> The government will announce that the electricity supply is to be cut off tomorrow.

The sentence leaves ambiguous the question of when the announcement will be made and when the electricity supply is to be cut off:

> ► Tomorrow, the government will announce that the electricity supply is to be cut off. (The announcement will be made tomorrow.)

> ► The government is going to announce that, tomorrow, the electricity supply will be cut off. (The announcement will be made now, the electricity will be cut off tomorrow.)

Syntactic ambiguities are sometimes more difficult than lexical ones to interpret on the basis of context. Also, the possible interpretations of a sentence may be closely related so that there may not appear to be a very wide difference in meaning. Often we assume that one interpretation is intended without giving any consideration to alternatives. But such differences can be very significant indeed. Suppose someone were to claim:

> We should not tolerate those homeless people living on our streets.

They might be saying that we should be intolerant of homeless people themselves. Or they might be saying that the people who do live on the

streets should not be allowed to live on the street. On the other hand, the intended proposition might be that we should not tolerate the fact that there are homeless people living on our streets. That is to say, the view expressed might be critical of a society in which people are forced to live on the streets rather than critical of such people themselves.

Vagueness

Vagueness is a property of words and phrases. It is not the same as ambiguity, but it is often mistaken for it. For instance, when former US President Clinton famously said, 'I did not have sexual relations with that woman . . .' he was not (as alleged) hiding behind the ambiguity of the phrase 'sexual relations', but rather behind its vagueness. As we saw when considering lexical ambiguity, a word is ambiguous when it has two or more possible and different meanings – thus two or more separate extensions. The vagueness of a word, on the other hand, is really a feature of its meaning: The meaning of a word or expression is vague if it is indefinite or uncertain what is conveyed by the word. Thus a word may be ambiguous without being vague – as in 'ball' (round plaything, formal dancing-party) – or vague without being ambiguous, as in 'sexual relations' (what exactly constitutes sexual relations?).

Sometimes, someone aware of the weakness of their own position will deliberately leave their meaning vague in order to camouflage that weakness and to evoke strong feelings of approval or disapproval in their readers or listeners. Many highly-charged words that wield rhetorical power in public discourse are used vaguely. Examples include: 'rights', 'liberal', 'harassment', 'racism', 'sexism'. It is hard to discern one perfectly exact meaning for each of these words and it would be unrealistic to expect them to have such a meaning. Their extensions tend to include a cluster of objects, beliefs or actions that are not necessarily unified in any precise way. Take 'liberal', for instance. This word conveys various characteristics including:

▶ Belief in a permissive society.

▶ Belief in freedom of speech, of association, of choice.

▶ Belief that certain restrictive laws should be relaxed (e.g. against drugs).

▶ Belief that the state should interfere as little as possible in citizens' lives.

▶ Belief in laissez-faire economic policies.

- ► Supports the Liberal Democratic Party.
- ► Not strict.
- ► Politically left-wing.
- ► Wishy-washy.
- ► Soft on crime.

One might be a liberal and not hold all of these beliefs, or have all of these characteristics. Indeed, someone might have some or even many of them and not be a liberal.

Here is a whole passage infected with vagueness of the kind we have in mind:

> Make no mistake, the researchers involved in the highly controversial project to map the human genome are involved in a radical project of unprecedented gravity and spiritual significance. Do they venture there with appropriate caution and humility? What they are doing is not even comparable to the research that made the atomic bomb possible, for it goes right to the essence of what we are. Like Dr Frankenstein, they are tinkering with life; they are travelling into unknown and sacred regions as no scientist previously has ever dared. The secret wellsprings of life, of our very being as *Homo sapiens*, have ever remained shut up, concealed by aeons of either blind but cunning and ultimately unfathomable natural processes, or, as some continue to believe despite the showy displays of science and technology, concealed by the very hand of its Author, the Author of Nature Himself.

What is the writer of this rather over-excited dose of hyperbole trying to argue? Clearly they think that there is something dangerous or otherwise ill-advised about the project to map the human genome. But they have not begun to make it clear what the danger is. The research is distinguished from atomic research by its concern specifically with life, but nothing is said as to why this is peculiarly dangerous beyond the use of extremely vague verbiage such as 'sacred', 'radical', 'gravity', 'spiritual significance', and so on. In a context such as this, with so much at stake, we need to have precise reasons why, despite the promise of medical benefits, the project is dangerous.

Words can also be vague in another, more philosophically technical respect. Philosophers of language use the term 'vague' to apply to words that have a clear meaning, but which have an indefinitely demarcated extension. Obvious cases are colour-words like 'orange': there is no precise

division between orange things and yellow things, for example. Things can often be precisely compared with respect to such attributes, however. For example X may be more bald, or fatter, sleepier, taller or faster than Y, even if it is not definite whether or not X is bald, fat, sleepy, tall or fast. Borderline cases can also arise in the case of nouns. In fact vagueness occurs in many more cases than we might at first think. Take 'city'. York is normally said to be a city, but is it really? What about Doncaster? Lancaster? Harrogate? Carlisle?

To a great extent, we take these sorts of vagueness in our stride, having become used unreflectively to interpreting these phenomena in ordinary language. But even the simplest cases can cause misunderstandings. Suppose your boss promises that you're going to receive a 'big pay rise' this year. When you receive the pay increase you discover that the rise is only ten pence an hour. When you complain, your boss defends their promise by saying that the rise is bigger than last year's and therefore big in comparison (see the section on **implicit relativity** for further discussion of such cases).

Primary and secondary connotation

The rich secondary connotation of some words provides a further source of vagueness. Every ordinary noun and every adjective – 'elephant', 'immoral', 'company', 'stupid' – has a range of things to which it applies: the **extension** of the term. The set of all bananas constitutes the extension of 'banana'; the set of all square things constitutes the extension of 'square'. A given thing falls within a word's extension if and only if it fits a certain rule associated with that word. For example the rule for the noun 'ram' is 'male sheep'. This rule is called the **primary connotation** of the term. This will be some set of characteristics, in this case being male and being a sheep which, by definition, everything to which the word applies must have. All of a term's primary connotation must apply to an object for that term to apply to it. The notion of a female ram, for example, is a logical impossibility, a contradiction.

However, when we are told that something is a ram, we tend to assume other things about that thing that are not included in the primary connotation: that it is woolly, has horns, lives on a mountainside or in a field, eats grass . . . So if you know that something is a ram, it is reasonable to suppose that it has these additional characteristics. These further characteristics that the term 'ram' also conveys make up its **secondary connotation**. Things which fall under the term will generally exhibit these characteristics, but there is no logical contradiction in supposing there to be a thing that falls under the term but lacks a characteristic

included under the secondary connotation. For instance, there is no logical contradiction in supposing that a thing might count as a ram – that is, fulfil the demands of the primary connotation – yet lack some or, indeed, all of these characteristics. It is not logically impossible that there could be a bald, hornless male sheep that lives in a barn and whose diet consists of potatoes.

Why should critical thinkers be interested in the distinction between primary and secondary connotation? The most immediate relevance was demonstrated in our examination of vagueness. It is difficult to pin down the precise meaning of a word such as 'liberal' because, on the one hand, its primary connotation is very difficult to pin down, and on the other, its secondary connotation is so rich. Take a look back at the list given earlier of characteristics conveyed by 'liberal': it is difficult to say which are part of its primary connotation, and which are only part of its secondary connotation.

A further reason for us to concern ourselves with this distinction is because it is the secondary connotation of many words that gives them their rhetorical power. Consider the noun 'feminist'. Its primary connotation is difficult to pin down and it is full of secondary connotations that can be used to the rhetorical advantage of both those who support and those who oppose feminism. Here are just some of the characteristics our critical thinking students have come up with when asked what the word 'feminist' conveys to them:

- Man-hating
- Lesbian
- Dungarees
- Unshaven
- Strong
- Political
- Fighter
- Staunch
- Left-wing
- Pro-abortion
- Pro-women

Rhetorical questions

Rhetorical questions take the form of a question but indirectly assert a proposition (like a declarative sentence does). That is, they are not really used to ask a question, but to make a point in an indirect way. Speakers and writers often use rhetorical questions when they're making a point they assume to be obvious, so the answer to the question 'goes without saying'. However, in many cases the point is neither obvious nor universally agreed. Rhetorical questions obfuscate speakers' and writers' intended meanings because they make it more difficult to interpret whether or not a speaker/ writer really does support a given claim. Rhetorical questions are common

in polemical newspaper articles and in readers' letters to an editor. If you encounter rhetorical questions in texts and speech that you are analysing, try to rewrite the question as a declarative sentence. For instance, if someone were to write:

> Should my right to freedom of speech be limited just because you disagree with me?

they probably wish to convey the proposition that their freedom should not be so curtailed, so they are not genuinely asking a question. They expect that the reader's response will be an automatic 'No, of course not'. To convey the proposition that seems to be intended, we could rewrite the rhetorical question as a declarative sentence:

> My right to freedom of speech should not be limited just because you disagree with me.

Irony

Speakers and writers sometimes express their claims using irony. This takes the form of language that, taken literally, would convey the opposite of what they wish to convey, or something otherwise very different from it. Consider the following instance:

> It is pouring with rain, very windy and cold. Mr I. Ronic says, 'Mmm lovely weather today'.

Mr Ronic is probably being ironic, and intends to comment that the weather is lousy.

It is important to be aware of the possibility of irony. In order to ridicule a position they are opposed to, speakers and writers sometimes sarcastically pretend to espouse that position; but it isn't always obvious that they are doing so.

Implicitly relative sentences

Consider the following examples:

► She earns an above average salary.

► He is of average intelligence.

► Great Aunt Edie is a fast runner.

- ▶ Taxes are high.
- ▶ The rent on our flat is low.

Sentences such as these represent another potential problem for the critical thinker striving to work out exactly what a speaker or writer intends to convey by their words. The sentences are **implicitly relative**. They make a comparison with some group of things, but that comparison is not explicitly mentioned. For instance, to understand what it is for someone to earn an 'above average salary', we need to know of what group the average to which it is compared is. Or consider the one about Aunt Edie? Does the speaker intend to convey that Great Aunt Edie is a fast runner such that she runs at world record pace or that she is a fast runner for a woman of her age? Or something in between, such as that she is faster than the average person? If such sentences are interpreted without the recognition of their implicit relativity, then there is the possibility that they will be interpreted as making a comparison with a group other than that intended by the writer or speaker. Great Aunt Edie is not a fast runner when compared with Linford Christie and, thus interpreted, the claim would be false. But when compared with other ninety-four year olds, many of whom haven't broken into a run for many years, she is a fast runner and the claim is true. Once we recognise such claims as implicitly relative and interpret them accordingly, they are more likely to have a definite truth-value. But not always. Implicit relativity is often compounded by other sources of vagueness. For example, even if we do know what comparison class is being invoked in the case of Aunt Edie, it is by no means clear just how much faster a person must be than the average person of that class in order to be fast relative to it.

Problems with quantifiers

Quantifiers are words that tell us how many/much of something there are/is, or how often something happens. As you will see, not all quantifiers specify an exact quantity of the thing, rather they provide a rough guide. In the following examples the quantifiers are underlined (this is not an exhaustive list of quantifiers):

- ▶ <u>All</u> men drive too fast.
- ▶ Members of Parliament are <u>often</u> self-serving.
- ▶ <u>Few</u> doctors support the health reforms.
- ▶ The lecturer awarded <u>almost all</u> the assignments an A grade.

Why should we become critical thinkers?

- ▶ <u>Nearly all</u> the students passed the course.

- ▶ She likes <u>hardly any</u> of her fellow students.

- ▶ <u>No</u> examiner should take bribes.

- ▶ <u>Lots of</u> computers develop faults.

- ▶ <u>Nine</u> hospitals will close at the end of this year.

- ▶ She <u>never</u> closes the door behind her.

- ▶ There are adequate computing facilities in <u>fewer than half</u> of the country's schools.

- ▶ He <u>always</u> writes his own speeches.

- ▶ <u>Most</u> women would choose to stay at home with their children if they could afford to.

There are four potential problems with quantifiers:

1 Speakers and writers don't always use quantifiers with sufficient precision, so that the proposition they intend to convey is unclear and open to misinterpretation and rhetorical abuse. Suppose your friend says, 'Premiership footballers all earn massive payments from sponsorship deals'. You don't agree and you mention an exception – Fergie Footballer receives only his footballer's salary, with no extra money from endorsing sports shoes or shirts. Suppose your friend defends her claim by saying that she didn't really mean that every single player in the Premiership earns big money from endorsements and sponsorship, but only that *most* or *nearly all* of them supplement their earnings in this way. Now that her claim is clear, you see that it is one with which you are more likely to agree.

2 Some quantifier words are themselves **vague**. Suppose, for instance, that someone claims

Some Members of Parliament support the decriminalisation of cannabis use.

What does 'some' mean here? It could mean that only a handful hold the view described, it could mean that a larger minority of members hold that view. Without a more precise understanding of how many Members of Parliament are intended to be conveyed by 'some', it is difficult to know how to respond to the claim. Moreover, the claim is open to abuse from people who hold views on both sides of such a debate. Advocates

of decriminalisation can use it in support of their cause; their opponents can use it to back up their anti-decriminalisation stance (the latter might say, 'Only some members of parliament support. . .').

3 Often people simply *omit* quantifiers. For instance, someone might protest:

> Lecturers don't give students a chance to complain.

At face value this might appear to convey the proposition that

> No lecturer (ever) gives a student a chance to complain.

Yet it is likely that what the speaker really wants to say is something like

> Most of the lecturers I've encountered haven't given students
> enough chance to complain.

Notice that once the appropriate quantifier is made explicit, the claim applies to a much smaller group of lecturers than one might have supposed when the quantifier remained implicit.

Consider another example:

> Today's students are dedicated to their studies.

If we interpret this as expressing this proposition:

> All of today's students are dedicated to their studies.

we are likely to want to challenge the claim as we will be able to cite exceptions to the generalisation. If, however, we interpret the claim as it is more likely to be intended, then the quantifier that we make explicit should be 'most' or 'almost all', thereby exposing the proposition really intended as

> Most of today's students are dedicated to their studies.

and this proposition has a greater likelihood of being true. Cases that we use to challenge the truth of a generalising claim are known as **counter-examples**. (The process discussed here should not be confused with that of refuting a complete argument by counter-example, which is dealt with in Chapter 6.)

Why should we become critical thinkers?

4   Quantifiers are used to express **generalisations**, and even where at least one quantifier is used explicitly, it can be unclear exactly what generalisation is intended. Consider the claim

> Everyone has tried drugs at some time in their lives.

Taken literally, the claim is a generalisation stating that every single person in the world has tried some drug or other at least once.[4] Even if we give 'drugs' the widest possible interpretation, the claim will be falsified by those cases of people who, for instance, have never even taken an aspirin for a headache. Of course, when challenged with such cases, the speaker or writer will most probably qualify their claim by saying that they didn't really mean literally everyone, but most of the people they know. So in order accurately to convey their meaning, the claim needs to be rewritten to make explicit to whom it applies. We do this by changing the quantifier to one that denotes fewer people than 'everyone' and by being more specific about to whom the claim is intended to apply. (This latter move is called 'restricting the scope of the generalisation'; we deal with it in more detail in Chapter 6.) The claim thus becomes,

> Most of the people I know have tried drugs at some time in their lives.

Notwithstanding the vagueness of 'drugs' (noted earlier), the claim now stands a better chance of being true than it did when it was expressed using an inappropriate quantifier.

Quantifiers: further issues

It is a commonplace for people to say that you can never really generalise. However, this is certainly not true just as it stands. When someone says this they might be understood as claiming, 'all generalisations are false'. But this is itself a generalisation; so if the claim is true, the claim is false! So that can't be what it means. It is in any case obvious that some generalisations are true (even if they are not very interesting ones). That is, there are counter-examples to the claim that all generalisations are false. For example, 'All cities in the UK have a bus service' is true and no case could really be raised to undermine its truth.

---

4   As you may have noticed, the word 'drugs' is used vaguely in this claim. Although in the context of most arguments about legislation and criminality it means illegal drugs, it could also include alcohol, prescription medicines, pain-killers, nicotine, and so on. Deliberately vague use of words in such a way constitutes either the rhetorical ploy of trading on an equivocation; see pp. 108–109 or the fallacy of equivocation; see pp. 140–141.

To get a better grasp of which types of generalisations may cause problems during the analysis and assessment of arguments, we need to distinguish between **hard** and **soft** generalisations. Consider the following generalisations (note that few of them have explicit quantifiers):

- ▶ Private schools attain better examination results than state schools.

- ▶ Traffic congestion is bad in Glasgow.

- ▶ Regular exercise benefits your health.

- ▶ Labour voters support a ban on hunting with hounds.

- ▶ People play less sport when they get older.

No doubt the counter-example fanatic will be able to provide us with plenty of exceptions – congestion-free short cuts across Glasgow; labour voters who are keen fox-hunters; the person who had a heart attack while doing their regular work-out at the gym. And they can cite this as a reason to accept the claim that all generalisations are false because one can always find an exception to them. However, to do so would be to misinterpret what people usually intend to convey when they say or write such things. It's rare for someone to mean that these sorts of generalisations are true without exception. The quantifier they intend to imply is probably one that is not synonymous with 'all' or 'every', but one such as 'in most cases', 'usually', or 'almost all'. These generalisations are **soft generalisations**. We use soft generalisations when we want to express the idea that such and such is true of certain things **normally**, **typically**, **generally**, **usually**, **on average**, **for the most part**. In the examples above, the speaker/writer could make her intended meaning much clearer by adding one of these words or phrases; for example,

Private schools generally attain better examination results than state schools.

On the other hand, someone using a **hard generalisation** does intend it to apply without exception. Such a generalisation is rightly conveyed by a quantifier such as '**all**', '**every**', '**no**', '**always**', '**never**'. For example

- ▶ <u>Every</u> passenger must hold a valid passport.

- ▶ <u>No</u> doctor who helps a patient to die should consider themselves to be above the law.

Why should we become critical thinkers?

If someone makes a claim that is intended as a hard generalisation and we can find a counter-example to it, then we have refuted their claim. But quantifier-free generalisations are not typically intended as hard generalisations. If the fanatical anti-generaliser does have a point, we believe, it is a point about rhetoric, not truth. What the anti-generaliser is justifiably worried about are generalisations about groups defined by race, ethnicity, nationality, gender, class and sexuality. Suppose that there are two social classes amongst Martians: the Zormons and the Ringons. And suppose the generalisation 'Ringons are more violent than Zormons' is true when taken as a soft generalisation, but false when taken as a hard generalisation. A Ringon anti-generaliser might object to someone's saying this. But we've seen that the point cannot be that the generalisation is not true, in spite of the fact that not every Ringon is more violent than every Zormon. This generalisation is, it must be admitted, true when taken as a soft generalisation.

However, it might be argued, it is rhetorically dangerous. There are two reasons: The first reason is that many people are not very clear about the possible ambiguities of such a statement. It might wrongly be taken as a hard generalisation, and furthermore it might wrongly be taken as asserting something about the innate or genetic qualities of Ringons. (For a similar case, refer back to the discussion of 'ability' in the section on ambiguity, earlier in this chapter.) In itself, it does not do this. So unless these possible misinterpretations are deflected by making the exact intended meaning perfectly explicit, this generalisation will remain very provocative and a likely cause of ill-feeling. The second reason is the brute fact that, even if these ambiguities are resolved, generalisations (even soft ones) about groups of people do often cause people to take offence. There are times when people take offence at a generalisation about a group and are simply irrational in doing so; no amount of explaining the difference between a soft generalisation and a hard one, or the difference between a generalisation about actual facts and one about alleged genetic qualities will change this. Like many kinds of irrationality, this is a natural kind of irrationality that cannot easily be overcome. No matter how factually true a generalisation may be, it is natural to feel that there is something dehumanising about it. So we cannot reasonably expect that people will always be able to overcome that feeling. Morality requires us to consider the consequences of our actions, and since speech and writing are types of action, natural (though irrational) responses to what we say and write must sometimes be taken into account in deciding what we ought to say. We should not say what is false, but that a proposition is true is not always enough to justify expressing it. This, we believe, is the grain of truth in the anti-generaliser's position.

Critical thinking enables us to ensure that we have **good reasons** to believe or do that which people attempt to persuade us to do or to believe. Attempts to persuade may be **argumentative or non-argumentative**. Most of the latter count as **rhetoric**, which is any attempt to persuade that does not attempt to give good reasons for the belief, desire or action in question, but attempts to motivate that belief, desire or action solely through the power of the words used. The former, on the other hand, persuade us by giving reasons for us to accept a claim or take the action suggested. Not all arguments are good arguments. Good arguments are those that provide us with **good reasons** to act or to accept a claim.

An argument consists of a **set of propositions**, one of which is its **conclusion** – the proposition argued for – the rest of which are its **premises** – the reasons given to accept the conclusion. Once we have determined that a text or a speech contains an argument, we must work out which sentence is intended to express the argument's conclusion and which are intended to express its premises. Words that serve as **conclusion indicators** and **premise indicators** offer a helpful (but not foolproof) guide to doing so successfully. We should also pay close attention to the context of the text or speech. Setting out arguments in **standard form** is a five-stage process that enables us to see the form of arguments better and hence, to compare, analyse and assess them more easily.

There are various **linguistic phenomena** that can make the task of identifying and interpreting arguments more difficult. In the case of **ambiguity**, **vagueness**, **metaphor**, **rhetorical questions** and **irony**, these can be problematic because they obscure speakers' and writers' intended meanings. In the case of **implicitly relative sentences** and sentences which use **quantifiers inappropriately**, they can be problematic because they fail to convey speakers' and writers' intended meanings in their entirety. Quantifying sentences can also cause problems for the interpretation of arguments when they are used inaccurately to express **generalisations**. There are two types of generalisation: hard and soft. **Hard generalisations** are true only if they are true without exception. To avoid misinterpretation they should be expressed in sentences that use quantifiers such as **all**, **every**, **no**, **none**, **always**, **never**. **Soft generalisations** are only true of the majority of the class that is the subject of the generalisation. They

Why should we become critical thinkers?

should be expressed in sentences that use quantifiers such as **most**, **almost all**, **in most cases**, **generally**, **typically**, **usually**. In all cases where linguistic phenomena prevent the intended meaning from being explicit, we should pay careful attention to context in order to render the most plausible interpretation of the attempt to persuade. Where appropriate we should **rewrite** sentences to make their meaning explicit.

EXERCISES **1** Decide whether each of the following cases contains an argument. If it does not, write 'N/A'. If it does, identify its premises and conclusion by underlining the appropriate propositions and writing 'C' under the conclusion and 'P' and the appropriate number under the premises.

Example
Bob is a dog and all dogs are black. So Bob is black.
         P1   P2   C

Notice that we have not underlined the words that connect or introduce the propositions, only the propositions themselves.

a   It follows from the fact that all cats are pests that this cat is a pest.

b   I'll never get to work if this traffic keeps up.

c   Whenever a person drinks instant coffee they end up with stomach ache and Jack is going to have stomach ache since he just drank a cup of instant coffee.

d   There is going to be a frost in the morning because the temperature has fallen below zero.

e   The biscuit tin is empty because the children ate all the biscuits.

f   We should not go to war against the Ringons because it would be unjust.

g   Since this animal is a fish, it can't be a mammal.

h   People really shouldn't drive so fast, I don't see what all the hurry is.

i   Leeds is north of Birmingham and Birmingham is north of Brighton. So Leeds is north of Brighton.

j   My ex-partner was always telling me to change my appearance, so I changed my partner.

k   If those chemicals are released into the river, thousands of fish will die.

l   Since inflation is increasing, the price of mortgages is sure to go up.

m   Everyone at the lecture is bored. No one who is bored is listening. Therefore no one at the lecture is listening.

n   He's been on crutches since he was injured in the accident.

o   On the basis of the fact that it includes scenes depicting drug abuse, the film should not be shown on prime-time television.

p   If we don't do something to control the level of car traffic now, air pollution will become so bad that our grandchildren will not be able to walk the streets for fear of asphyxiation.

q   The government proposes to reform the benefits system. Whenever such reforms occur someone loses out, so the government's proposals are unfair.

r   Something must be done to regulate the cultivation of genetically manipulated foodstuffs. Uncontrolled production of these crops will lead to a collapse of the ecosystem.

s   If we hit our children, they will learn that violence is acceptable, so we shouldn't physically discipline our children.

t   It is, therefore, an impractical solution to the problem of homelessness.

---

2   Write out the following arguments in standard form. You need not supply missing premises or change the words used unless it is absolutely necessary to retain the sense of a sentence.

Example:
The government should ban hunting. Hunting causes suffering to animals and anything that causes suffering to animals should be banned.

Why should we become critical thinkers?

P1) Hunting causes suffering to animals.
P2) Anything that causes suffering to animals should be banned.

_____

C)   The government should ban hunting.

a   If Leeds win against Arsenal, Manchester Utd will go to the top of the Premier League. Leeds have beaten Arsenal so Manchester Utd will be top of the league.

b   Children should not watch television programmes that lack educational merit. *Pokemon* fails to promote linguistic and cognitive development and programmes only have educational merit in so far as they promote linguistic and cognitive development. Children should not watch *Pokemon*.

c   I put it to you that Ms White killed Colonel Mustard in the ballroom with the candlestick. The reason I say this is that on the night of Colonel Mustard's death Lady Scarlet saw Ms White in the ballroom beating Colonel Mustard over the head with a candlestick, which was later found to have Ms White's fingerprints and Colonel Mustard's blood on it.

d   The team manager should be sacked. Whenever the team manager is sacked, team spirit is revitalised and this team's spirit certainly needs revitalising.

e   Excessive consumption by consumers in the developed world causes poverty and disease in the developing world and that's simply unjust. So if we care about the rest of the world, we should curb our consumption.

f   History will show President Clinton to have been a successful president after all. The reason is that he managed to maintain the strength of the US economy and that's the most important criterion by which to judge a president.

3 Without looking back at the relevant section, write a paragraph explaining the difference between lexical and syntactic ambiguity, then give a plausible example of each and explain their possible interpretations.

4   In the following sentences indicate the words or phrases that are lexically ambiguous and explain their possible meanings.

a   The last time I saw them they were on their way down to the bank.

b   Happiness is the end of life.

c   Archbishop of Canterbury praises organ donor for their humanity.

d   Museum visitor attacked by mummy.

e   Stolen car found by statue.

f   British left waffles on Ireland.

g   Iraqi head seeks arms.

h   An intense depression swept over the British Isles today.

i   Blair leans further to the right.

j   Chancellor wins on budget, but more lies ahead.

5   The following sentences are syntactically ambiguous. Rewrite them so as to give the most plausible interpretation. If two or more interpretations are equally plausible, give them all. You may need to rearrange the word order and/or add words.

Example:
A former professional dancer was accused of assaulting a 33-year old woman with her daughter.

A former professional dancer and her daughter were accused of assaulting a 33-year old woman.

a   The two suspects fled the area before the officers' arrival in a red Ford Escort driven by a woman in black.

b   I was invited to go to the movies yesterday.

c   Mary left her friends depressed.

d   People who use cocaine often die early.

e   Smith had a pair of boots and a pair of slippers that he borrowed from Jones.

f   That was John Major, the replacement for Thatcher, who was deposed for being too unpopular with the voters.

Why should we become critical thinkers?

g   Jones left the company in a better state.

h   Glasgow's first commercial sperm bank opened last Friday with semen samples from twenty men frozen in a stainless steel tank.

i   They were exposed to someone who was infected with the virus a week ago.

j   The police would like to speak to two women and a van driver who fled the scene of the accident.

6 Without looking back at the relevant sections, write a paragraph explaining the difference between vagueness and ambiguity. Give examples to illustrate your explanation.

7 For each of the following, identify the quantifier and say whether the generalisation is soft or hard.

Example:
Almost all students have contemplated cheating in an examination.
  Almost all
  Soft generalisation

a   No one may leave the room until the culprit owns up.

b   Few of the applicants are sufficiently qualified for the job.

c   Most Members of Parliament are committed to their constituents.

d   A majority of our members are prepared to go out on strike in support of their pay claim.

e   All passengers must fasten their seatbelts for take-off and landing.

f   Usually birds can fly.

g   Almost all of the patients are ready to be discharged.

h   Hardly any of the people surveyed were in favour of the proposed law change.

i   Every doctor must abide by the Hippocratic Oath.

j   Almost none of the candidates have the charisma to succeed in politics.

8 Each of the following sentences expresses a generalisation, but its quantifier is missing. For each sentence, if you think it is uncontroversially true as a hard generalisation, add an appropriate quantifier to make it a hard generalisation. If you think that it could only be true as a soft generalisation, add an appropriate quantifier to make it a soft generalisation.

Example:
Passengers must hold a valid ticket before boarding the train.

All passengers must hold a valid ticket before boarding the train.

a   Cats have tails.

b   Children like to eat ice cream.

c   Voters voted for Labour Party candidates at the last general election.

d   Owls are mammals.

e   Cars run on petrol or diesel.

f   Citizens of a democratic country should be free to come and go as they please.

g   Members of Parliament are male.

h   Universities in the UK have mathematics departments.

i   People care enough about the environment to change their lifestyles.

j   British people can speak a foreign language.

Chapter 2

Logic:
deductive validity

Our attempts to engage in critical thinking are sometimes frustrating. This is often because, even when we feel certain that there is something wrong with an argument, we find it hard to explain exactly what is wrong with it. Often this is frustration with ourselves; but it can easily look like frustration with the person giving the argument (it can certainly be interpreted as such by that person!) One of the primary aims of training in critical thinking is to learn concepts and techniques which will help us to express clearly what is wrong with an argument, thereby dispelling that frustration. By helping us to assess arguments more efficiently, this helps us in the pursuit of truth. But also, by becoming more articulate in our criticisms, we become less frustrated, and thereby less bad-tempered. This can help to smooth out our relationships with other people (whereas you might have thought that improving your skill at critical thinking would make you into a disagreeable, captious quibbler).

This frustration derives from two sources: First, confronted with an argument, we find it hard to hold the whole thing clearly before our mind's eye; we find it hard to say exactly what reasoning it is that we think must be mistaken. Second, even when we do succeed in laying the argument out before us clearly, we find it hard to describe or explain what is wrong with it.

- The first issue is addressed by techniques and strategies for **argument reconstruction**: the representation of arguments in standard form, so as to give us a clear and comprehensive view of them.
- The second issue is addressed by techniques and concepts of **argument assessment**: the determination of whether or not arguments provide good reasons for accepting their conclusions.

We discuss practical details of argument reconstruction in Chapter 5. This chapter and the next are mostly concerned with argument assessment. Ordinarily, we speak of arguments as being good or bad, strong or weak, valid or invalid, sound or unsound, persuasive or unpersuasive, intelligent or stupid, without having a clear idea of what we mean by these terms, and without clearly distinguishing their meanings. So not only are we vague when we use one of these terms to criticise the argument; our attempts to explain ourselves by means of the others are still vague. Thus our primary task in this chapter and the next is to explain the basic logical concepts in terms of which assessment is carried out – **validity**, **soundness** and **inductive force**.

You may be surprised that detailed discussion of argument assessment precedes the detailed discussion of argument reconstruction in Chapter 5; surely you have to reconstruct an argument before you can assess it? In fact, it is slightly less straightforward than that: although the final assessment of an argument must await its reconstruction, good reconstruction practice must be informed by a good grasp of the concepts used in assessment. The purpose of the next section is to explain this important point.

The Principle of Charity

An argument is a system of **propositions**: a set of premises advanced in support of a conclusion. People succeed in expressing the propositions they have in mind with varying degrees of clarity. In addition, an argument may depend upon premises that the arguer does not state at all, but which he or she is implicitly assuming.

For example, if someone argues: 'Sally is taking drugs; therefore she is breaking the law', the arguer is probably using the rather vague term 'drugs' in the narrow sense of 'recreational drugs', or perhaps in the sense of 'narcotics'. In the wider sense of 'drugs' that includes medicinal drugs, this would obviously be a bad argument. Furthermore, the arguer is assuming, without explicitly stating, that it is illegal to take such drugs. So two sorts of thing are left implicit in this argument: first, the arguer assumes a more precise meaning than is explicitly expressed by the word 'drugs'; second, the arguer fails to make explicit all the facts from which he or she infers the conclusion. A premise is left implicit.

Since the purpose of argument reconstruction is to determine exactly what argument has been given, it follows that part of the task of argument reconstruction is to **clarify** what the arguer actually said, and to **supplement** what the arguer actually said (to make explicit what was merely implicit in the arguer's statements). That is, we try to represent

the argument in such a way as to create a perfect match between the propositions upon which the argument actually depends and the sentences which represent the argument in standard form.

Two important consequences follow from this:

◆ The sentences we use in a reconstruction of the argument need not be the very same sentences used by the arguer in giving their argument. We may employ sentences which more clearly or precisely express the propositions that constitute the argument.
◆ Our reconstructed version of the argument may contain premises which are not expressed by any of the sentences actually used by the arguer.

Argument reconstruction is essentially a task of **interpretation**. What we are trying to reconstruct, to represent as clearly as we can, is a certain train of thought, of reasoning – however well or badly the arguer may have succeeded in expressing it. This cannot be an exact science. It cannot be mechanical or foolproof. It calls for judgement, a critical but sympathetic eye or ear, and even a certain degree of intuition, of understanding of people – of the ways people tend to think in given sets of circumstances, and of some typical ways in which people fail to express themselves clearly.

Nevertheless, the process can be undertaken in a systematic way, and there are general guidelines to follow. One of the most general of these is what we call the **Principle of Charity**, which we now explain.

We have just said that the reconstruction of arguments is often a task of surmising what the arguer had in mind, and was trying to express. Our main evidence for this, naturally, is the specific words actually used by the arguer. Beyond this, we look to various sorts of facts about the **context** or **circumstances** in which the person employed the words that he or she did. For example, consider this argument:

> But he is still in Paris! Therefore, he cannot possibly be in St Petersburg by tomorrow.

Of course nowadays St Petersburg is only a few hours from Paris by aeroplane. In the context of today, someone's being in Paris today would not prevent their being in St Petersburg tomorrow. If someone were to give a similar argument regarding some presently living person – the Russian president Vladimir Putin for example – then we would be puzzled. Since everyone is aware of jet travel, nobody thinks that it is impossible to get from Paris to St Petersburg in one day. So in the context of today, we would have to inquire further to discover what the arguer thinks is

preventing Mr Putin's journey. But suppose these words were given by someone in 1807, referring to Napoleon. Then surely the arguer would be assuming that it is not possible to get from Paris to St Petersburg at such a speed. That certainly would have been an appropriate assumption in Napoleon's day, so much so that it would have gone without saying. The fastest way to travel then was by horse.

Such facts pertaining to the context in which the argument is given, together with the specific words used by the person, will constitute the total evidence you have for reconstructing the argument. In some cases, the context is known, and makes it obvious what the arguer was implicitly assuming. In other cases, we may have to learn more about the context; this happens especially when interpreting historical documents.

In other cases, however, we may learn all the relevant contextual factors, yet it remains possible to represent the person's argument in more than one way. And it may happen that one reconstruction represents the argument as a good one, another as a bad one. In such a case, which reconstruction should you prefer? Which should you advance as *the* reconstruction of the argument?

It depends upon your purpose. If you are hoping to convince others that the person is wrong, you are most likely to succeed if you represent it as a bad one. Indeed, this is a very common ploy. If your aim is to defeat your opponent – or to make it seem as if you have defeated them – then success is more likely if you attack a weakened form of your opponent's argument. In a context like that of a public debate, this is often a good strategy to some extent. For what you are trying to do is to appear, to the audience, to get the upper hand. By representing your opponent's position as weaker than it really is, you are more likely to appear to be the victor. You also put your opponent on the defensive, forcing him to scramble, saying things like 'that's not what I meant'. If your aim is to persuade, or to appear to be the victor, then you may be well-advised to choose the weaker version, especially if your audience is not aware that a stronger version is available.

However, *if what matters to you is whether or not the conclusion of the person's argument is true, then you should choose the best representation of the argument.*

Why? Suppose you are wondering whether some particular proposition is true. You are wondering, for example, whether increasing taxes for the wealthy would lead to a rise in unemployment. Suppose further that you honestly have no idea whether or not this is true; you cannot see any reasons at all either in favour or against this proposition. Now suppose that someone attempts to persuade you that this proposition is true by giving you an argument for it. But you find that this argument admits of being reconstructed in either of two ways. On one reconstruction, the

argument is good, that is, it provides a good reason for accepting the proposition as true. On the other reconstruction, however, it is no good at all; you find that the reasons you have represented the person as giving in favour of this proposition do not support it at all. Suppose you decide on this latter representation of the argument, the one which represents it as bad. Can you now conclude that the proposition is true, or that it is false? Since, reconstructed that way, the argument was no good, you certainly cannot conclude, on the basis of it, that the proposition is true. But nor can you conclude that the proposition is false. The fact that someone has given a bad argument for some proposition is not, in itself, a reason to reject the proposition as false. For example, someone might argue that since three is a lucky number, there will not be a third world war. That is a bad argument; it gives you no reason to believe that there will not be a third world war. But (fortunately!) its being a bad argument provides no reason to believe that there will be a third world war. In short, the fact that someone has given a bad argument for the proposition in question leaves you in precisely the same position as you were when you started. If you began with no evidence either for or against the proposition, then your position is unchanged – you've no reason to accept the proposition as true, and none to reject it as false.

Suppose, then, that you accept the first reconstruction of the argument. Since this constitutes a good argument, you are now in a different position; now you do have some indication as to whether or not the proposition is true. In particular, you have a reason for its being true. On this first reconstruction, then, you represent the person as having made a useful contribution to the debate. You now have reasons that you lacked before. Thus, insofar as we engage in critical thinking – insofar as our interest is in discovering the truth of things, and not just in persuading or refuting people – we are most interested to discover good arguments, not bad ones. So we should always choose the best reconstruction of a given argument. That way, we discover reasons for accepting or rejecting particular propositions, advancing the cause of knowledge. This is an application of the **Principle of Charity**.

There is a further reason for observing the Principle of Charity, which has more to do with ethics than with logic. When you give an argument, you may or may not succeed in expressing yourself clearly, but you do want your listener to try to understand you. If your listener impatiently seizes upon your words in order to refute your argument as swiftly as possible without taking the trouble to understand you, naturally you feel ill-used, that the person is not being fair to you. You think it wrong, unjust to be treated that way. If so then we ought to try to be equally receptive to others – to try to understand them, rather than be too eager to refute them or discredit them. When people give arguments, they

almost always have some reason or other for what they are saying (although, of course, sometimes people do try to persuade us of things – especially to do things like buy Coke – without actually trying to give us good reasons). People are very seldom completely illogical. But they are seldom very well-practised at expressing their reasons clearly either, and often they are not so interested in clarity as in persuasion or eloquence. Still, beneath it all, they will usually have genuine reasons of some sort in mind, so it seems only right and proper that we should try to bring them to light, to understand what the person is really trying to say. If we do not attempt this, then we are not really doing the person justice; we are not being as receptive to his or her attempts at communication, as we would surely wish others to be to ours.

The Principle of Charity, however, has a certain limit, beyond which the nature of what we are doing changes somewhat: If our task is to reconstruct the argument actually intended by the person, then we must not go beyond what, based upon the evidence available to us, we may reasonably expect the arguer to have had in mind. Once we go beyond this point then we are no longer in the business of interpreting their argument. Instead, we have become the arguer.

If our concern is with how well a particular person has argued, then we should not overstep this boundary. *However, if our concern is simply with the truth of the matter in question, then to overstep this boundary is perfectly all right.* It often happens that, in reconstructing an argument, we hit upon another, similar or related argument for the same conclusion which is better than the one we are reconstructing. *If what concerns us is simply to find the best arguments on either side of an issue, then we will want to give a representation of this better argument.*

Truth

If your aim is to give the best possible reconstruction of an argument, then you have to know something of what makes an argument good or bad. Fortunately, logic gives us some very clear answers as to what does make arguments good and bad.

The fundamental concept of logic is the concept of **truth**.[1] For one thing, the overarching concern of the critical thinker is typically with the truth, or lack of it, of the conclusions of arguments. Further, truth is the

---

1   The great German logician Gottlob Frege – who is generally agreed to be the inventor of the modern science of logic – said that the Laws of Logic are really the 'Laws of Truth', in something like the way that the Laws of Physics are the laws of the physical world.

concept in terms of which the logician attempts to explain everything else. We thus begin our discussion of the concepts of logic by saying a little bit more about this uniquely important concept.

Many people are put off by the word 'truth'. This is usually the symptom of a philosophical worry that one cannot speak simply of 'truth'. One might worry that perhaps there is no one truth: that what is true for one person or group need not be true for another person or group. Or one may worry that truth is in some way beyond us, unapproachable by mere fallible human beings. But for our purposes, we can leave aside those sorts of abstruse philosophical worries as irrelevant. As noted in Chapter 1, the way in which the logician uses the word 'truth' is really very simple and down-to-earth. Properly understood, the word should not invite those sorts of controversies.

Consider the following proposition:

(A) Fish live in water.

This proposition is true. What does it mean to say that this proposition is true? It means, simply, that that is the way things are. To say that that proposition is true is to say nothing more than that, yes, fish *do* live in water. Thus consider the proposition that says that (A) is true:

(B) It is true that fish live in water.

(A) and (B) are equivalent in the sense that, necessarily, if (A) is true then so is (B), and if (B) is true then so is (A). In other words, to say that it is true that fish live in water comes to the same thing as saying that fish live in water. Used this way – which is all that is needed for logic or critical thinking – the word 'true' is no more mysterious than the words occurring in the sentence 'Fish live in water'. In this sense, you cannot doubt that there is 'really' such a thing as truth, or that truth is knowable, any more than you can doubt that fish live in water, or that the sky is blue, or that the Earth is bigger than a grapefruit. For these are all known truths.

Discomfort with the word 'true' is sometimes due to a failure to distinguish truth from belief. If John says 'Fish live in water', then he does, of course, show that he believes that fish live in water (presumably he knows that fish live in water). Likewise, if Mary now refers to what John said, and says 'That's true', then she also shows that she believes that fish live in water. Despite their having done so by different means, both John and Mary have asserted the proposition that fish live in water. Mary, unlike John, has used the word 'true'. But they have asserted the same proposition; they have expressed the same belief. Yet clearly the

truth of this proposition has nothing to do with what Mary or anyone else believes. That depends only on how things stand as regards fish, and what fish do does not depend upon what people think. So despite the fact that Mary has used the word 'true' to assert something, the truth of what she asserts does not depend on her beliefs in any way.

Of course, what Mary believes depends on her, and it is possible that people could have different beliefs as regards fish. But that has no effect on fish (we will, however, return to this issue in the final chapter).

The reverse side of this is that to say that a proposition is **false** is just to deny it – in this case, for example, some misinformed person who thought that snakes are fish might say, 'That's false; not all fish live in water'. 'It is false that fish live in water' is equivalent to 'Fish do not live in water'.

Sometimes we will speak of the **truth-value** of a proposition. This just means the truth of the proposition, if it is true, or its falsity, if it is false. There are two truth-values, true and false. For example we can say that the truth-value of 'Fish live in water' is truth, that that of 'Fish live in the sky' is falsity, and that the truth-value of 'It is now Tuesday' must always differ from that of 'It is now Friday'. To ask 'what is the truth-value of that proposition?' is the same as asking whether or not that proposition is true.

A question might have occurred to you: If to say that a proposition is true amounts to the same as asserting it, then why do we have the terms 'is true' and 'is false'? What is their purpose? Why are they not just redundant, superfluous appendages? One reason is convenience; saying 'that's true' is quick and easy, like saying 'yes', or nodding one's head. But a more important reason is that we sometimes want to generalise about propositions in terms of truth and falsity. That is, we sometimes wish to speak about true or false propositions in general, without specifying any propositions in particular. This is crucial in the formal study of logic, but less technical examples are no less important. For example, we have characterised critical thinking as aiming at truth. This means that we undertake it because we want to know whether capitalism is the fairest economic system, whether so-and-so committed the crime, whether the danger of war is increasing or decreasing . . . and so on, for everything we might want to know. That critical thinking aims at truth is a generalisation that sums this up.

Logic: deductive validity

Deductive validity

We also need this sort of generalisation in order to define the important concept of **deductive validity**, to which we now turn. For brevity, we will sometimes call it simply 'validity'. In studying it, you should forget whatever the word might mean to you ordinarily. We mean logical validity, the concept of validity that concerns logic, the study of reasoning.

Consider the following arguments:

A P1) The Prime Minister's dog is infested with fleas.
   P2) Fleas are bacteria.
   _____

   C) The Prime Minister's dog is infested with bacteria.

B P1) Colette owned a dog.
   P2) All French Bulldogs are dogs.
   _____

   C) Colette owned a French Bulldog.

Argument A speaks of 'The Prime Minister's dog', but it is not made clear who that is, for we are given no indication of when, or even in what country, this argument was given. So we have no idea what dog, if any, has been referred to. Furthermore, P2 of A would be false under any circumstances in which the argument might have been given – fleas are insects, not bacteria. But scrutinise these arguments carefully. You can easily recognise that there is something right about A, and something wrong with B. The conclusion of A does follow from its premises, and the conclusion of B does not follow from its premises. What you are recognising is that A is **valid**, and that B is **invalid**.

Now what does this mean? What, exactly, are you seeing when you see that A is valid and B is invalid? Consider A. When you recognise its validity, you do not need to know whether or not P1 or C of A is true or even precisely which dog of which Prime Minister of which country is intended; nor do you care about the fact that P2 of A is positively false. For what you are seeing is that if the premises of A _were_ true, then the conclusion would be have to be true as well. In short, _it would be impossible for the premises to be true but the conclusion false_. The truth of the premises, in any possible or imaginable situation, would necessitate the truth of the conclusion. If fleas were bacteria, and the dog

being referred to were infested with fleas, then it would, in that case, be infested with bacteria.

On the other hand, consider B. When you recognise that it is not a valid argument, what you are recognising is that even if the premises were true, it would still be *possible* for the conclusion to be false. The conclusion does not follow. Whether or not the premises are *in fact* true, it would be possible or conceivable for the premises to be true and the conclusion false.

Now as it happens, the premises of B are true. The French author Colette did have a dog, and of course French Bulldogs are dogs. Indeed, the conclusion of B is also true; Colette's dog was a French Bulldog. But that is beside the point, so far as validity is concerned. It may be true that Colette had a French Bulldog, but this does not follow merely from the fact that she had a dog (along with the fact that French Bulldogs are dogs). A person given only the premises of the argument, and lacking any further information about Colette, would be in no position to infer that Colette had a French Bulldog. If this point seems strange, remember that you saw that the argument is invalid *before you knew the truth-values of its premises or of its conclusion*. And that is as it should be. You can tell that an argument is valid or not without knowing the truth-values of the propositions it comprises, because the validity of an argument (or lack thereof), does not depend upon the actual truth-values of those propositions.

This shows you that the concept of validity pertains to the *connection* between the premises and conclusion of an argument, not their truth-values considered individually. This is indeed the crucial lesson about the concept of validity: it pertains to *whole arguments* (more exactly: it pertains to *inferences*; extended arguments may contain more than one inference, and each one is subject to being valid or invalid).

Thus, it should be clear that it would be nonsense, to say of a single proposition, that it is valid. That would be like saying, of a single word, that it rhymes (a rhyme requires a relation between words). By the same token, it would be nonsense to say, of an argument, that it is true. That would be like saying, of an entire jigsaw-puzzle, that it doesn't fit (this could be said of some or even of every piece, but not of the puzzle itself). A single proposition can be true or false, but not valid or invalid; an argument can be valid or invalid, but not true or false.

These two points should be borne in mind, as it is a common mistake to confuse the notions of truth and validity, applying them to the wrong sorts of things.

Here then are two definitions of **validity**; they are equivalent (they come to the same thing), so you are free to make use of the one you find easier to work with:

Logic: deductive validity

> **To say that an argument is *valid* is to say: It would be impossible for all the premises of the argument to be true, but the conclusion false.[2]**

And the second one:

> **To say that an argument is *valid* is to say: If the premises are (or were) true, the conclusion would also have to be true.**

If the condition specified by the definition does not hold, then the argument is **invalid**. A consequence of these definitions is that the following cases of valid arguments are all possible:

1   The premises are all (actually) true, and the conclusion is (actually) true.

2   The premises are all (actually) false, and the conclusion is (actually) false.

3   The premises are all (actually) false, and the conclusion is (actually) true.

4   Some of the premises are (actually) true, some (actually) false, and the conclusion is (actually) true.

5   Some of the premises are (actually) true, some (actually) false, and the conclusion is (actually) false.

The only case in which an argument *cannot* be valid is the case when the premises are all (actually) true, but the conclusion is (actually) false. For if that is so, then obviously there *is* a

---

2   This definition has the consequence that if any premise of an argument is a necessary falsehood, or if the conclusion is a necessary truth, then the argument is valid (a necessary falsehood is a proposition that could not possibly be true; a necessary truth is a proposition that could not possibly be false). In such cases the premises may be entirely irrelevant to the conclusion. For example, 'There is a married bachelor, therefore the moon is made of green cheese' is valid, as is 'The moon is made of green cheese, therefore there is no married bachelor'. Our definition, then, is quite useless as a guide to reasoning, where necessary truths and necessary falsehoods are concerned. We believe this a reasonable price to pay, for the alternative – a definition of validity whereby an argument is valid by virtue of its form – is too difficult, for our purposes, to apply profitably to ordinary language. Further, it is really very seldom that necessary truths or falsehoods figure as the conclusions or premises of arguments encountered ordinarily.

possible case in which the premises hold true when the
conclusion is false – the actual case.

This is easier to grasp by looking at some examples. The following
are, respectively, examples of cases 1–5 given above; the 'T's and 'F's in
parentheses to the right of each premise or conclusion indicates its actual
truth or falsity, as the case may be:

1 P1) Janet Baker is an opera singer.   (T)
    P2) All opera singers are musicians.   (T)

    C)   Janet Baker is a musician.   (T)

2 P1) Janet Baker is a baritone.   (F)
    P2) All baritones are Italians.   (F)

    C)   Janet Baker is an Italian.   (F)

3 P1) Janet Baker is a baritone.   (F)
    P2) All baritones are English.   (F)

    C)   Janet Baker is English.   (T)

4 P1) Janet Baker is a soprano.   (T)
    P2) All sopranos are English.   (F)

    C)   Janet Baker is English.   (T)

5 P1) Janet Baker is a soprano.   (T)
    P2) All sopranos are Italians.   (F)

    C)   Janet Baker is an Italian.   (F)

Logic: deductive validity

And here are some *invalid* arguments:

6	P1)	Janet Baker is a soprano.	(T)
	P2)	Janet Baker is a musician.	(T)
	C)	Janet Baker is an Italian.	(F)

7	P1)	Janet Baker is a woman.	(T)
	P2)	All baritones are women.	(F)
	C)	Janet Baker is a baritone.	(F)

8	P1)	Janet Baker is a singer.	(T)
	P2)	All sopranos are singers.	(T)
	C)	Janet Baker is a soprano.	(T)

9	P1)	Janet Baker is a baritone.	(F)
	P2)	All singers are baritones.	(F)
	C)	Janet Baker is a singer.	(T)

How to test for validity

The way to determine whether or not an argument is valid, then, is really to ignore the actual truth-values of the premises and the actual truth-value of the conclusion (of course, if you know that the conclusion is actually false, yet the premises are all true, then the argument must be invalid. But normally when assessing arguments we do not know the truth-value of the conclusion, because the whole reason for considering the argument is that we want to find out the truth-value of the conclusion). The way to do it – and this is what the definition of validity should lead you to expect – is to reason as follows:

> **Whether or not they are actually true, suppose or pretend
> that the premises were all true; then in that situation –
> aside from how things really are – could the conclusion
> conceivably be false? If it could not be, then the argument
> is *valid*. If it could, then the argument is *invalid*.**

The systematic study of validity is the concern of **logic**. Logicians are concerned to devise perfectly reliable procedures for detecting validity, or the lack of it, even in the case of extremely complex arguments such as those occurring in mathematical proofs. Since the validity of an argument is independent of the truth-values of its premises, logic has a unique status amongst the sciences; for other sciences are concerned to find out the truth-values of particular propositions about its characteristic subject matter. Ichthyology, for example, seeks to know which propositions about fish are true, and which false. The logician has no particular concern with fish, nor with the truth as regards anything else in particular. Logic has no concern with particular truths. In a sense, the logician does not have to know anything. The logician is concerned only with relations between propositions, not with their truth-values. These are the sorts of relations displayed between premises and conclusion in valid arguments such as 1–5. There are sophisticated technical procedures involving special symbols for exploring these relationships in a systematic, detailed way. Such is the concern of what is known as 'formal logic', 'mathematical logic' or 'symbolic logic'. Courses on the subject are taught in philosophy, mathematics, linguistics and computer science.

In critical thinking we are doing what you might call practical logic. We want to learn to identify the reasoning in commonly encountered attempts to persuade us, and to assess it as good or bad. For this, we need the concept of validity, but we do not need elaborate technical procedures for detecting validity. The reason is that the logic of the vast majority of arguments in everyday life is rarely of any great complexity. Once we know exactly what the argument is, whether or not it is valid can, almost always, simply be seen by applying the definition given above. Most of the work goes into the reconstruction. And as we have seen – and will see in more detail later – we cannot profitably reconstruct an argument without knowing what makes a good argument, hence not without grasping the concept of validity.

Logic: deductive validity

Further examples

The examples so far of valid arguments have all been of the same type or form. Here are some further, still quite simple types of valid arguments. The examples are all fictional. Their being fictional helps us to realise that the actual truth-values of premises and conclusions are usually irrelevant to determining whether or not arguments are valid.

P1) No Zormons are ticklish.
P2) Trozak is a Zormon.

C)   Trozak is not ticklish.

P1) Either Trozak is on Mars, or he is on Jupiter.
P2) Trozak is not on Mars.

C)   Trozak is on Jupiter.

P1) If Trozak is on Mars, he will visit Ichnik.
P2) Trozak is on Mars.

C)   Trozak will visit Ichnik.

P1) If Trozak ate all the biscuits, then the biscuit tin is empty.
P2) The biscuit tin is not empty.

C)   Trozak did not eat all the biscuits.

Conditional propositions

This is a good point at which to distinguish arguments from a certain kind of statement, which both logicians and grammarians call the **conditional**. Conditionals are most characteristically expressed using the 'if–then' form of declarative sentence. For example:

If it is raining, then the barometric pressure is low.

Another example is P1 of the last example of the preceding section of this chapter. Conditionals can also be expressed in other ways, however. For example, the following statements express the very same proposition as the example just given; they represent exactly the same connection between rain and barometric pressure:

▶ It is raining only if the barometric pressure is low.

▶ Either the barometric pressure is low, or it is not raining.

▶ It is not raining unless the barometric pressure is low.

▶ If the barometric pressure is not low, then it is not raining.

▶ It is not raining if the barometric pressure is not low.

▶ The barometric pressure is low if it is raining.

▶ There cannot be rain without low barometric pressure.

You may find the first one puzzling, because in it the word 'if' precedes the bit about the barometric pressure rather than the bit about the rain. But with a little reflection it is not difficult to see that it means the same as 'If it is raining then the barometric pressure is low'. In general, a conditional is a compound proposition consisting of two parts, each of which is itself a proposition, where these two parts are joined by some connecting words (they are called 'logical connectives') such as 'if–then', 'either–or', 'unless', or 'only–if', or something similar. Sometimes the presence of two whole propositions is somewhat concealed, as in the last example.

Especially notable is the second example, using the words 'either–or'. Wherever two statements are joined by 'either–or', or just by 'or', to form a compound statement, the compound is equivalent to a statement using 'if–then', and vice versa. But to pass from the version using 'or' to the version using 'if–then', or vice versa, we have to insert the word 'not'. For example, the following pairs of statements are equivalent:[3]

---

3   A complication is that the word 'or' is used in either of two ways, known as the 'inclusive' and 'exclusive' senses. In the inclusive sense, 'P or Q' means that either P is true, or Q is true, or they are both true; so the compound sentence is false only if both P and Q are false. This is the sense intended if one says something like 'United will lose if either Beckham receives a red card or Keane receives a red card'. In the exclusive sense, 'P or Q' means that either P is true or Q is true, but not both. This is the sense normally intended if one says something like 'Either you go to bed now or I will not read you a story' (the child would rightly feel cheated if he or she goes to bed immediately but doesn't get a story). A conditional statement 'If P then Q' is equivalent only to 'Either not-P or Q' only in the inclusive sense of 'or'. Where it is clear that the exclusive sense is intended, the reconstructed argument should contain conditionals running in both directions. A sentence such as 'Either the murder took place here or there was an accident here', intended in the exclusive sense, would be represented as 'If no accident took place here then the murder took place here' and 'If an accident took place here then the murder did not take place here'.

Logic: deductive validity

> Rangers will win or Celtic will win.
>
> If Rangers do not win then Celtic will win.
>
> Either Jane will practise diligently or she will fail her exam.
>
> If Jane does not practise diligently then she will fail her exam.

However they are joined, what a conditional says is that if one proposition is true then another one is true also. Usually in logic this relation is represented by a single device, an arrow:

> It is raining → The barometric pressure is low.

The one *from* which the arrow points is called the **antecedent**; one *to* which the arrow points is called the **consequent**, for obvious reasons. We will use this terminology a lot, so you should memorise it.

Now, here is tricky and important point. In one way or another, the examples above all express the very same conditional proposition. They all express the same relation between rain and barometric pressure. Thus the antecedent and consequent in all the above examples are the same – in all of them, the logical antecedent is the proposition that it is raining, and the logical consequent is the proposition that the barometric pressure is low. In this sense, the fact that one proposition the antecedent and another the consequent of a conditional statement is a matter of the *logic* of the statements. It is not a matter of the grammar of the sentences. It does not matter in what order, in the whole sentence, the two smaller sentences occur; what matters is the logical relationship asserted by the sentence. You can see this from the fact that in the sentences, the bit about rain sometimes occurs before, sometimes after the bit about barometric pressure.

Our example of a conditional is a true statement. Some conditionals are false, such as 'If the barometric pressure is low, then it is raining' (since, sometimes, the barometric pressure is not low but it is not raining). A conditional is said to be true or false, rather than valid or invalid. For a conditional is not itself an argument. A conditional is one proposition that comprises two propositions as parts, joined by 'if–then' or a similar device. An argument cannot be just one proposition. It needs at least two.

The following, however, would be an argument:

> It is raining. Therefore, the barometric pressure is low.

This is not a conditional, but an argument composed of two propositions. Moreover, this argument actually asserts that it is raining, and that the

barometric pressure is low. A person giving it would actually be asserting those things. Not so for the corresponding conditional: To say 'If it is raining then the barometric pressure is low' is not itself to assert that it is raining, or that the barometric pressure is low. People sometimes make a mistake on this point; we sometimes witness conversations like this:

> Mary: If Edna gets drunk, then her graduation party will be a mess.
>
> Jane: Why do you say Edna's going to get drunk? You're always so unfair to her.
>
> Mary: I didn't say Edna is going to get drunk, I said *if* Edna gets drunk . . .
>
> Jane: Well, you implied it.
>
> Mary: No I didn't!
>
> Jane: Don't try to get out of it! You always think the worst of Edna, and you're always trying to take back what you say!

What Jane seems not to understand is that a conditional does not assert either its antecedent or its consequent. An argument asserts its premises and its conclusion, but Mary is not arguing that the party is going to be a mess; she is only saying that it will be if Edna gets drunk (perhaps because she knows that Edna doesn't handle alcohol very well). Strictly speaking, she is not even suggesting that Edna is likely to get drunk (of course, she would not have brought the whole thing up if she did not think there was some danger of Edna's getting drunk).

Finally, it is important to recognise that many arguments have conditional conclusions – that is, conclusions which are themselves conditionals. For example:

> P1) If Labour does not change its platform, it will not attract new supporters.
>
> P2) If Labour does not attract new supporters, it will lose the next election.
>
> ———————————————————
>
> C) If Labour does not change its platform, then it will lose the next election.

Here both premises as well as the conclusion are conditionals. This particular pattern is very common, and is called a 'chain' argument. A chain of conditionals is set up, like a row of dominoes. What the argument is saying is that if the antecedent of P1 comes true, then the consequent of P2 will come true. Chain arguments can have any number of links.

Logic: deductive validity

Deductive soundness

Normally, if you wish to assess an argument, you wish to do so because you wonder whether or not the conclusion is true. You want to know whether the arguer has given you a reason for thinking that the conclusion is true. Now if you find that the argument is invalid, then you know that the conclusion could be false, even if the premises are true. Therefore the reasons given by the arguer – the premises – would not suffice to establish the conclusion, even if they are true. But suppose you find that the argument is valid. Then there are two possibilities:

A   One or more of the premises are (actually) false.

B   All of the premises are (actually) true.

Now, as illustrated by examples 2 and 5 on p. 53 knowing that the argument is valid is not enough to show you that the conclusion is true. In order to determine that, you need a further step: you must determine the truth-values of the premises. You might already know them. But if you don't, then of course, logic is no help. If one of the premises is that the Octopus is a fish, then unless you know already, you have to consult a book or ask an ichthyologist. Suppose now that you have done this, and what you have is a case of (A), i.e. one (or more) of the premises is false. In that case, you can draw no conclusion as to the truth-value of the conclusion (as illustrated by arguments 4 and 5 on p. 53, a valid argument with one or more false premises may have either a true or a false conclusion). But now suppose you find the argument to be a case of 1 – that is, you have found it to be a valid argument with true premises. Eureka! For now, according to the definition of validity, you may infer reliably that the conclusion of the argument is true. The argument has accomplished its purpose; it has demonstrated its conclusion to be true. We call this a **deductively sound** argument. Argument 1 above about Janet Baker, for example, is a deductively sound argument.

> **To say that an argument is *deductively sound* is to say:**
> **The argument is valid, and all its premises are (actually) true.**

This reveals the importance of the concept of validity. Given the definition of validity, it follows from the definition that the conclusion of a deductively sound argument must be true. There cannot be a deductively sound argument with a false conclusion.

An argument which is not deductively sound – which has one or more false premises, or is invalid, or both – is said to be **deductively unsound**. Deductive soundness, like validity, pertains to whole arguments, and not to single propositions.

It is also important to recognise what follows if you happen to know that the conclusion of an argument is not true. Suppose someone gives you an argument, the conclusion of which is that there are platypuses at the local zoo. And suppose you know that this is not true. You know that the local zoo has no platypuses. Therefore you know that this argument is not deductively sound. You should make it clear to yourself why this is so; check the definitions again if this is not clear.

If you really do know – and do not merely have the opinion – that there are no platypuses at the zoo, then you know it is possible to give a deductively sound argument for that conclusion. *But there cannot be deductively sound arguments on both sides of the issue*. For deductively sound arguments have true conclusions. If there were deductively sound arguments on both sides of this issue, it would follow that there are platypuses at the zoo, and also that there are not, which is impossible. This is important to recognise, because frequently we do say that there can be 'good' arguments on both sides of a given issue (especially a controversial one); we say this, perhaps, out of a wish to show respect for different opinions, or simply to express our own indecision over the issue. But in saying this, we cannot mean that there are *deductively sound* arguments on both sides of an issue. Later, we will explain in exactly what sense there can be 'good' arguments on both sides of an issue (for, to be sure, there can be).

If we know that the conclusion of an argument is false, then we know that the argument is deductively unsound. What follows from that? Look at the definition of deductive soundness. If the argument is deductively unsound, it follows that *either* the argument has (at least) one false premise, *or* the argument is invalid (or perhaps both – perhaps it is invalid and it has one or more false premises). Suppose then that you determine the argument to be valid. Then you know that at least one premise must be false. On the other hand, suppose that you find that the argument is invalid. What can you conclude about the truth-values of the premises? Nothing! For you know that an invalid argument may have either true premises or false premises.

Similarly, if you perhaps do not *know* (for you are not certain), but you do believe, or hope, that the conclusion of a given argument is false, then you must run through the same procedure. Suppose you merely hope that there are no platypuses at the zoo, because you fear they would not be happy there (this might be reasonable; platypuses are notoriously shy, so they might not like to be put on display). Then you must hope

that the argument is deductively unsound, in which case you must hope to find either that the argument is invalid, or that it has a false premise. This is the sort of thing you might do if you were a courtroom barrister, hoping to refute the opposing side's arguments. You might want to show that the prosecution's argument for your client's guilt is unsound.

CHAPTER SUMMARY

The aim of argument **reconstruction** is to **clarify**, and make fully **explicit**, the argument intended by an arguer. We do this by putting the argument into **standard form**. If our main concern is whether or not the conclusion of the argument is **true**, then our reconstruction should be guided by the **Principle of Charity**: we should aim for the best possible reconstruction of the argument.

In order to do this, we need precise concepts in terms of which to assess arguments. We need, to begin with, the concepts of **truth**, **deductive validity**, and **deductive soundness**. A deductively valid argument (a **valid** argument for short) is one whose premises could not possibly be true without the conclusion being true also. A deductively sound argument is a valid argument with true premises. It follows that a deductively sound argument must have a true conclusion.

Individual propositions can be true or false, but not valid or invalid. Arguments can be valid or invalid, but not true or false. **Conditional propositions** are not arguments, but single propositions: a conditional proposition is a single proposition made up of two propositions, the **antecedent** and **consequent**. Typically the antecedent and consequent of a conditional proposition are joined by 'if–then', but many other devices can do the same job. The crucial thing is the logical relationship between them.

Arguments do stand in a certain relationship to conditional propositions. If the argument is valid, then this conditional proposition is true: **If** the argument's premises are true, **then** its conclusion is true.

1   Study the section that describes the Principle of Charity. Then, without looking at the section again, and in your own words, write a short essay of about 250 words that explains the principle, and why it should be observed.

2   Suppose that someone says, 'The purpose of argument is to defeat the opponent'. Write a short essay (about 250 words) commenting on this.

3   Again, without consulting the book, and in your own words, explain the purpose of argument reconstruction.

4   Suppose that Mr Smith argues as follows:

Mr Jones argues that the unemployment rate will rise this year. However, as we have explained, Mr Jones' argument is clearly invalid. Furthermore, we have shown that the premises of the argument are false. Therefore Mr Jones is wrong. The unemployment rate will not rise this year.

Criticise Mr Smith's argument.

5   Suppose that Bob says:

Sarah is so arrogant. She's entitled to her opinions, of course. I don't mind her *saying* that landlords have an ancient right to prevent walkers' access to their property. That's her opinion. But she has to say it's *true*. I hate that. She always thinks she's right. She always thinks *her* opinion is *the* truth.

Does Bob have a point, or is he confused? Explain your answer in an essay of about 150 words.

Logic: deductive validity

6 For each argument, decide whether or not it is deductively valid. If it is not valid, briefly describe a possible situation in which the premises are true, but the conclusion false.

Example:

P1) Every Roman emperor before Constantine was a pagan.
P2) Julian was a Roman emperor.

C) Julian was a pagan.

Invalid. P2 says that Julian was a Roman emperor, but not that he ruled before Constantine. It would be possible for the premises to be true but the conclusion false, then, if Julian were a non-pagan Roman emperor after Constantine.

a P1) Constantine ruled Rome before Constantius.
   P2) Constantius ruled Rome before Julian.

   C) Constantine ruled Rome before Julian.

b P1) The only just political systems are democracies.
   P2) Rome's political system was unjust.

   C) Rome's political system was not a democracy.

c P1) Unless some historians have told lies, there were miracles during the first century.
   P2) Some historians have told lies.

   C) There were no miracles during the first century.

d   P1) The accident on Sunday at Cleever's junction was caused by the driver coming too fast down the hill.

       C) Accidents at Cleever's junction are always caused by drivers coming too fast down the hill.

e   P1) If Constantius' professed Christianity was genuine, then history has represented him unfairly.
      P2) We can never know that Constantius' professed Christianity was *not* genuine.

       C) History has represented Constantius unfairly.

f   P1) If Germanicus had not died at the German front, he would have become emperor.
      P2) If Germanicus had become emperor, then Tiberius would not have become emperor.
      P3) If Tiberius had not become emperor, then Caligula would not have become emperor.
      P4) If Caligula had not become emperor, then he would not have been murdered.
      P5) If Caligula had not been murdered, then Claudius would not have become emperor.

       C) If Germanicus had not died at the German front, then Claudius would not have become emperor.

g   P1) If any Roman emperor was wise, then Marcus Aurelius was wise.
      P2) Augustus was a Roman emperor.
      P3) If Augustus was wise, then Marcus Aurelius was not wise.

       C) Marcus Aurelius was wise.

Logic: deductive validity

h   P1) If John and Susan are late, then Mary will be
                disappointed.
      P2) John is late.
      P3) Susan is late.

      C)   Mary will be disappointed.

i   P1) If John and Susan are late, then Mary will be
                disappointed.
      P2) John is late.
      P3) Susan is not late.

      C)   Mary will not be disappointed.

j   P1) If John and Susan are married, then Mary will be
                disappointed.
      P2) John is married.
      P3) Susan is married.

      C)   Mary will be disappointed.

k   P1) If the Prime Minister does what the opinion polls say
                the people want, then he is cowardly.
      P2) If the Prime Minister does not do what the opinion
                polls say the people want, then he is arrogant.

      C)   If the Prime Minister is not arrogant, then he is
                cowardly.

7 Which of the following statements are true, and which false?

a It is impossible for the premises of a valid argument to be false.

b It is possible for the conclusion of a valid argument to be
false.

c   It is possible for a deductively sound argument to have a false conclusion.

d   A valid argument cannot have true premises and a false conclusion.

e   A valid argument cannot have false premises and a true conclusion.

f   It is impossible for the premises of a valid argument to be true, if the conclusion is false.

g   A valid argument must have true premises.

h   If the conclusion of a valid argument is false, then all the premises must be false.

i   If an argument has more than one premise, then if one of the premises is true, then the others must also be true.

j   If a valid argument is sound, then one of its premises must be false.

k   A sound argument cannot have a false premise.

l   An argument cannot be both valid and sound.

m   If an argument has true premises and a true conclusion, then it is valid.

n   If an argument does not have true premises and a false conclusion, then it is valid.

o   An argument may have true premises and a false conclusion.

---

8   In exercise 7, a–f, various scenarios are said to be possible or impossible. In cases where the thing is possible, give an example.

---

9   Rewrite the following sentences using 'if–then', so as to express the same relation between propositions. Sometimes you may have to insert or remove a 'not'.

a   Either Trajan was great or Hadrian was great.

b   Augustus Caesar was the greatest emperor, unless Marcus Aurelius was the greatest emperor.

c   Galerius and Maximin were not both admired in Rome.

Logic: deductive validity

10 Can there be more than one deductively sound argument for the same conclusion? If so, give an example. If not, explain why not.

11 Consider the following two sentences:

Lee Harvey Oswald murdered John Kennedy.

Lee Harvey Oswald did not murder John Kennedy.

Suppose we have two arguments, one with the first as conclusion and the other with second as conclusion. Could both arguments be valid? Could both arguments be deductively sound? Why or why not?

Chapter 3

Logic:
inductive force

Inductive force

Consider this argument:

P1)   Fiona lives in Inverness.

_____

C)   Fiona owns at least one item of woollen clothing.

Is this argument valid? One might think so. Inverness, in northern Scotland, is a pretty cold place. And there are a great many sheep, hence a great deal of wool, in the vicinity of Inverness. So surely one may conclude that, if Fiona does live in Inverness, then she's got at least, say, a woollen jumper or two. Nevertheless, the argument is clearly invalid: it would not be impossible for someone to live in Inverness, yet have no woollen items of clothing. Indeed, it is probable that there really are some people there without any (some people are allergic to wool, for example).

Now we could make the argument valid by adding the premise, 'Everyone in Inverness owns at least one woollen item of clothing'. Thus:

P1)   Fiona lives in Inverness.
P2)   Everyone in Inverness owns at least one woollen item of clothing.

_____

C)   Fiona owns at least one woollen item of clothing.

But this argument, though valid, is probably not sound, since P2 is probably false. Yet the original argument is surely a good argument, in some sense. The truth of the premise would be a good reason for expecting the conclusion to be true; one would be surprised to find it was false. Certainly if you had to bet on whether or not the conclusion is true, then, given no relevant information except for P1, you would bet that it is true, not that it was false. You would reasonably infer the truth of C from P1 rather than its falsity.

However, we must notice something very important about the inference from P1 to C. A logician would not be in a position to recognise the forcefulness of this argument as we have written it so far. For in order to recognise it, one has to know certain facts about Inverness – the facts which make it different from, say, a hot place around which they grow cotton like Cairo. The logician, as noted earlier, does not know anything as specific as that. Someone giving the argument is implicitly relying on a proposition which is similar to P2 in the argument given directly above, but which is much more likely to be true – namely, that very few people in Inverness own no items of woollen clothing, which is the same as saying that *almost everyone* in Inverness owns at least one item of woollen clothing. That is to say, the use of the quantifier 'everyone' would be inappropriate here, but use of the weaker quantifier 'almost everyone' is appropriate. Given what we know about Inverness, we can be almost completely certain that almost everyone in Inverness owns at least one woollen item of clothing. (Recall from Chapter 1 that expressions such as 'most', 'almost all', and 'few' are called 'quantifiers'.) We can represent the argument, then, as follows:

P1) Fiona lives in Inverness.
P2) Almost everyone in Inverness owns at least one item of woollen clothing.

C) Fiona owns at least one woollen item of clothing.

This argument is still not valid, since Fiona still might conceivably be one of those few who do not have any woollen clothing. Still, these premises, just by themselves, do provide a good reason for accepting the conclusion: The logician, knowing nothing but the truth of P1 and P2, could happily accept that if the premises are true, then, probably, so is C. We recognise this by calling such an argument **inductively forceful** (the word 'inductively' is meant to contrast with 'deductively'), and inserting the word 'probably', in parentheses, before the conclusion:

P1) Fiona lives in Inverness.

P2) Almost everyone in Inverness owns at least one item of woollen clothing.

C) (Probably) Fiona owns at least one woollen item of clothing.

The word 'Probably', here, is not to be regarded as part of C. It is not, strictly speaking, part of the argument. What it is, rather, is an indication to the reader that the argument has been judged, by the person doing the reconstruction, to be inductively forceful. If we were to remove it from the above reconstruction, the inductive force of the argument, and its status as not being deductively valid, would be unaffected.

Roughly, then, an inductively forceful argument is one that is not deductively valid – the truth of the premises would not ensure the truth of the conclusion – but whose premises provide good reason to expect the conclusion to be true rather than false. Before we characterise the concept of inductive force more accurately, however, it will be useful to look briefly at the more basic concept of **probability**.

Probability

We express the probability that a given proposition is true (or that a given event has occurred or will occur) on a numerical scale between 0 and 1, expressed either as a decimal or as a fraction. For example the probability that a tossed coin will land heads-up is 0.5 or 1/2. Perhaps surprisingly, there are different ways in which to explain probability. We'll briefly consider three: **proportion**, **frequency**, and **rational expectation**.[1]

First, **proportion**. Many arguments contain a premise that says something like 'Most X are Y' or '7/8 of Xs are Ys'. Such quantifiers as 'most' and '7/8 of' indicate proportions, and are importantly related to probability. Suppose you want to know the probability that the card you have drawn from an ordinary, complete deck of playing cards is an ace. One way to do this would be to assume that this probability is equal to the proportion of aces in the deck to the total number of cards in the deck.

---

1   Another common basis for probability-estimates is the use of **models**. For example, if an engineer wants to estimate the probability that a lorry will tip over in high winds, he or she would consider such things as the exact design of the lorry and the winds it would be likely to encounter in order to deduce a probability using the laws of physics.

Since this figure is 1/13 (there are four aces and fifty-two cards, and 4/52 = 1/13), you assume that the probability that you'll draw an ace is 1/13 (about 0.077).

Now **frequency**. Suppose you want to know the probability that it is going to snow in December in London, when December is still several months away. One simple way to do this would be to find out how frequently this has actually happened over the past, say, one hundred years. Suppose you find that, out of the past one hundred Decembers, it has snowed during fourteen of them. Then you might infer that the probability that it will snow in London this coming December is 14/100 (0.14).

These strategies are extremely important in the general theory of probability and statistics, but they are not sufficiently general for our purposes. There seem to be cases for which neither proportion nor frequency will serve as a direct indicator of probability. For example, bookmakers sometimes give odds that a given politician will become the next leader of his or her party. Suppose they say that the odds of Mr X becoming the next leader of the Labour Party are 1:1 (i.e. the probability is 1/2). Estimates such as these are often perfectly reasonable. But clearly the bookmaker, in this case, is not basing the probability on the frequency with which Mr X has become leader of the Labour Party in the past; that frequency is zero! Nor is it based upon the proportion of times that the politician has become leader. Probability-estimates of this kind are sometimes quite reasonable, but there is no immediate and simple recipe for basing them upon frequencies or proportions.[2] In such cases, we cannot simply convert a proportion or a frequency into a probability.

Because of these complications, we shall take **degree of rational expectation** as our general concept of probability.[3] A person's degree of rational expectation in a given proposition is the degree to which he or she is entitled to believe it, given the evidence he or she has. Besides the fact that this concept of probability is more widely applicable than others, it has the further advantage that it corresponds to the way that the word 'probably' is typically used: when we say, 'That's probably true', we typically mean that our total sum of evidence makes it reasonable to expect the proposition in question to be true.

---

2   This is not to say that such estimates cannot be based upon proportions or frequencies. In fact, they typically are: it is precisely the task of those whose profession is to estimate probabilities (in the insurance industry, for example) to base the estimates upon relevant proportions, frequencies and other data. More generally, to say that a degree of expectation is rational, it seems, is precisely to say that some such statistical facts are available in terms of which the expectation can be justified. However, such justifications often turn out to be extremely complex, and in many cases we seem to know that probability-estimates are well-founded even when we are unable to explain them adequately. Since degree of rational expectation is arguably the most inclusive or general characterisation of probability, our strategy here is to proceed without requiring that probability-claims always be justified in terms of statistics. The science of such justifications – taught in courses in Statistics and Probability Theory – is extremely interesting and of ever increasing importance.

3   There are various terms in the existing literature for this; another common one is 'epistemic probability'.

This can best be appreciated by thinking again of something like cards. Suppose George has a card face down on the table before him; he doesn't know what it is, but he does know with certainty that it's a red card (perhaps he caught enough of a glimpse to see that it's red, but couldn't tell whether it's a heart or a diamond). Since George knows that clubs are black, he can be perfectly certain – rationally completely certain – that the card is not a club. So his degree of rational expectation that the card is not a club is 1 (equivalently: his degree of rational expectation that the card is a club is 0). Provided he knows that hearts and diamonds are the only red suits, his degree of rational expectation that the card is a heart is 1/2.

Notice two things about this characterisation of probability in terms of rational expectation. First, it is clear that the basis for assigning degrees of rational expectation may consist in proportions. It is because George knows the relevant proportions that we can assign degrees of rational expectation in this case. In other cases, frequencies provide the basis for assigning degrees of rational expectation. In still other cases, we may assign degrees of rational expectation without knowing either relevant proportions or relevant frequencies.

Second, notice that when assigning degrees of rational expectation, we spoke of the degree to which one is entitled to believe something given such-and-such evidence. Since we are taking degree of rational expectation as our concept of probability, what this means can be expressed by saying that our key concept is the concept of **conditional probability**. That is to say, what we are interested in is the probability that a proposition is true, given that, or on the assumption that, some given set of propositions is true. More exactly, this is the degree to which it would be reasonable to accept a certain proposition, given no other relevant information except that contained within a certain set of propositions. In George's case, his evidence is that the card is red. So the conditional probability of the proposition 'the card is a heart', relative to the proposition 'the card is red', is 1/2.

We can now say more precisely what an inductively forceful argument is:

Let [P] stand for one or more premises, and let A stand for a conclusion. Suppose we have an argument:

[P] . . .

_____

C) A

To say that such an argument is _inductively forceful_ is to say that the conditional probability of A relative to the set [P] is greater than one-half, but less than 1.

Logic: inductive force

This is our 'official' explanation, but you may find it more helpful to think of inductive force along the lines of the following:

> **To say that an argument is *inductively forceful* is to say:**
> **The argument is not deductively valid, but: if the premises are true (or were true), then, given no other information about the subject-matter of the argument, it would be more reasonable to expect the conclusion to be true than it would to expect it to be false.**

There are several further points to bear in mind as regards probability and inductive force.

1 It is true that we do not always express probabilities as conditional probabilities. For example, if we simply pick a card at random from the deck, it seems we can say outright that the chance of its being an ace is 1/13. So it seems we do ordinarily attribute probability to a single proposition without stopping to specify any further information. Usually, however, this is only because the relevant further information upon which the probability-claim is based is left implicit; it goes without saying. In this case, the relevant information that one would normally take for granted is that the deck is standard and complete, in which case four of its fifty-two cards are aces. Thus, given the information that the deck is standard and complete, it would be reasonable for you to conclude that the card is probably not an ace. But so long as it is kept in mind that there must be some relevant set of information or premises in the picture, it is perfectly harmless to attribute probability to single propositions, and we will sometimes do so.

2 Probability, of the kind we are speaking of, is not an alternative to truth or falsity, in the way that finishing in the middle of the league table is an alternative to finishing at the top or at the bottom. Nor is probability a kind of truth. To say that a proposition is probable, in this sense, is to say that probably it is true. It is to express an expectation that a proposition is true that falls short of perfect certainty. What one says is correct if the evidence upon which the claim is based really does make the expectation rational. The idea is not that there is some third thing between truth and falsity, namely probability. For example, if I know that Reggie has taken his driving test today, and I say 'Probably, Reggie failed his driving test', I am not saying 'It isn't true that Reggie failed his driving test, and it isn't false; it's probable that he failed it'.

No: either it is true that Reggie failed his driving test or it is false. That is, there are two possible states of affairs: either he failed the test or he did not. There is no mysterious third state of affairs, i.e. that he probably failed it. Rather, the function of the word 'probably' is to indicate that the evidence makes it rational to believe, but does not make it certain, that the proposition is true.

3   Unlike truth, probability is a matter of degree; different propositions may have various degrees of probability (relative to a given body of information). For example, given the information we have, it is very highly probable that in the year 863 BC, someone ate a rabbit. But there is a very tiny possibility that no one did. It is somewhat less probable, but still probable, that someone will swim the English Channel during the year 2046. It is possible, but *im*probable, that someone swam the English Channel in the year 863 BC. We do sometimes specify probabilities in terms of numerical values, but often the probability of something being the case cannot be estimated with enough precision to justify assigning an exact numerical value. Sometimes the most we can do is to *rank* probabilities; for example we can say with confidence that a certain man is more likely to eat a cucumber than he is to visit the moon during the coming year, but we cannot assign precise numerical probabilities to these propositions, in the way that we can with the cards. And sometimes we cannot even rank them with confidence. Of single propositions, the most we can say in such cases is, 'that's probable', or 'that's very probable', or 'that's improbable'. (By 'probable', remember, we mean the case in which the probability is greater than 1/2; by 'improbable' we mean the case in which it is less than 1/2; something can be neither probable nor improbable, if its probability is exactly 1/2).

Because of this, inductive force, unlike deductive validity, is also a matter of degree. We cannot say that one argument is more valid than another, but we can say that one argument is more inductively forceful than another. Validity was defined in terms of impossibility, in which case the probability is zero (look at the definition of validity again). Validity is thus all-or-nothing. Indeed, one could define validity simply as the 'limiting case' of inductive force – the case in which the conditional probability of the conclusion relative to the premises is 1. Thus we can think of arguments as being arranged on a scale of conditional probability, ranging from deductively valid to inductively invalid (see Figure 3.1).

Because we defined inductive force the way we did, an argument may be inductively forceful but only to a very small degree. For example, if you know that fifteen of the twenty-nine children in the class are wearing white shoes, then you have an inductively forceful argument for the conclusion that the child who got the highest mark in spelling is wearing

Logic: inductive force
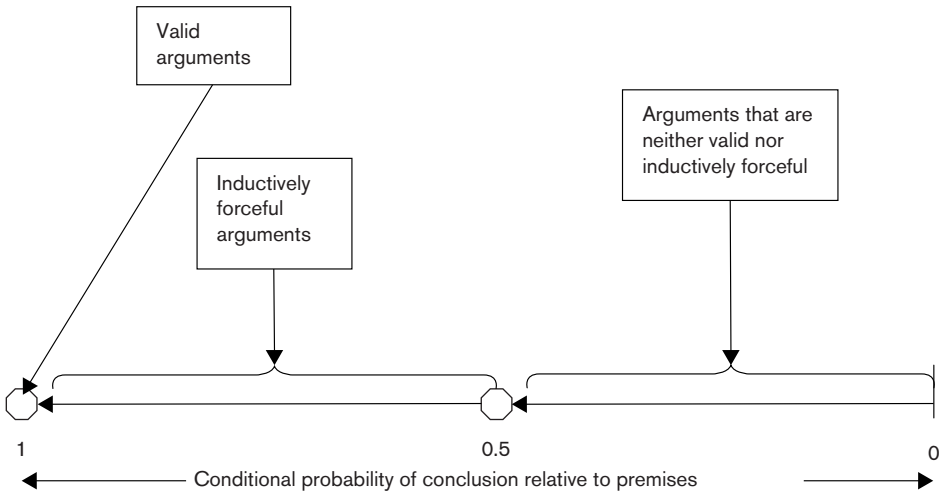
Figure 3.1 Arguments on a scale of conditional probability

white shoes. But the probability of that conclusion, relative to the premise, is only slightly better than 1/2; so the argument is just barely inductively forceful. It is inductively forceful, but only to a low degree.

4   Whether or not an argument is deductively valid does not depend on whether anyone thinks it is. Likewise, probability, in our sense, is the degree to which it is rational or reasonable to expect something to be true (given a certain set of premises), irrespective of how likely we actually think it to be. When we claim that an argument is inductively forceful, we are not always correct. We may think that a set of premises (our evidence) makes it reasonable to accept a conclusion, when in fact it is not. Rational expectation concerns what is in fact reasonable, not what any particular person thinks is reasonable.

We can reinforce this point by considering a case in which everyone will agree that rational expectation depends on proportion. Suppose that Fiona is asked to pick a card at random. Her chance of getting, say, a heart, is exactly one in four. Her chance of not getting a heart, then, is three in four, or 3/4. So probably she will not get a heart. If she is perfectly rational and well-informed, she will be 3/4 certain that she won't get a heart. Suppose, however, that Fiona is not perfectly rational. She thinks that since she has recently fallen in love, she probably will choose a heart. She might express her thinking so by saying, 'I'll probably choose a heart'. Nevertheless, the probability that she will not choose a heart is

greater – indeed three times as great – as the probability that she will. Fiona's actual degree of expectation is higher than the degree of rational expectation. More generally, we can observe that different people, going on the same set of premises, can have different degrees of expectation that something is or will be true; but there is only one correct answer as to how reasonable it is to infer the conclusion from the premises.

5   Like the validity or invalidity of an argument, the degree of inductive force of an argument is independent of the truth-values of the premises. This can be seen most easily in the case of a completely fictional example, in which the question of the truth-values of the premises cannot even arise:

> P1)   Almost all Zormons play chess.
> P2)   Trozak is a Zormon.
> _____
> C)   (Probably) Trozak plays chess.

Of course this is not even a real argument, since there are no such beings as Zormons, and no such being as Trozak. Nevertheless, if there were, then this certainly would be an inductively forceful argument: if there were a being named 'Trozak' who was a member of some class of beings called 'Zormons', and P1 were true, then it would be reasonable to infer C if P1 and P2 were all one had to go on. If you did not know whether or not these beings exist, you would still be right to say that if this did express a real argument, it would be inductively forceful. Whether or not there is such a being as Trozak or such beings as Zormons, and whether or not P1 is true, are thus irrelevant to the estimate of inductive force.

6   The same point goes for conclusions: estimates of inductive force should not be affected by information relevant to the truth-value of the conclusion, except for that given in the premises. Consider this argument:

> P1)   Sweden is a Scandinavian country.
> P2)   One Scandinavian country exports cars.
> _____
> C)   Sweden exports cars.

Almost everyone knows that Sweden – home of Volvo and Saab – exports cars. Furthermore, both premises of this argument are true. However it

is clearly not an inductively forceful argument. If the only propositions you knew relevant to the conclusion were P1 and P2, then you would be wrong to think that C is probably true.

'All', 'most' and 'some'

To say that all rodents have tails is the same as saying: 'Every rodent has a tail', or 'Any rodent has a tail', or 'No rodent has no tail'. To say that most rodents have tails is to say that more than half of all rodents have tails, or that there are more rodents with tails than there are without them. But what about 'Some rodents have tails'? What proportion of rodents must have tails for that to be true? In fact, that is a bad question. Such a statement does not tell you what proportion of rodents have tails. Consider this argument:

> P1) Some patients who have been treated with X have developed liver disease.
>
> ───
>
> C)   If I am treated with X then I will develop liver disease.

This argument is not inductively forceful. P1 could be true even if, say, three patients developed liver disease upon being treated with X, yet thousands were treated with X without developing liver disease.

That much should be fairly obvious. But the word 'some', as ordinarily used, can be rather tricky. Often, when we actually say something of the form 'Some A are B' (like P1), we say it when we believe that more than one A is B, but that not all A are B. Often when using the word this way, we say, 'Only some A are B'. Other times, we say it when we believe that more than one A is B, but do not know whether or not all A are B. For example, a medical researcher who discovers a few patients that have developed liver disease after being treated with X might use the sentence P1 to announce the discovery – but this would be to leave open the possibility that perhaps all patients treated with X develop liver disease. He would not be ruling it out.

For the purpose of reconstructing arguments it is most convenient to assume the latter understanding of 'some', whereby 'some A are B' does not rule out that all A are B. It does not say 'Some, but not all A are B'. In this sense, if it is true that some A are B, it might also be true that all A are B.

Further, when using 'some' in our argument-reconstructions, we will take it to mean 'at least one'. That is: if only one A is B, then it will be

true that some A are B. On this understanding of 'some', what 'Some rodents have tails' means is simply that it is not the case that no rodents have tails. Similarly, 'Some rodents do not have tails' means the same as 'Not all rodents have tails'. This departs slightly from what ordinary language typically suggests, but it is much more convenient for our purposes. If, when reconstructing an argument, it is clear that by 'some' the arguer means 'at least two', then we can simply make this explicit in the reconstruction, writing 'at least two'.

Inductive soundness

You can probably guess what this is, by analogy with the definition of deductive soundness:

> **To say that an argument is *inductively sound* is to say: it is inductively forceful and its premises are (actually) true.**

The important thing to note here is that an inductively sound argument, unlike a deductively sound argument, may have a false conclusion. That possibility is precisely what is left open by the definition of inductive force – an inductively forceful argument is the case in which the truth of the premises would probably make the conclusion true, but not necessarily make it true. Look again at the last argument given concerning Fiona, who lives in Inverness. Suppose it is true that Fiona lives in Inverness, and that almost everyone there has some woollen garments. An argument with those two facts as premises, and that Fiona has at least one woollen garment as a conclusion, would be inductively sound – even if, as it happens, Fiona has no woollen garments.

Note that we have not said that one should always be convinced by inductively sound arguments. You could know that an argument is inductively sound, but also know, for independent reasons, that the conclusion is false (review the discussion of inductive soundness if this is surprising). We will return to this point in Chapter 6.

Probability in the premises

So far in this chapter we have mostly considered premises expressing proportions, such as 'Most Zormons plays chess'. Proportions are expressed using quantifiers such as 'Most', '95 per cent of', and the like. But consider this case:

Logic: inductive force

> P1) If Napoleon is not ill then the French will attack.
> P2) Probably, Napoleon is not ill.
> _____
> C1) (Probably) The French will attack.

This argument is inductively forceful, but it does not contain any quantifiers. The reason that the truth of the premises would not guarantee that the French will attack, of course, is the presence of the word 'Probably' in P2; it is only said to be probable that Napoleon is not ill. (We would get the same result if in place of P2 we wrote 'Napoleon is probably not ill', or 'It is unlikely that Napoleon is ill', or some other sentence synonymous with P2 as written.)

Note that words such as 'probably' can also occur in the antecedents and consequents of conditionals. For example:

> P1) If Napoleon is not ill then, probably, the French will attack.
> P2) Napoleon is not ill.
> _____
> C1) (Probably) The French will attack.

This too is an inductively forceful argument.[4]

Arguments with multiple probabilistic premises

Sometimes, probabilistic elements can occur in more than one premise. Assessing such arguments can be tricky. For example, we may have the following sort of case:

_____

4   Some readers may wonder why we enclose the word 'probably' in parentheses before the conclusion of an inductive argument, but do not enclose it in parentheses when it occurs in a premise. The reason is somewhat complicated, but what it boils down to is that we want to preserve the simple and common-sense contrast between deductive and inductive arguments according to which the inductive case is that where the premises provide good but not water-tight reasons to accept the conclusion. Thus when we write '(Probably) Such-and-such' beneath the inference-bar, the conclusion is 'Such-and-such', not '(Probably) Such-and-such'. If the probability-indicator were regarded as part of the conclusion, then the question of how to distinguish deductive arguments from inductive becomes more complicated; we would also have to delve further into the theory of probability than is genuinely needed for our purposes here.

P1) Most people in Glasgow live in council housing.
P2) Most council housing is substandard.
P3) Ian lives in Glasgow.

C) Ian lives in substandard housing.

This argument is not inductively forceful. 'Most' means 'more than half', or 'more than 50 per cent'. The premises could well be true, then, in the following circumstances: slightly more than half the residents of Glasgow live in council housing; slightly more than half the council housing units in Glasgow are substandard, yet no other housing in Glasgow is substandard. If that were so, then the proportion of substandard housing in Glasgow would be slightly more than half of slightly more than half the total housing units in Glasgow. Since half of a half is $1/2 \times 1/2 = 1/4$, this means that, in these circumstances, the proportion of housing units in Glasgow that are substandard would be a little bit more than 1/4. So the mere fact that Ian lives in Glasgow, along with the truth of the other two premises, would not make it reasonable to expect that Ian lives in substandard housing.

Of course, if almost everyone in Glasgow lived in council housing, and almost all of it were substandard – then, provided that we knew these facts, Ian's living in Glasgow would be a good reason to infer that he lives in substandard housing.[5]

The same sort of thing applies when a word such as 'probably' occurs more than once in the premises: inductive force is not 'transmitted' across multiple inferences. Here is a somewhat more difficult case involving an explicitly conditional probability:

P1) Edna will probably come to the party.
P2) If Edna comes to the party, then probably she'll get drunk.

C) Edna will get drunk.

Is this argument inductively forceful? No. All that is required for something to be probable is that its probability be greater than 1/2. Suppose, then, that the chance of Edna coming to the party is exactly 51 per cent,

---

5   Of course, these premises are not true, and nor is it true that the only substandard housing in Glasgow is council housing!

and that if she were to come to the party, then the chance of her getting drunk would be 51 per cent. Then the premises of the above argument would be true. But the conclusion would not thereby be probable. It might be that if she does not come to the party then she certainly will not get drunk, in which case the probability of her getting drunk will be, again, only slightly greater than 1/4. Of course, if the probability of Edna's coming the party were higher, or the likelihood of her getting drunk if she did were higher, or both, then the probability of her getting drunk would be higher.

Some arguments contain a premise with a word like 'probably' *and* a premise with a quantifier such as 'most':

> P1) Probably, James will gain a science degree.
> P2) Most people with science degrees earn better-than-average salaries.
>
> ---
>
> C) James will earn a better-than-average salary.

Once again, this argument is not inductively forceful. However, if the argument were changed so as to indicate a sufficiently high probability of James' gaining a science degree, and/or to indicate a sufficiently high proportion of science degree-earners earning better-than-average salaries, then the argument would be inductively forceful. Of course the argument becomes inductively forceful if the word 'probably' is simply deleted from P1. It would also be inductively forceful if the word 'Most' in P2 were changed to 'All'; but it would not be inductively sound, since P2 would be false. Not everyone with a science degree earns a better-than-average salary.

Another tricky situation that sometimes arises is illustrated by the following:

> P1) Nine out of ten Green Party members are vegetarians.
> P2) Only one in ten non-Green Party members are vegetarians.
> P3) Alistair is a vegetarian.
>
> ---
>
> C) Alistair is a Green Party member.

It is tempting to think that this argument is inductively forceful, but this would be a mistake (the mistake is known as the 'base-rate fallacy'; see the

discussion in Chapter 4). In fact, even if P1 said that all Green Party members are vegetarians, the inference would be mistaken. For suppose that only one out of every 101 people (in the UK, for example) is a member of the Green Party. If the remaining (non-Green) population is 10 per cent vegetarian, then, since every Green is vegetarian, eleven out of every 101 people are vegetarian. But only one of those eleven is a Green. Thus ten out of every 101 people are non-Green vegetarians, and the non-Green vegetarians outnumber the Green vegetarians by ten to one! So if you meet a vegetarian such as Alistair, then, unless you had other reasons to think he is a Green, the reasonable expectation would be that he is not a Green. Summing up, then, the premises do not give you a good reason to think that Alistair is a Green.

The problem is that the premises seem to tell us that if someone is a vegetarian, then, probably, they are a Green. But they do not. Similarly: All plumbers own pipe wrenches, and few non-plumbers own them; but these facts do not entitle us to say that if someone owns a pipe wrench, then probably, they are a plumber. As a matter of fact, many DIYers own pipe wrenches. So the number of non-plumbers owning pipe wrenches may well outnumber the plumbers.

Inductive force in extended arguments

The conclusion of an inductively forceful argument may serve as a premise for a further argument, which may itself be deductively valid, inductively forceful, or neither. However, if any sub-argument of an extended argument is not deductively valid, then the argument as a whole is not deductively valid. At most, it is inductively forceful. Here is a hybrid extended argument which illustrates the point:

P1) If Napoleon is not ill then the French will attack.
P2) Probably, Napoleon is not ill.

C1) (Probably) The French will attack.

P3) If the French attack, then the Prussians will be routed.

C2) (Probably) The Prussians will be routed.

If the first argument were valid, entitling us to assert C1 categorically – that is, without the qualifier 'Probably' – then since the argument from C1 and P2 to C2 is deductively valid, there would be no need to write

Logic: inductive force

'Probably' beneath the second inference bar. As it stands, however, C2 inherits the qualifier 'Probably' from its ultimate source in P2. The simplest way to think of this is to think of C2 as a proposition inferred from three premises: P1, P2 and P3. If the word 'probably' were deleted from P2, then that argument would be deductively valid. As it stands, however, that argument is only inductively forceful.

Note that in carrying out the inference from C1 and P3 to C2, we think of that sub-argument just as if it were presented like this:

C1) Probably, the French will attack.

P3) If the French attack, then the Prussians will be routed.

C2) (Probably) The Prussians will be routed.

That is, we think of it just as if there were no parentheses around the word 'Probably' in C1. When the conclusion of an inductively forceful argument is regarded as a premise in a further argument, the parentheses can be ignored.

Conditional probability in the conclusion

Consider the following argument:

If you start a new internet company, then probably it will fold, or get bought by another company within three years. Most new internet companies do.

The conclusion of this argument is not 'Probably, your new internet company will fold or get bought by another company within three years'. It is, rather, 'If you start a new internet company, then, probably, it will fold or get bought by another company within three years'. Its conclusion, then, states a conditional probability. We can reconstruct it like this:

P1) Most new internet companies either fold, or get bought by another company, within three years.

C)   If you start a new internet company, then, probably, it will either fold, or get bought by another company, within three years.

You can see that the inference is correct. But is this argument deductively valid, or only inductively forceful? The conclusion, it seems, can be asserted on the basis of P1 with perfect confidence; surely there is no need to insert a 'Probably' before the conditional conclusion as well as before its consequent. It might seem, then, that the argument is deductively valid. However, there is a slight complication that one should be aware of. We said that statements of probability are usually statements of conditional probabilities, even if the evidence upon which the claim is based is left implicit. When we say something like 'Probably, your dog has fleas', we usually mean given everything we know, it is reasonable to accept the proposition. If so, then the consequent of C in the argument above really means something like:

> Relative to all the evidence, the conditional probability is
> greater than 1/2 that your new internet company will either fold,
> or get bought by another company, within three years.

If so, then, strictly speaking, the argument is not deductively valid: we might have other reasons to think that the person's company will not fold. In that case P1 would be true but C false. To make the argument valid, then, we would need the further premise that there is insufficient evidence that this person's new internet company would be significantly less likely than others to fold or be taken over. Alternatively, one could change the conclusion's consequent to 'unless there is other evidence to suppose otherwise, your company will fold or be taken over within three years'.

Evidence

Arguments presented as being inductively forceful may have several premises. Consider the following:

> The case against Dr X can be summed up as follows. Dr X, a
> medical doctor, has the knowledge, and had the opportunity, to
> carry out the type of gradual programme of strychnine
> poisoning that killed the victim. He had a strong motive for the
> murder. And finally, Miss Y, a reliable witness, has testified that
> Dr X expressed to her his wish that the victim were dead.

Let us grant that the case against Dr X is compelling. Here are four facts that seem jointly to condemn Dr X as the murderer. The case does not take the form of a deductively valid argument: It would not be impossible for these facts to be as the arguer says, yet Dr X innocent. But it

Logic: inductive force

is crucial to recognise that the case should not be reconstructed as four separate arguments, each presented as inductively forceful in its own right. The prosecutor's case is based on four premises:

P1) Dr X has the knowledge to carry out the type of gradual programme of strychnine poisoning that killed the victim.
P2) Dr X had the opportunity to carry out the murder.
P3) Dr X had a strong motive for the murder.
P4) Miss Y, a reliable witness, has testified that Dr X expressed to her his wish that the victim were dead.

These constitute the evidence against Dr X. But an argument for the conclusion that Dr X murdered the victim based on P1 alone (along with a suitable connecting premise), would not be inductively forceful. If it were, then the argument would point with equal force to the guilt of a great many other people with the requisite knowledge of strychnine. Nor would an argument similarly based only upon P2, P3 or P4. If the prosecutor were to advance any of these four separate arguments by itself, then the jury could only conclude that the prosecutor lacks a compelling case: none of the arguments would be inductively forceful. It is essential to the prosecutor's case that these four items of evidence point jointly to the defendant's guilt. The argument, then, should be constructed like this:

P1) Dr X has the knowledge to carry out the type of gradual programme of strychnine poisoning that killed the victim.
P2) Dr X had the opportunity to commit the murder.
P3) Dr X had a strong motive for the murder.
P4) Miss Y, a reliable witness, has testified that Dr X expressed to her his wish that the victim were dead.

C) Probably Dr X murdered the victim.

This is how it is with items of evidence: If both B and C are evidence for A, the conjunction of B and C is (typically) stronger evidence than B or C alone.[6] Put in terms of arguments, the argument with B and C as premises is more forceful than either the argument with B as a premise

---

6   This holds so long as B and C are independent: whether or not B occurs has no bearing on whether or not C occurs, and vice versa. Otherwise the situation is more complicated.

or that with C as a premise. The degree to which this is so, however, cannot be estimated without looking at the facts of the particular case. For example, we might know that the murderer is a man from Harrogate, is left-handed and walks with a limp. Obviously a Harrogate man's being left-handed would not be strong evidence against him, and nor would his walking with a limp. His being both left-handed and walking with a limp would be stronger evidence against him, but to what degree? It depends on how many left-handed men there are in Harrogate who walk with a limp. If there happened to be exactly one, the evidence would be conclusive (the argument would be deductively valid). If there were more, then the evidence would be much weaker.

Inductive inferences

Consider the following argument:

> So far, every guppy I've had has died when I fed it cat food.
> Therefore, if I feed my new guppy cat food, it will die.

This inference seems reasonable, but how do we reconstruct the argument? We might try:

P1 (a) All guppies die if fed cat food.
_____

C)   My new guppy will die if I feed it cat food.

Or:

P1 (b) Most guppies die if fed cat food.
_____

C)   (Probably) My new guppy will die if I feed it cat food.

The first argument is valid and the second inductively forceful, and perhaps both are sound. But there is a logical problem. The arguer has not asserted either version of P1. What he has asserted is that all the guppies which he has fed cat food have died. The sample of guppies with which the arguer has had this experience is only a tiny fraction of the total number of guppies in the universe. So the arguer has not said

anything about all or most guppies. Presumably the arguer does believe either P1(a) or P1(b), but if so, he does so on the basis of his previous experience with guppies and cat food. Somehow, the arguer must get from the proposition 'Every guppy I've had has died when I fed it cat food' to 'Most (or all) guppies die when fed cat food'. So far, in our discussions of arguments, we have said nothing to help us with this sort of inference.[7]

However, this kind of inference is very common, both in science and in ordinary life. We call it an **inductive inference**. This is our name for the case when you extrapolate from a sample of a total population of things either to something outside the sample, or to a generalisation about the population as a whole (so far, we have always gone the other way, either from a generalisation about a whole population to a claim about some smaller sample, or from the whole population to a claim about a single member of it). In the above case, the arguer needs to extrapolate directly from the sample of guppies they have experimented with to a proposition about the new guppy:

P1)   Every observed guppy has died when fed cat food.

C1)   My new guppy will die if I feed it cat food.

If we judge this inference to be inductively forceful – let as assume that it is – then, just as with any inductively forceful argument, we write 'Probably', in parentheses, before the conclusion:

P1)   Every observed guppy has died when fed cat food.

C1)   (Probably) My new guppy will die if I feed it cat food.

---

7   As we pointed out, many probabilities can be based upon proportion. One of the tasks of statistical analysis is to show how this is so, even in cases that might initially seem to resist this strategy (see Note 2). Inductive inferences represent a class of cases where proportion certainly cannot be applied straightforwardly. If we know that, say, 74 per cent of the past 956 cases of A have been cases of B, we cannot without further ado assume that the probability that the next case of A will be a case of B is 74 per cent. The next case of A is not amongst the set of cases for which that proportion is known to have held. So it is not like the case of a card face down on the table, which is drawn from a set of cards for which the relevant proportion is known. In the inductive case we have to make the further assumption that probabilities taken from proportions of known cases can be transferred to unknown cases. Of course such inferences do often seem to be rationally justified; the question of why this is so – the 'Problem of Induction' – was first raised by David Hume in his *Treatise of Human Nature* (1739), and remains an open problem.

Alternatively, the arguer could have obtained this conclusion by a slightly longer route. The arguer might have reached the conclusion by means of an extended argument, comprising an inductive inference to a generalisation about all guppies, together with a second, deductively valid inference from that generalisation to the conclusion concerning the arguer's new guppy:

> P1) Every observed guppy has died when fed cat food.
> _____
> C1) All guppies die if fed cat food.
> _____
> C2) My new guppy will die if I feed it cat food.

However, assuming the truth of P1, one could not assert the conclusion here with quite as much confidence as in the first case. We may assert C2 with as much, but no more, confidence than that with which we may assert C1. But since C1 is a generalisation about all guppies, it stands a greater chance of being false than C1 of the first guppy-argument – a statement about a single guppy. The reason is that the second argument depends upon a more ambitious extrapolation than the first, one that is not needed for the desired conclusion. Indeed, the inferences to C1 in either case are fundamentally of the same type. They are both extrapolations from observed cases to unobserved cases. Each can be thought of as an extrapolation from a sample to a larger population of which the sample is a part. In the first case we can think of the larger population referred to as comprising the sample (the observed guppies) plus this one new guppy. In the second case, obviously, the larger population is simply the total population of guppies. Some inductive inferences are concerned to make a broad generalisation (all guppies), others are concerned only to extend to a few new cases, or to a single new case (the new guppy). But they are fundamentally the same type of inference. Thus we define:

> To say that an inference is an *inductive inference* is to say:
> (a) it is not deductively valid; (b) its premise is a
> generalisation about a sample of a given population, and
> (c) its conclusion is a generalisation about the total
> population from which the sample is drawn.

Requirement (a) is to avoid counting the following sort of inference as inductive: 'Some black cats have no tails; therefore not all cats have tails'.

Logic: inductive force

Inductive inferences frequently extrapolate from past to future. For example, if we infer, from Ireland's never having won the football World Cup, that probably it never will, our sample is all the World Cups that have so far been contested, and the total population is the set of all World Cups – past, present and future.

Sometimes the conclusion of an inductive inference is a more precise statistical generalisation such as '37 per cent'. For example, when a pollster finds that 37 per cent of a sample of the adult British population supports a given political party, the pollster might conclude that 37 per cent of the adult British population supports that party. How probable this conclusion is, however, will depend on how **representative** the sample is. This point requires further explanation.

How representative is the sample?

Suppose you live near the North Pole, and every bear you've ever seen is white. So you argue:

P1) Every observed bear is white.

___

C) All bears are white.

Formally, this argument is just like the one about the guppies. But, whereas the guppy-inference seemed to be correct, this one seems not to be. Only a small minority of bears are white (Polar Bears, and perhaps the occasional albino). The difference is that, whereas the guppy-arguer can reasonably assume that the sample of observed guppies was **representative** of the total population of guppies, the bear-arguer cannot. That is, the guppy-arguer can reasonably assume that the sample of guppies was **relevantly similar** to the total population, but the bear-arguer cannot. Polar bears are only one of many species of bears, and if you know anything about zoology and adaptation, you know that a trait like colour, though somewhat likely to be similar across a single species, is not nearly so likely to be similar across different species, even within the same genus. Further, you can reasonably assume that mammals living in non-snowy regions are much less likely to be white than ones living in perpetual snow. Being a different species, which lives in a completely different climate, is certainly a relevant difference when drawing inductive inferences about colour with respect to classes of mammals, even if the species is a closely related one.

How do you know whether a sample is representative of the total population, i.e. relevantly similar to it? There is no simple rule for this; our estimate of relevant similarity must be based upon our knowledge of the subject-matter in question. For example, we drew upon some basic biology in criticising the bear-inference.

In general, the larger and more representative the sample used for a given generalisation, the more inductively forceful, or the **stronger** the inductive inference. An inductive inference that is not forceful – one whose premise does not really support its conclusion – is a **weak** one. For example, suppose someone argues, 'My brother is a fool. Therefore, probably, all boys are fools.' That is a very weak inductive inference; it is totally unwarranted. The bear-arguer's inference is not quite as bad as that, but it is also very weak. Also, from a given sample, an inference to a generalisation about *most* of the totality is a stronger one than an inference to a generalisation about *all* of the totality, since obviously the latter is less certain.

A programme for assessment

We are now in a position to outline a basic procedure for the assessment of arguments represented in standard form. We first consider arguments that have only one inference. When you represent an argument in standard form, do not include the word 'probably', or any similar word, in the conclusion – not yet. Put such words in the premises only. Once you have reconstructed the argument in this way, you should proceed as follows (and you should, looking back at the definitions given in this chapter, stop to see why the chart is arranged as it is):

1   Is the argument deductively valid?
    If not, proceed to 2.
    If yes, are all the premises true?
    If yes, the argument is deductively sound. Stop.
    If not, the argument is valid but unsound. Stop.

2   Does the argument consist of an inductive inference?
    If not, proceed to 3.
    If yes, write 'Probably', or suitable variant, before the conclusion of the inference. Is the inference a forceful one?
    If yes, proceed to 3.
    If not, the argument is neither valid nor inductively forceful. Stop.

Logic: inductive force

3 Is the argument inductively forceful?

If not, the argument is neither valid nor inductively forceful. Stop.

If yes, how forceful? Write 'Probably', or suitable variant, in the appropriate place in the conclusion. Are the premises true?

If yes, the argument is inductively sound. Stop.

If not, the argument is neither deductively nor inductively sound. Stop.

In the case of extended arguments, this procedure is to be carried out first with respect to the sub-arguments whose conclusions are premises of the extended argument. For example, suppose we have an extended argument whose form is like this:

P1) ...
P2) ...
_____
C1) ...

P3) ...
_____
C2) ...

In this case the argument whose conclusion is C1 is a sub-argument for an extended argument whose conclusion is C2. You would first assess the argument from P1 and P2 to C1; you would then assess the argument from C1 and P3 to C2. Finally, you would use these results to assess the total argument from P1, P2 and P3 to C2.

Assessment is to be distinguished from **reconstruction**, which is discussed in more detail in the next chapter. Broadly speaking, assessment may be said to fall into two categories: The first we may call **logical assessment**, which is that whereby we make judgements of deductive validity and inductive force. The second we may call **factual assessment**; here is where we ask after the truth-values of premises, in determining the soundness of arguments already found to be deductively valid or inductively forceful. However, as noted, the evaluation of inductive inferences is not purely the one nor purely the other: it is partly logical and partly factual.

It is important to realise that if your verdict is that the argument, as you have reconstructed it, is not *de*ductively sound, you might be able to reconstruct it so as to make it *in*ductively sound. Look again at the example involving Fiona and wool. An early attempt at reconstructing it

used the premise 'Everyone in Inverness owns at least one article of woollen clothing'. So rendered, the argument was valid, but not sound. We then weakened the premise to 'Almost everyone in Inverness owns at least one article of woollen clothing'. The result was not deductively valid, but it was inductively forceful. In general, it is best to sacrifice validity to soundness when reconstructing.

Later, we will learn some techniques which aid logical assessment. There is, you will appreciate, no readily definable procedure for factual assessment; for this is simply the task of determining whether particular propositions are true – e.g. whether there are blue whales in the Black Sea, whether inflation rose in 1993, and so on. For that we must look outside logic books!

CHAPTER SUMMARY

The **conditional probability** of a proposition relative to a set of premises is the degree to which it would be rational to expect it to be true, given no information relevant to the conclusion except that given in the premises. Estimates of **rational expectation** are often based upon other measures of probability such as **proportion** and **frequency**, but need not be.

A deductively valid argument enables us to be certain of the truth of its conclusion, if we are certain of the truth of its premises. An **inductively forceful** argument does not allow this, but it does allow us, in the absence of other information relevant to the truth-value of the conclusion, to think its conclusion more likely to be true than not: an inductively forceful argument is one for which the conditional probability of the conclusion relative to the premises is greater than 0.5.

Like deductive validity, both probability and inductive force are objective: the conditional probability of a conclusion relative to a given set of premises is independent of people's actual estimates of conditional probability. But unlike deductive validity, both probability and inductive force are a matter of degree. An argument may be inductively forceful but only just barely.

Inductively sound arguments are inductively forceful arguments with true premises. Unlike deductively sound arguments, they may have false conclusions.

Some arguments contain **inductive inferences**. An inductive inference is a deductively invalid inference whose premise is a generalisation about some **sample** of a given population, and whose conclusion is a generalisation about the population as a

Logic: inductive force

whole (the population may comprise simply the original sample plus one new case). In general, the force of an inductive inference depends on the degree to which the sample is **representative** of the population.

1 Each argument below contains a premise with a missing quantifier. Reconstruct the argument, adding an appropriate quantifier. If possible, make the premise true. Use the strongest quantifier you can, so long as the premise still seems to be true (that is, don't use 'most', if a phrase such as 'almost all' seems to you to result in a true premise). If the resulting argument is deductively valid, say so. If the argument is inductively forceful, insert 'Probably' before the conclusion.

a   Sex crimes are committed by people who have themselves been victims of child abuse. The defendant has committed a sex crime. Therefore the defendant was a victim of child abuse.

b   Children love the Harry Potter books. Therefore your child will love the Harry Potter books.

c   If doctors in this country are overworked, then the general health of the population will decline. Doctors in this country are overworked. Therefore, the general health of the population will decline.

d   Old people have poor eyesight. No one with poor eyesight should be allowed to drive. Therefore, most old people should not be allowed to drive.

e   In countries colonised by Spain, there are many people of Spanish descent. There are very few people of Spanish descent in Laos. Therefore Laos was not colonised by Spain.

2 Consider the following pair of arguments.

P1) Most AIDS patients in California are homosexual.
P2) Mr X is a California AIDS patient.

C)   Mr X is a homosexual.

P1) Most residents of Orange County, California, are not homosexual.

P2) Mr X is a resident of Orange County, California.

C) Mr X is not homosexual.

Assume that this is the same Mr X referred to in each argument. Are these arguments inductively forceful? Could both arguments be inductively sound? If not, explain why not. Otherwise, describe a scenario in which they would both be sound. If we knew the premises of both arguments to be true, would we have reason to believe, or not to believe, that Mr X is homosexual?

3 Two statements are *incompatible* if it is impossible for both of them to be true (if they are incompatible, it may or may not be possible for both to be false). Which of the following pairs of statements are incompatible?

a   Some hamsters are black. All hamsters are white.

b   Some hamsters are black. Some hamsters are white.

c   Some hamsters are black. Most hamsters are not black.

d   Some hamsters are black. All hamsters are black.

e   No hamsters are black. Most hamsters are black.

f   Every hamster is black. This hamster is white.

g   This hamster is black. Some hamsters are white.

h   No hamster is black. No hamster is not black.

i   No hamster is black. No hamster is white.

j   Some hamsters are black. There is exactly one black hamster.

k   Most hamsters are black. Many hamsters are white.

l   Most hamsters are black. All my hamsters are white.

m   Most hamsters are black. Every hamster in England is white.

Logic: inductive force

---

4 Reconstruct the following as either three separate arguments or as one, depending on which strategy would amount to a more forceful case.

Lewis is probably not going to win another gold medal in the 100 metres. He is now older than any previous winner of the event. His form has been poor this year, and he has a hamstring injury.

---

5 Reconstruct the following arguments. Say whether or not they are inductively forceful. If the argument is not inductively forceful, change one premise so that it is.

a   If alcohol advertising is banned, then, probably, drinking would decline. If drinking declines, then, domestic violence would decline. Therefore, if alcohol advertising is banned, then, probably, domestic violence would decline.

b   If alcohol advertising is banned, then, probably, drinking would decline. If drinking declines, then, probably, domestic violence would decline. Therefore, if alcohol advertising is banned, then, probably, domestic violence would decline.

c   Most children who go without breakfast have trouble concentrating at school in the morning. Johnny concentrates well at school in the morning. Therefore, probably, he does not go without breakfast.

d   If this meat was grown in Scotland, then it is extremely unlikely that it is infected with BSE. And even if it is infected, then it is very unlikely that eating it will make you ill. Therefore, if this meat was grown in Scotland, then it is not probable that eating this meat will make you ill.

e   Brazil is more likely to win the World Cup than Argentina. Therefore, probably, Brazil will win the World Cup.

f   Scotland has never won the World Cup. Therefore, probably, Scotland will never win the World Cup.

g   Probably, English football fans will make trouble at the next World Cup. If English football fans do make trouble at the next World Cup, then there is a good chance that England will be expelled from the European Cup. Therefore, it is very likely that England will be expelled from the European Cup.

h   Very few habitual cannabis users are violent. The murderer is
    clearly very violent. Mr X is a habitual cannabis user. Therefore,
    probably, Mr X is not the murderer.

i   UN intervention in local conflicts usually leads to a more stable
    political situation. A more stable political situation usually leads to
    economic growth. Therefore, probably, UN intervention in this local
    conflict will lead to economic growth.

j   Almost all successful athletes have trained long and hard to
    become successful. Therefore, if you train long and hard, you will
    become a successful athlete.

k   Some teachers at Lenman secondary school are known to smoke
    cannabis. Mr X is a teacher at Lenman secondary school.
    Therefore, probably, Mr X smokes cannabis.

l   Most vegetarians eat eggs. A great many people in Berkeley are
    vegetarians. Roger is from Berkeley. Therefore, probably, Roger
    eats eggs.

m   If Hansen is leaving early tomorrow, then, probably, he isn't in the
    bar. Hansen isn't in the bar. So probably he is leaving early
    tomorrow.

---

6   The following passages contain inductive inferences. First, rewrite
them as arguments in standard form. (You may have to supply a
missing quantifier.) Can you think of any reasons why the sample might
not be representative? If so, list them.

a   The elimination of sweets and other junk food from school lunch
    menus has resulted in improved student performance in schools all
    over the country, and in every case in which it has been tried. If we
    do this at our school, then, we can look forward to improved
    performance.

b   No communist system can succeed. Since communist regimes
    began early in the twentieth century, every one has either
    collapsed, or is near collapse.

c   Surprisingly, musicians have a higher average IQ than either
    doctors or lawyers. A study of seventeen conductors of major
    orchestras found their average IQ to be 17 points higher than that
    of UK doctors, and 18 points higher than that of UK lawyers.

Logic: inductive force

d   Most teenagers who take illicit drugs have serious family problems. I know, because as a counsellor for teenagers I have found that most of those who have come to me and confessed to taking illicit drugs do have serious family problems.

e   Manchester United is not going to win the treble again. No English team has won it twice.

f   Studies have shown that people who take vitamins regularly live longer than average. So if you take vitamins regularly, then, probably, you will live longer than average.

Chapter 4

Rhetorical ploys
and fallacies

Sometimes we are moved to accept or reject claims when we have been given no good grounds for doing so. Often this is because speakers or writers attempt to persuade us in ways which appear to provide good reasons but which do not really. We'll call these persuasive devices **sham-reasons** and the process of employing them **sham-reasoning**. As critical thinkers we should be alert to the possibility of sham-reasoning, take care to avoid being persuaded by attempts to persuade which rely upon it, and avoid using sham-reasons in our own attempts to persuade others. We are interested in two types of sham-reasoning: **rhetorical ploys** and **fallacies**. This chapter aims to equip you to distinguish between rhetorical ploys and fallacies; to familiarise you with various common types of sham-reasoning, and to develop strategies of reconstruction and evaluation that will enable you to deal with them when analysing, assessing and constructing attempts to persuade.

Neither rhetorical ploys nor fallacies provide us with good reasons to accept the claim they are intended to support. **Fallacies** are **argumentative sham-reasoning**. That is, they are still arguments in the sense that fits our definition of a set of propositions, some of which are premises, one of which is a conclusion, the latter intended to follow from the former. But in one way or another, they are bad arguments. **Rhetorical ploys**, on the other hand, are **non-argumentative sham-reasoning**: some of these persuasive devices may pretend to provide reasons for accepting a claim, but their real persuasive capacity depends on something non-argumentative. Recall our earlier definition of rhetoric as:

Rhetorical ploys and fallacies

> *Rhetoric* – Any verbal or written attempt to persuade someone to believe, desire or do something that does not attempt to give good reasons for the belief, desire or action, but attempts to motivate that belief, desire or action solely through the power of the words used.

The difference between fallacies and rhetorical ploys is understood most easily as a difference in the **function** of the language being employed. As we saw in Chapter 1, politicians, advertisers and newspaper columnists tend to be experts when it comes to using rhetorical ploys. Rhetorical ploys typically make a more or less direct appeal to feeling and emotion rather than to reason, which is the domain of argument. Fallacies, on the other hand, are simply defective attempts at argument (they may be defective in any of various different ways, as explained below). They may fool us into thinking they are not defective, but they are still presented as attempts at argument. Of course, many writers and speakers will use a mixture of rhetorical ploys, fallacies and genuine arguments when attempting to persuade us of the truth of their claim. In fact, it is possible for a given form of words, as advanced by a would-be persuader, to constitute a fallacy, yet, at the same time, function as a rhetorical ploy. For example, the following sort of advertisement is ubiquitous:

> More mothers use Namby-pambies for their babies than any other disposable nappy. Shouldn't you?

As we will explain in more detail below, this is an example of the fallacy of **majority belief**: The advertisement wants mothers to accept the argument: X is used by more people than any other product of its kind; therefore X is the best such product. That argument is fallacious because its implicit premise – that the most popular thing is the best, or most popular belief is true – is unjustified. But the advertisement also exemplifies the rhetorical ploy of **appeal to popularity**: the mere fact that something is popular often **causes** us to desire it (possibly by awakening a fear of being left out). The point can be put this way: the detection of fallacies is an exercise in logical and factual assessment: it involves the assessment of reasons. The detection of rhetorical ploys, on the other hand, is an exercise of psychology: we consider ways in which our desires,

fears, beliefs and actions could be non-rationally influenced by uses of language which are intended to persuade us to hold beliefs and perform actions.

As critical thinkers, we aim to become adept at distinguishing and identifying these different types of sham-reasoning. We want to understand how they work, and how to avoid being taken in by them. In the next section, we consider specific types of rhetorical ploy and examples thereof. Later we turn to a comprehensive consideration of fallacies.

Rhetorical ploys

Appeals to specific feelings

There is a range of rhetorical ploys that attempt to tap into specific feelings in order to influence our behaviour and opinions (especially our consumer behaviour). Here we discuss a number of the most common. Some of them are not strictly or specifically linguistic ploys, but this will not affect the points we need to make.

Appeal to novelty

Here someone attempts to persuade us to try or buy something because the item is new and, by implication, different from and better than existing related items. Often this ploy appeals to our desire not to miss out on a new trend, or arouses our fear of appearing outdated in our tastes. Sometimes it appeals to our (vain) sense of our selves as flexible and willing to try out new experiences. It may also persuade us that because the product is new, it must be an improved version of the existing product. Familiar examples of the appeal to novelty include the re-naming of the Labour Party as the New Labour Party; advertising that attempts to persuade us to upgrade to frequently newly released, but little-changed or improved computer software packages, and attempts to persuade us to switch to allegedly new and improved washing detergents. The appeal to novelty may also be employed to persuade us to adopt new ideas or beliefs. Again, this operates on our desire not to appear inflexible or stick-in-the-mud. Thus we might be enjoined to give up a belief in the value of trade union participation on the grounds that trade unionism is 'old-fashioned' and 'has no place in the workplace of the twenty-first century'. Notice that we have been given no reason to

reject trade unionism, we are just told that it is 'old-fashioned' and thus, by implication, something undesirable that we wouldn't want to be involved with.

Appeal to popularity

Like the appeal to novelty, this ploy appeals to our desire to run with the crowd, not to appear different from the norm and not to miss out on what others have. Again, it is commonly used to persuade us to buy things, but also occurs frequently as a means to persuade us to adopt a belief or to follow a certain course of action. Consider the following attempt to persuade students to buy a computer software package.

> Get Hotstuff! – the best-selling comprehensive software package among today's students – and turn your assignments into hot stuff.

Such an advertisement can work in various ways. It makes a straight-forward appeal to our desire to have what others have and not to miss out on the benefits enjoyed by those who already use the package in question. But in addition, it tends to lead us to make some unjustified assumptions about why the product is the most popular ('best-selling'). If we take insufficient care, we are inclined to think that the most popular product must be the most effective, but its popularity might stem purely from its competitive price or the success of the marketing campaign. The software package's being the best-selling does not give us a compelling reason either to buy the product or to conclude that it must be the most suitable software package for student use.

Like many rhetorical ploys, the appeal to popularity can sometimes be construed as presenting an argument – in this case, as presenting a reason to buy something. As we can see from the following reconstruction, the reason given is not typically a good one (so the argument is fallacious):

P1) The software package that sells best amongst students must be the best software package for students.
P2) The software package that sells best amongst students is Hotstuff.

C1) Hotstuff is the best software package for students.

This argument would be valid but almost certainly not sound because the assumption made explicit in P1 is not plausible; it is not the case in general that the best-selling product of a given kind is the best of that kind (see the discussion of the fallacy of majority belief, below). Certainly if you were a student intending to invest in an expensive software package, you would want much stronger evidence that Hotstuff is the best (and furthermore, you might well think that since different students have different needs, the claim that a given software package is the best for students in general must be pretty vague, and without any direct implication for you, with your particular needs; we will discuss this kind of vagueness in more detail below).

Appeal to compassion, pity or guilt

This common rhetorical ploy operates by attempting to move us to do something purely by evoking a feeling of compassion towards the recipients of the suggested act or belief, or a feeling of guilt about their plight. The feeling alone does not provide a good reason for us to perform the act in question. Examples of this ploy are myriad. Charity appeals and advertising provide plenty of instances. Pictures of hungry children, suffering animals and so forth accompanied by simple slogans are designed to tweak our feelings of compassion, pity and guilt about the situation of the people or animals featured. Remember that the primary purpose of our examination of rhetorical ploys is a critical one; that is, we encourage critical thinkers not to be persuaded by rhetoric and to avoid rhetoric in their own attempts rationally to persuade others.

As we pointed out in Chapter 1, however, rhetoric can play a useful role when put to good ends. In the case of the appeal to compassion, it serves a positive role by pricking our conscience and opening us up to rational argument regarding how we should act in the given situation. Feelings of compassion can prompt us to look for arguments that do provide good reasons for the course of action recommended by the rhetorical ploy. Suppose for example you read a charity advertisement that appeals to your sense of compassion towards hungry children. You may go on to reason that if one can do something to alleviate hardship, one should; and if you are in a position to do so then you should make a donation to the charity. Such an argument might look thus when reconstructed and rendered in standard form:

Rhetorical ploys and fallacies

> P1) A donation to the Worldwide Fund For Children would probably help to alleviate the suffering of children from extreme hunger.
>
> P2) I should try to alleviate extreme suffering where it's possible for me to do so.
>
> P3) It is possible for me to make a donation to the WWFC.
>
> ___
>
> C) (Probably) I should make a donation to the WWFC.

Notice that the argument is inductively forceful, its premises make it probable that I should act in the way prescribed. It is certainly not implausible to say that this argument is sound. Suppose then that you were moved to donate to the WWFC by a message simply stating that many children throughout the world suffer from extreme hunger, along with instructions how to donate to the WWFC. The message certainly does not present the above argument, nor any other. But it might have moved you to act by causing you to think of the argument yourself.

Appeal to cuteness

This rhetorical technique uses children, animals or animated characters to deliver a message. The product that we are urged to buy or the action we are urged to take is made to seem attractive by its association with the cute character that urges us to buy it or take it. The appeal can also work by helping us to remember a product (what advertisers would call making us brand-aware) via its association with the cute figure delivering the sales pitch.

Appeal to sexiness

This is similar to the appeal to cuteness, except that it does not use children or animals. But it has a further dimension. To those who would desire the sexy person depicted in the advertisement, the product is made to seem desirable by its association with the sexy person. But it is also made to seem desirable to those who would like to think of themselves as sexy in the way that the sexy person is. The ad seeks to flatter us. It may even invite us to reason, fallaciously: All sexy people buy or do this; therefore if I buy or do this, I'm sexy. This would be the **fallacy of affirming the consequent**, discussed below.

Appeals to wealth, status, power, hipness, coolness, etc.

The reader can easily work out what these are by analogy with the appeals to cuteness and sexiness.

Appeal to fear (also known as scare tactics)

This is the tactic of trying to elicit a fear in one's readers or listeners in order to influence their behaviour or attitudes. A frequent example of the appeal to fear occurs in discussions about immigration into countries such as the UK. Many politicians and other opinion-formers such as journalists use the tactic of eliciting citizens' fears of economic destitution and cultural demise by constructing deliberately exaggerated images of 'waves of immigrants'[1] entering a country illegally and generally living a life better than that they deserve; taking jobs, education, health care and state benefits to which they have no rightful entitlement and thereby making Mr and Ms Average Citizen worse off. This familiar scare tactic (which almost certainly also makes appeal to some people's racist attitudes) is often used by politicians to persuade people to support draconian immigration policies and infringements of people's civil rights, and to demonstrate that support by voting for them at election time. But no reason has been given for the belief that disaster would result if such extreme policies were not enacted. Instead, it is hoped that describing these disastrous scenarios will alarm people so severely as to disturb their reason, prompting the confused supposition that the apparent disastrousness of the worst possible scenario should be matched by the severity of the preventive measures taken.

The appeal to fear should be distinguished from genuine warnings. In instances of the former, there is no warranted connection between the fear elicited and taking the suggested course of action or accepting the claim. Whereas in the case of a warning, we are given a good reason to act. This is usually because the circumstances of the warning are themselves such as to warrant the belief that the warning is well-founded. For example, the warning 'Don't touch the dog, it may look cute and friendly but it bites!' would normally be given only by someone who knows that the dog bites; since such warnings are very seldom given insincerely or with a deceitful motive, the fact that such a warning has been given is good reason to heed it. Such an inference, needless to say, would be much more doubtful in the realm of advertising or political discourse.

These are the most common instances of rhetorical ploys intended to manipulate specific feelings and thereby to influence our attitudes and behaviour, but there are others which occur frequently. Try to think of some and find some examples of your own.

---

1   Notice the rhetorical power of the phrase 'waves of immigrants', connoting as it does a flood that will sweep people away.

Rhetorical ploys and fallacies

The direct attack and hard sell

The direct attack is the simplest of all rhetorical ploys. It occurs most frequently in advertising, though it also appears in political campaigning. It often takes the form of a very simple slogan. For instance, 'Say no to tuition fees!', 'Drink Cola!' Notice that we are given no reason to say 'no' or to drink cola. The belief of those who employ the direct attack is that the more we hear or read these commands and internalise them, the more likely we are unreflectively to do as they advocate, despite having been given no reason to do so. We often talk about 'giving someone the hard sell'. The hard sell is simply the direct attack repeated persistently. Children are notably effective with it. The persuader just keeps it up until the subject of their attack gives in and does as they want, the persuader thereby having influenced their target by verbal means without giving reasons for doing as they command.

Buzzwords

This is the technique of using fashionable or otherwise currently 'hot' words or phrases which are loaded with rhetorical power due to their rich secondary connotation. (If you don't feel familiar with the concept of secondary connotation, have another look at the relevant section of Chapter 1.) Buzzwords can be enormously provocative and therefore hard to tame, and this makes them especially problematic for the critical thinker. If we want to make an objective analysis of a passage or speech act, we should rephrase what is said or written in such a way as to eliminate the buzzwords, and then embark on the analysis. Here is an example of a passage containing several buzzwords:

> The Prime Minister's solid stance against European Union bureaucrats' latest attempts to create yet more employment rights for European workers and yet more financial burdens for European employers sent a message to business that his Government would continue to stand tall in its commitment to the free market and to wealth creation.

The writer uses the terms 'bureaucrats', 'business', 'free market' and 'wealth creation' as buzzwords to manipulate readers' sympathies towards the Prime Minister's anti-European stance. In combination with the rhetorical import of some of the other words used in the passage, these buzzwords have the effect of showing the Prime Minister to be protecting these (allegedly) uncontroversially good things in the face of an unwarranted attack on them by the European Union, which the writer casts in a bad light by their use of the negative connotations of the buzzword 'bureaucrats'.

Scare quotes

This tactic is a means of influencing opinion against a view that one opposes. The speaker/writer takes key words in terms of which their opponent expresses their views and attempts to discredit those views by making them appear ridiculous or suspicious through the use of scare quotes. No reason is given for rejecting the view, we are simply manipulated into rejecting it because it has been made to appear ridiculous or suspicious. Suppose someone makes the following claim about people trying to settle in the UK for reasons of political asylum:

Almost all asylum seekers are economic migrants.

Now consider the effect of using scare quotes around the term 'asylum seekers' so that the claim becomes:

Almost all 'asylum seekers' are economic migrants.

As you can see, the addition of scare quotes has the same rhetorical effect as putting the phrase 'so-called' before the crucial term or phrase. The claim becomes much more explicitly negative in respect of its questioning of the legitimacy of people's claims for asylum. Indeed, it has virtually the same effect as the rhetorically explosive phrase 'bogus asylum seekers'.

This tactic can also be used more subtly but to similar effect. In such cases an opponent's opinion is made to seem dubious by placing scare quotes around words used to describe what would be perfectly normal, acceptable facts about them. This can turn a perfectly innocuous statement into one that casts doubt on someone's credibility. Compare

My opponent does of course have his reasons for what he believes to be right.

with the scare-quoted,

My opponent does of course have his 'reasons' for what he believes to be right.

While the first sentence suggests disagreement with the opponent's view, the use of scare quotes in the second lends it increased rhetorical power, not only expressing disagreement with the opinion in question, but also casting doubt on the legitimacy of the justification provided for those views.

Rhetorical ploys and fallacies

Care should be taken not to confuse the rhetorical ploy of scare quoting with the legitimate use of quotation marks to demarcate a direct citation of what someone has said or written. Sometimes we use quotation marks in this way even if the view we are quoting is one with which we disagree. For instance, in responding to a letter one has read in a newspaper, one might write,

> Ms Long is simply mistaken when she claims that 'men are better critical thinkers than women'.

Here the writer is simply using quotation marks to show that she is using the original writer's words verbatim. Sometimes we do use citations of what people actually say or write to rhetorical effect and we use quotation marks to demarcate their actual words, but this is not the same as the tactic of scare quoting. A good example of this latter technique is the use of someone else's words taken out of context in order to give rhetorical support to one's own opinions.

Trading on an equivocation

This ploy deliberately exploits the **ambiguity**, and in some cases the **vagueness**, of a word or phrase in the given context. Although nothing false is claimed, the speaker or writer manages to influence our actions or beliefs by misleading us. It is generally used when someone is attempting to persuade us of the benefits and virtues of their product or policy; hence it is common in advertising, particularly when an advertisement uses a superlative such as 'best', 'biggest', 'most successful', and so on. So imagine we come across the following advertisement:

> *Britburgers*: Britain's favourite hamburger restaurant!

To equivocate is misleadingly to use the same word in more than one sense; the ploy is to get us to interpret a message in a way that favours the product or view being advanced, when it is only under another interpretation that the message is true or well-founded. In this case, the word 'favourite' is ambiguous, having at least two possible meanings: It could mean (i) that Britburgers is the hamburger restaurant that most British people prefer or like best, or (ii) that Britburgers is the hamburger restaurant with the most customers in Britain (or possibly: that it's the one that sells the most hamburgers in Britain, or that has the most franchise outlets in Britain). Obviously the advertisers hope that we'll understand the message according to (i); this is much more favourable to the restaurant than merely having the most customers, since its market domination might be due to some factor other than actual customer preference.

Now we might realise that Britburgers would risk prosecution if they made this claim knowing it wasn't true, and thus conclude that Britburgers really is actively preferred by more British hamburger-eaters. But there's the rub with equivocation: the slogan could be true without the proposition it expresses giving us a good reason to expect Britburgers' burgers to be any good, for the truth expressed by the slogan might not be the one we take it to be. The slogan might be true only under inter-pretation (ii), and not under interpretation (i). In that case the advertisers could not rightly be convicted of having said something false; they can always claim that (ii) is all that was intended.

Smokescreen (changing the subject)

This is the tactic of avoiding discussion of an issue or acknowledgement of a point through diverting or distracting one's opponent from the issue at hand by addressing a different (possibly related) issue. The issue irrel-evant to the discussion thereby acts like a smokescreen by obscuring our view of the real issue. The more subtle the smokescreen, the more effec-tive at distracting the listener or reader it tends to be. Consider, for instance, a court case brought against a company that has allegedly been pursuing illegal monopolistic practices. When giving evidence, the company's Chief Executive Officer responds to counsel's questions about those practices, by launching into a speech about how the court case constitutes an attack on traditional values of hard work and innovation. Instead of addressing the issue at hand – whether or not the company was indeed involved in the alleged practices – the CEO is attempting to distract the court's attention from it by obscuring it in the rhetoric of their speech about hard work and innovation, a speech that is rhetori-cally powerful because hard work and innovation are the sorts of things that right-minded people are supposed to admire and want to protect.

The smokescreen tactic is very similar to the **red herring** fallacy, discussed in more detail below. The crucial difference is that while the smokescreen works via its rhetorical power, the red herring is still an attempt to persuade by argument: it gives **reasons** for accepting a claim, just not good reasons. When we are trying to distinguish the smoke-screen ploy from the red herring fallacy, we must first work out whether or not the writer or speaker is attempting to persuade by argument.

Fallacies

Strictly speaking, a fallacy is a mistake in reasoning. One commits a fallacy when the reasons advanced or accepted in support of a claim fail

to justify its acceptance. A fallacy can be committed either when one is deciding whether to accept a claim on the basis of a fallacious argument with which one has been presented, or when one is presenting the fallacious argument oneself.

A fallacious argument or inference is one in which there is an **inappropriate connection** between premises and conclusion. Almost all fallacies fall under one of the following two types:

- Formal fallacies. Sometimes the inappropriate connections are failures of **logical connection**, for example in the case of the fallacy of affirming the consequent; here the argument or inference is neither deductively valid nor inductively forceful, even where all implicit premises have been made explicit. It is simply a logical mistake.
- Substantive fallacies. Sometimes the inappropriate connections involve reliance on some very general unjustified assumptions or inferences. We need only make these **premises** explicit in order to see that they are false and unjustified. What distinguishes a fallacious argument of this kind from an ordinary unsound argument is that the implicit, false premise will be of a very general nature, having nothing specifically to do with the subject-matter of the argument. What this means will become apparent when we turn to examples.

The vast majority of fallacies that we encounter in everyday texts and speech are substantive fallacies, but some are formal fallacies, and a few are neither formal nor substantive. Arguments embodying formal or substantive fallacies are necessarily unsound, but not all fallacious arguments are actually unsound. Since fallacies are many and various in this way, it would be tedious and distracting to formulate a precise definition. We will simply take fallacies to be mistakes in reasoning that arise because of inappropriate connections between premises and conclusions; we will identify the most commonly encountered fallacies, and we will demonstrate different strategies for dealing with them.

It is important at the outset to note the general point, however, that a fallacious argument can have true or false premises: *Simply having false premises does not make an argument fallacious.* Nor does having *true* premises guarantee that an argument is *not* fallacious.

Furthermore, a proposition accepted on the basis of a fallacious argument may turn out to be true as a matter of actual fact. Suppose that someone reasons as follows:

There is blood on the candelabra. And if Colonel Mustard killed the victim with the candelabra, then there would be. So Colonel Mustard must be the murderer.

Laid out in standard form the argument is thus:

> P1) If Colonel Mustard killed the victim with the candelabra, then there is blood on the candelabra.
> P2) There is blood on the candelabra.
> _____
> C) Colonel Mustard killed the victim.

Now suppose that, in fact, the conclusion is true: Colonel Mustard is the murderer. Still, the reasoning that takes us to that conclusion is fallacious (the fallacy is the fallacy of **affirming the consequent**). It is not legitimate to infer on the basis of P1 and P2 that Colonel Mustard is the murderer. Even if both premises were true, it would be possible for the conclusion to be false: for example, someone else could have killed the victim with the candelabra. Knowledge of the truth of those premises, just by themselves, would not be sufficient to infer the conclusion, even if that conclusion were actually true. If you were to be fooled by such an argument, you would end up with a true belief, but for mistaken reasons.

Like linguistic phenomena and rhetorical ploys, the best way to become acquainted with the different types of fallacies considered in this chapter is to practise identifying and analysing them. As they are attempts to persuade by argument, you need to reconstruct them in standard form and then use techniques of argument analysis and assessment to demonstrate the ways in which they are fallacious.

And one last thing to bear in mind before we work through specific types of fallacies: as we have already mentioned, many types of fallacious argument are effective as rhetorical ploys. Someone might be aware that their argument commits a fallacy, but will use it to try to persuade us because they are aware of its rhetorical power: they are aware, that is, that it does tend to persuade people. Common examples include the use of *ad hominem*, majority belief and slippery slope fallacies to rhetorical effect.

Formal fallacies

The first group of fallacies we want to discuss here are strictly formal fallacies; these are patterns of argument whose reasoning makes purely **logical** mistakes. Each type of fallacy constitutes an **invalid argument** and once you are familiar with the patterns, the fallacies will be recognised by the presence of the particular invalid pattern.

Rhetorical ploys and fallacies

Affirming the consequent of a conditional

Or 'affirming the consequent' for short. This occurs when we argue from the conditional premise that if P (the antecedent), then Q (the consequent) together with the premise that Q to the conclusion that P, as in the following example:

> If mortgage rates go up, then house prices fall. House prices have fallen. Therefore mortgage rates have gone up.

Reconstruction of the argument reveals its invalidity:

> P1) If mortgage rates go up (P), house prices fall (Q).
> P2) House prices have fallen (Q).
> _____
> C) Mortgage rates have gone up (P).

If the premises were true, the conclusion would not have to be true. There are a variety of circumstances other than a rise in mortgage rates under which house prices might fall. So the fact that they have fallen would be insufficient for us to draw the conclusion that mortgage rates had gone up. That the inference is fallacious is obvious if we consider a true conditional such as 'If it is raining, then there are clouds'; it certainly does not follow, from this and the premise 'there are clouds', that it is raining. Notice that affirming the **antecedent** of a conditional premise does not make for an invalid argument:

> P1) If mortgage rates go up (P), house prices will fall (Q).
> P2) Mortgage rates have gone up (P).
> _____
> C) House prices will fall.

This is of course valid because the conditional gives one condition under which house prices will fall and P2 asserts that the condition is met; so it is legitimate to conclude that house prices will fall.

Denying the antecedent of a conditional

Or 'denying the antecedent' for short. This is the fallacy that occurs when we argue from a conditional premise (if P then Q) together with the

negation of its antecedent (not P) for the conclusion that the consequent is also negated (not Q). The invalidity of this pattern can be seen clearly using a version of the example above:

P1) If mortgage rates go up (P), house prices will fall (Q).
P2) Mortgage rates have not gone up (not P).

C)   House prices will not fall (not Q).

The invalidity here has similar grounds to that of the fallacy of affirming the consequent. There are conditions other than mortgage rates going up that could precipitate a fall in house prices, so the fact that mortgage rates have stayed the same (or decreased) does not give us sufficient grounds for concluding that house prices will not fall; other factors may lead to their falling. Or again: from 'If it is raining then there are clouds' and 'it is not raining', it does not follow that there are no clouds.

Fallacy of deriving ought from is

The Scottish philosopher, David Hume, famously argued that an ought cannot be derived from an is. The claim can be understood in two ways: first, as a claim about motivations to act or refrain from acting. The fact that something is thus-and-so is, argued Hume, insufficient as a reason for thinking that one ought to act in such-and-such a way. The fact that torturing animals causes them suffering is not, according to this argument, sufficient reason to refrain from torturing animals. An additional motivating force – a desire to avoid harm to animals – must also play a role in our motivation to act in such cases. Although an interesting and, indeed, controversial thesis, this way of interpreting Hume's claim is not the one that interests us here.[2] Rather we are interested in Hume's claim understood as the claim that a prescriptive conclusion cannot be validly derived from purely descriptive premises, such an inference is fallacious. Thus the fallacy of deriving an ought from an is occurs when a prescriptive conclusion – a conclusion making a claim about something that should or ought to be done or avoided or believed or not believed – is deduced solely on the basis of *descriptive* – fact-stating – premises. Inferences

---

2   For a detailed discussion of this debate see J.L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin, 1977).

from descriptive to prescriptive propositions are considered fallacious because the fact that something happens to be the case, or happens not to be the case, is insufficient grounds for concluding that it ought or ought not to be the case. If we want to make a valid argument for a prescriptive conclusion, we must always do so from premises at least one of which is prescriptive. Thus the following commits the fallacy of deriving ought from is:

> How can anyone claim that the monarchy should be abolished? The monarchy as we know it has been central to British life for nearly a thousand years, arguably longer.

Although it is not expressed as such, the conclusion of this argument is really a prescriptive one – that the British monarchy should be retained – but the conclusion is arrived at purely on the basis of the premise that it exists, and has done so for a long time. As such, the argument is plainly invalid:

P1) The British monarchy has existed for nearly a thousand years.

C) The British monarchy should be retained.

In this particular case we can make the argument valid by adding a prescriptive premise as follows:

P1) The British monarchy has existed for nearly a thousand years.
P2) Anything that has existed for nearly a thousand years should be retained.

C) The British monarchy should be retained.

But the argument is clearly unsound: poverty has existed for at least a thousand years, but it is surely false to claim that it should be retained.

A pattern of argument that occurs frequently in the context of reasoning about problem-solving and policy making can easily fall into

the trap of committing this fallacy. Suppose someone attempting to find a solution to the problem of global warming were to argue:

> Something must be done to save the planet from the environmental crisis that will be caused by global warming. I say we should ban the use of private vehicles, car exhaust is a major cause of global warming.

A reconstruction shows that the argument as given is fallacious:

P1) Something must be done to prevent global warming.
P2) Car exhaust is a major cause of global warming.

C) We should ban the use of private vehicles.

But the argument can be improved and made valid by the addition of a prescriptive premise that provides the link between the problem and the proposed course of action (P4). We also need to add a premise making explicit the claim that the proposed action is indeed a solution to the problem (P3):

P1) Something must be done to prevent global warming.
P2) Car exhaust is a major cause of global warming.
P3) Banning the use of private vehicles would reduce the effects of a major cause of global warming.
P4) We should do anything we can to prevent global warming.

C) We should ban the use of private vehicles.

Notice that although valid, the argument is unsound. The hard generalisation in P4) is open to counter-examples; suppose that we could prevent global warming by completely giving up any activity that involves burning fossil fuels, that would probably be too high a price to pay. We call arguments such as these practical arguments, and consider their analysis in more detail in Chapter 5.

Rhetorical ploys and fallacies

The base rate fallacy

This fallacy has already been mentioned in Chapter 3 (see p.82). It is committed when one mistakenly infers from a proposition that the majority of a minority group are thus-and-so together with a proposition that the minority of a larger majority are not thus-and-so that someone who is thus-and-so must be a member of the first minority. The inference is mistaken because the proportion of the majority who are thus-and-so may still be numerically larger than the majority of the minority. The fallacy commonly occurs when arguers resort to gender or racial stereotypes to make a point. Suppose that Ringons are a minority group outnumbered by non-Ringons by ten to one and someone argues:

> P1) Almost all Ringons have a criminal record.
> P2) Few non-Ringons have a criminal record.
> P3) Apex has a criminal record.
> _____
> C) (Probably) Apex is a Ringon.

The argument fails to be inductively forceful because even though there is such a high incidence of criminality amongst Ringons, there are so many more non-Ringons than Ringons that someone with a criminal record is still more likely to be a non-Ringon than they are to be a Ringon. Thus the premises, if true, would not make it likely that the conclusion is true.

Substantive fallacies

The first two substantive fallacies we consider involve an illegitimate inference from the commonality of a belief or an action to its acceptability. Like many, but not all of the fallacies we'll discuss, we can expose the fallacy by making explicit the hidden assumption that generates the illegitimate inference.

The fallacy of majority belief

This is the fallacy of concluding, on the basis of the fact that the majority believe a certain proposition, that the proposition is true. The following reasoning commits it:

> Of course the government must crack down on drug trafficking, after all that's what most reasonable, law-abiding people believe.

The only reason the arguer gives for a government crack-down is that most people think such a crack-down would be a good thing. We can see this more clearly by laying out the argument in standard form:

> P1) Most reasonable, law-abiding people believe that the government should crack down on drug trafficking.
> _____
> C) The government should crack down on drug trafficking.

So rendered, the argument is invalid. But we can make the cause of the faulty reasoning clearer by making the hidden assumption explicit thus:

> P1) Most reasonable, law-abiding people believe that the Government should crack down on drug trafficking.
> P2) Any belief shared by most reasonable, law-abiding people is true.
> _____
> C) The Government should crack down on drug trafficking.

Once we add the implicit premise, the argument becomes valid; but we can now see that the assumption underlying the argument, expressed by the hard generalisation P2, is false and that the argument is therefore unsound. Even if a majority of people do believe a proposition and even if they are of sound character, etc., their believing it is not sufficient to make it true. Imagine, for example, if most reasonable, law-abiding people believed that the Earth was flat, that would not establish that it is! More controversially, if a majority of reasonable, law-abiding people believe that capital punishment should be restored as a punishment for murder, the fact that it is a majority belief is not sufficient to make it true.[3] This is the principal strategy for dealing with substantive fallacies: by exposing the hidden assumption generating the fallacy, we reveal that the same (or a very similar) false proposition, usually a generalisation, figures in all instances of a particular fallacy.

The fallacy of majority belief often places us in an apparent dilemma when we have to take decisions on behalf of others. Suppose you are a

---

3   We consider the relationship between truth and belief in more detail in Chapter 7.

member of the jury in an important and high-profile trial that has received a great deal of pre-courtroom publicity. On that basis you know that most people believe that the defendant was the victim of a conspiracy on the part of the investigating police officers; nevertheless, the evidence presented in court gives you good reasons to believe that he is guilty. Thus you might be tempted into reasoning as follows:

> P1) Most people believe that the defendant was set up and is not guilty of the crime he is alleged to have committed.
> _____
> C) The defendant is innocent.

Such reasoning would commit the fallacy of majority belief and should be avoided as a basis for a juror's decision. Again we can make the faulty reasoning clear by exposing the hidden false assumption, which will be a very similar generalisation to that which figured in the previous example, thus:

> P1) Most people believe that the defendant was set up and is not guilty.
> P2) Any belief held by the majority is true.
> _____
> C The defendant is innocent.

So the argument is valid but not sound.

The fallacy of majority belief is similar to the rhetorical ploy of appeal to popularity, because both use the fact of something's popularity or commonality to attempt to persuade us to do or believe it. Sometimes, the very act of presenting an argument that commits the fallacy of majority belief is, at the same time, an example of the rhetorical ploy of appealing to popularity. Still these are distinct concepts: again, to call it a (substantive) fallacy is to say that the argument presented implicitly embodies a certain sort of unjustified assumption; to call it a rhetorical ploy is to say that it attempts causally to induce us to accept a certain belief by activating our social instincts, desires and fears. Some cases exemplify the fallacy without exemplifying the rhetorical ploy: someone might, for example, be quite sincere in thinking that if something is believed by the majority it must be true, and accept a conclusion on that basis. On the other hand, some cases exemplify the rhetorical ploy without exemplifying the fallacy: an

advertisement, for example, that shows someone being humiliated because they are the only person not drinking a certain brand of cola is clearly trying to influence us by way of our social instincts and feelings, not by presenting an argument.

Common practice

This is the tactic of attempting to persuade someone to do something they shouldn't do by giving them the justification that 'everyone does it'. The implication of this sham-reasoning is that if everyone does X, X must be acceptable. We often use this tactic to provide justification to ourselves for doing things we ought not to do. For example: 'It won't hurt to call in sick today, everyone does it once or twice a year.' Like the fallacy of majority belief, the common practice fallacy is driven by a false assumption concerning the connection between what is commonly believed or done and what it is morally, socially or rationally acceptable to believe or do. The following reasoning is an instance of the common practice fallacy:

> Of course it's OK to fiddle your expenses once in a while, everyone does.

This standard form rendition of the argument shows it to be invalid because there is no appropriate connection between the premise and the conclusion:

P1) Everyone fiddles their expenses occasionally.
———————————————————————
C)   It is acceptable to fiddle one's expenses occasionally.

Once again, exposing the false assumption produces the following argument, the unsoundness of which rests on the actual falsity of both P1 and P2:

P1) Everyone fiddles their expenses occasionally.
P2) Any act that everyone occasionally performs is acceptable.
———————————————————————
C)   Fiddling one's expenses is acceptable.

Rhetorical ploys and fallacies

Surely there are some people who have never cheated when making a claim for expenses from their employer; and we can think of actions, meting out violence to an innocent stranger, say, which, if performed occasionally by everyone would still be unacceptable. Even if the quantifier in P1 and P2 were qualified by 'almost', then although P1 might be true, it becomes even less plausible that P2 is true.

The next group of fallacies involve using alleged facts about the person(s) putting forward an argument as the basis for inferring a conclusion that their argument should not be accepted. In each case the fact about the person is irrelevant to the issue of whether or not we should accept their argument. As critical thinkers we are only interested in the argument, not in the person giving it. It is also helpful to remember this if we feel uncomfortable criticising arguments put forward by people we like, respect, fear or want to impress. If the argument is a good one, it makes no difference by whom it was given, likewise if it is a bad one.

Ad hominem

This fallacy (from the Latin, meaning 'to the man') can be committed in either of two ways: either by responding to someone's argument by making an attack upon the person, rather than addressing the argument itself, or by rejecting a claim because of disapproval of, or dislike for, the person who makes it. The following reasoning commits the *ad hominem* fallacy by citing a claim about a minister's speaking style as a reason for rejecting the legislation they propose:

> I don't see why we should accept the new Criminal Justice Bill when the minister presented it to us in such an imperious way.

An initial reconstruction displays the argument as invalid:

P1) The minister presented the Criminal Justice Bill in an imperious way.

C)   The Bill should be rejected.

Our strategy of exposing the crucial underlying assumption results in the following argument, which is unsound because P2 is false:

P1) The minister presented the Criminal Justice Bill in an imperious way.
P2) Any legislation presented by a person with an imperious manner should be rejected.

C) The Bill should be rejected.

All instances of the *ad hominem* fallacy will depend upon similar underlying general assumptions referring to certain characteristics or beliefs of arguers.

Ad hominem *circumstantial*

This is a sub-species of the *ad hominem* fallacy and occurs when someone's argument in favour of doing or believing something is discounted on the grounds that they would allegedly benefit from our doing or believing it. Someone would be committing the *ad hominem* **circumstantial** fallacy if they were to argue:

> Of course academics argue in favour of the proposed expansion of university education: the more aspiring graduates there are, the more job opportunities there are for people like them.

Notice that the conclusion of this argument is left implicit, but given the arguer's tone, it is reasonable to conclude that they do not agree with the expansion proposal. When reconstructing we need to add the conclusion for ourselves:

P1) Academics argue in favour of the proposed expansion of university education.
P2) They would benefit from such an expansion.

C) We should reject academics' arguments in favour of the proposal to expand university education.

The result, however, is an argument that is neither valid nor inductively forceful. Once again, we can see the true shape of the reasoning that constitutes this fallacy by exposing the hidden assumption thus:

121

Rhetorical ploys and fallacies

> P1) Academics argue in favour of the proposed expansion of university education.
> P2) They would benefit from such an expansion.
> P3) Whenever someone would benefit from something, we should reject their arguments in favour of it.
>
> ---
>
> C) We should reject academics' arguments in favour of the proposal to expand university education.

Note that P3 is really quite ludicrous: If it were true, then one could never hope to argue successfully for what one wants! Normally, that is exactly what we do, and there is nothing intrinsically illegitimate about it. So it is unreasonable to reject an argument because the arguer desires or would benefit from the truth of the conclusion. What matters is the strength of the reasons given for the claim, irrespective of the arguer's motives for making the claim. Reason places no strictures on arguing in favour of things from which one would benefit.

However, this is not to say that the issue of a speaker's or a writer's character is entirely irrelevant in matters of argument analysis and evaluation. A person's character and actions are certainly relevant to their **credibility**: the degree to which someone's having said something constitutes a reason to think it true. We should, for example, be on our guard against believing the claims of people whom we know to be dishonest. Even those who are not habitually dishonest are more likely to attempt deception if they are arguing for something that is strongly in their own interest. There is certainly a higher probability that someone who is not disinterested will deliberately resort to techniques of sham-reasoning in their attempts to persuade us of the truth of their claims; for they have more to lose if their arguments are not accepted. This probability increases with the degree to which the person is unscrupulous. In such cases we should check the reasoning carefully, and we should not take their having advanced a premise as a reason, in itself, to think it true. This is not, however, to say that we should positively assume their reasoning to be faulty or their premises false. That would be committing the *ad hominem* fallacy. We should never lose sight of the need to assess the argument on its own account. If the argument is found to be valid or inductively forceful, or the premises true, the character of the arguer is irrelevant.

There's one final observation to make concerning the relationship between people's characters and reputations and the strength of their arguments. This point does not directly concern a type of fallacious reasoning, but is relevant to considerations about the credibility of arguers and their arguments: Sometimes we are tempted to ignore or reject criticism of a

person's arguments because of the way that they have been labelled or treated by the media. The sorts of cases we have in mind are those in which someone has received racist, sexist or homophobic treatment in the media and this has been given as a reason for avoiding criticism of, or for accepting, their views. A recent example was provided by arguments that it was acceptable to allow the convicted rapist and former heavyweight boxing champion Mike Tyson into the UK on the grounds that much of the media coverage of the case was racist. It may well be the case that a good deal of what was written and said about Mr Tyson was racist and also that there were good reasons to grant him entry. It is crucial to recognise, however, that the existence of the racist discourse was irrelevant to the strength or otherwise of arguments that he should or shouldn't have been granted leave to enter the UK.

Tu quoque

In common with *ad hominem* fallacies, the *tu quoque* ('you too') fallacy occurs when we make unwarranted connections between someone's alleged lack of credibility and the strength of their argument. Here the alleged lack of credibility ensues specifically from their being hypocritical: an inconsistency between the arguer's actions and their claims. The fallacy is committed when we reject someone's claim that a behaviour or proposal should be refrained from or rejected on the grounds that they themselves practise that behaviour, or when we reject someone's claim that a behaviour or proposal should be adopted on the grounds that they fail to follow it themselves. This fallacy occurs frequently in discussions of the gaps between politicians' policy and private decisions. Suppose someone were to argue the following:

> Labour's education policy is a joke, how can we take their promises to improve state school standards seriously when even the Prime Minister sends his children to a grant-maintained school?

An initial reconstruction yields an invalid argument:

P1) The Labour Party promises to improve standards in state schools.
P2) The Prime Minister's own children do not attend a state school.

C) We should not take Labour's education policy seriously.

Rhetorical ploys and fallacies

Exposing the hidden false assumption demonstrates that the argument is unsound:

P1) The Labour Party promises to improve standards in state schools.

P2) The Prime Minister's own children do not attend a state school.

P3) Whenever someone's behaviour is inconsistent with their policies, we should not take those policies seriously.

C) We should not take Labour's education policy seriously.

The arguer has spotted the inconsistency between what the Prime Minister does and what he says, but that inconsistency does not render false the Prime Minister's views about the need to improve standards in state schools and it does not provide a reason to reject the policies in question. Whether it would be a good idea to improve standards in state schools does not depend on whether or not the Prime Minister sends his children to a state school. When we commit this fallacy we are going against the old adage: 'don't do as I do, do as I say'. We might well think that people ought to follow their own advice and principles, in fact in most cases it is irrational for them not to do so, but it is equally irrational of us to discount their arguments solely on the grounds that they themselves don't heed the conclusions of those arguments. Of course, the fact that someone in the public arena behaves inconsistently or hypocritically does undermine their credibility and may lead us to withdraw our trust and respect, but it is not in itself a reason to reject their arguments.

Appeal to authority

This fallacy also involves mistaken assumptions about the people mentioned by an argument. It is committed when an argument makes an unjustified appeal to an alleged authority. This can occur either because the authority appealed to is not in fact authoritative on the matter in hand or because there is good reason to doubt that the claimed authority is adequately informed of the facts of the matter. For example, the fallacy is committed when someone in power such as a prime minister or other political leader is unjustifiably invoked as an expert. For instance:

It is always better to drink white wine with fish. Tony Blair says so, he must know what he's talking about, he's the Prime Minister.

An initial reconstruction shows the argument to be valid:

P1) Tony Blair says that it is always better to drink white wine with fish.
P2) Tony Blair is prime minister.
P3) If someone is prime minister, then they must always be knowledgeable about all the subjects they talk about.

C)   It is always better to drink white wine with fish.

It is unsound, however, because of the falsity of the conditional in P3. Someone's being Prime Minister is not sufficient reason to think that they are knowledgeable about everything on which they express an opinion. That does not mean, of course, that all appeals to authority are fallacious; only those which are mistaken about someone's claim to be authoritative about the matter in hand. Prime ministers *can* claim authority on those matters on which they are best qualified to speak. If Tony Blair were to say that the best way to win an election is to make realistic promises, then one might accept that claim on the basis of his evident expertise in such matters.

The appeal to authority can also function as a rhetorical ploy, the pull of the alleged authority being used to lure us into accepting the proposition argued for. A common example of this use of the appeal to authority is the use of celebrity endorsement to sell things. While one might take Tiger Woods' endorsement of a brand of golf clubs as a good reason for buying that brand, one ought to be more wary about taking him to be an authority on makes of car, for example. Whereas, Tommy Hakinnen's endorsement of a car would be authoritative, but his recommendation of a particular brand of golf clubs would not.

The perfectionist fallacy

This fallacy occurs when we place excessive demands on an idea or a proposal and then reject it purely on the grounds that it will not completely solve a problem. Someone would commit the perfectionist fallacy if they were to argue or to accept the argument that:

Rhetorical ploys and fallacies

> The government should give up its plans to spend ten billion pounds on extra surgery in order to reduce hospital waiting lists. It's just not possible to get rid of waiting lists that way.

An initial reconstruction shows that the argument is invalid:

P1) The government's plan to spend ten billion pounds to reduce hospital waiting lists will not get rid of waiting lists completely.

C) The government's plan should be abandoned.

The assumption driving this instance of the perfectionist fallacy is that governments should only pursue plans that completely solve the problems they are intended to solve; more generally, the perfectionist assumption is that no measure aimed at solving or reducing a problem is justified unless it solves or reduces it completely. This is really quite silly if you think about it. It is obvious that many measures are intended to reduce a problem, not completely eradicate it, and are justified if they do reduce the problem sufficiently. For example fences are put around cow pastures to keep the cows in; once in a while a cow escapes, but no one would say that the occasional escape shows that the fences are unjustified.

When we include the perfectionist premise in the argument, in this case expressed as a conditional, we see that it is deductively unsound because the perfectionist premise is false:

P1) The government's plan to spend ten billion pounds to reduce hospital waiting lists will not get rid of waiting lists completely.
P2) If a government proposal will not completely solve the problem it is intended to solve, it should be abandoned.

C) The government's plan should be abandoned.

Conflation of morality with legality

This is the mistake of assuming that anything legal must be moral, or conversely, that anything illegal must be immoral. For something to be

legal within the boundaries of a given political entity (nation, city, district, the United Nations) is to say that there is no law in the statutes of that entity that prohibits it. However, the fact that something is legal does not automatically make it morally acceptable. For example, it is not illegal to cheat on your lover, or to be rude to a shy person at a party just out of cruelty, but these things are immoral. More gravely, in some countries, slavery, though obviously immoral, was legal until the nineteenth century. In others, even today, people may legally be denied human and civil rights because of their race, ethnicity, gender or religious belief.

In some cases the laws ought to be changed in order to reflect what is morally right. Indeed, it is tempting to argue that what is legal should include only what's moral. But to try to outlaw everything immoral would be unworkable, and indeed it would overstep what are generally considered to be the proper bounds of governmental authority. Some aspects of behaviour that we find morally unacceptable are not necessarily appropriate subjects for legislation – conduct in personal relationships, for example: we would not want it to be against the law to tell a lie, break a promise, make fun of someone, fail to show up for a date. Conversely, the fact that something is illegal does not automatically make it immoral. In some countries, Australia for example, it is illegal not to cast a vote in elections, but it is not obvious that there is a moral issue at stake here. Similarly, you are breaking the law when you park on a double yellow line, but it is arguable whether or not you are doing anything immoral. And sometimes laws that are meant to legislate against immorality are mistaken; throughout the ages there have been laws against all manner of things, thought by the authorities of the time to be immoral, that in retrospect we think are not immoral, and that never were – the worship or non-worship of certain gods, the education of women, the wearing of certain clothes. Indeed, to think that everything legal is moral is to nullify the possibility of criticising existing laws on moral grounds.

Consider this example:

> I don't see why people are so hard on Donald Mirving. After all he's done nothing wrong; in this country there's no law against denying the Holocaust.

The 'There's no law against it, so it's acceptable' argument is an extremely common instance of the fallacy of conflating legality and morality. The first reconstruction shows that the argument is invalid without the conflating assumption and the second that it is valid but unsound with it:

P1) Donald Mirving has denied the extent and facts of the Holocaust.
P2) It is not illegal to do this in the UK.

C) Donald Mirving's denial of the Holocaust was not morally wrong.

P1) Donald Mirving has denied the extent and facts of the Holocaust.
P2) It is not illegal to do this in the UK.
P3) Anything legal is moral.

C) Donald Mirving's denial of the Holocaust was not morally wrong.

So pointing out to someone whom you are trying to persuade that the thing in question is not against the law does not show that it is not wrong. Of course, many of the questions surrounding whether or not an issue is a moral one are controversial (the examples given above are by no means clear cut); they may require reflection and argument in their own right.

Weak analogy

Analogies are often interesting and may be **illustrative** of points one wishes to make, but arguing on the basis of analogy is often unsuccessful and often turns out to be fallacious either because the analogy is too weak to sustain the argument or because the analogy itself has not been argued for. (This makes the argument **question-begging**. We discuss the fallacy of begging the question later in this chapter.) In the case of the fallacy of weak analogy, it will be helpful first to see the form that instances usually take, and then to consider an example. The fallacy of weak analogy usually argues on the basis of a proposition that if one thing is similar to another in one respect, then it is similar in a further respect. This mistaken inference is based on the false assumption that if something is similar to another thing in one respect, it is similar in all respects. Hence we construct an argument of the following valid but unsound form:

P1) An object X is similar to an object Y in respect of characteristic A.
P2) Whenever an object X is similar to an object Y in one respect, it is similar in all respects.
P3) Y has characteristic B.

C) X has characteristic B.

An instance of this fallacy occurs frequently in debates about legislation to control the ownership and use of firearms. So let's take such an instance as our example:

> I don't see what all the fuss is about guns. Of course gun ownership shouldn't be prohibited, you can kill someone with a cricket bat, but no one proposes to ban ownership of cricket bats.

A reconstruction demonstrates that this argument takes the form characteristic of the fallacy of weak analogy except that to render the arguer's thinking fully, we have to add a further premise to the effect that things that are the same should always be treated the same.

P1) Guns are like cricket bats in that both can be used to kill people.
P2) Whenever an object X is similar to an object Y in one respect, it is similar in all respects.
P3) Objects that are similar to each other in all respects should be treated identically.
P4) We would not ban ownership of cricket bats.
_____
C) We should not ban ownership of guns.

The argument is unsound because it is obviously false to assume that similarity in one respect implies similarity in all respects. While it is of course true that cricket bats and guns have some shared similarities, those similarities are not sufficient for the analogy to hold. The dissimilarities – cricket bats have a different primary purpose; it is difficult to commit mass murder with a cricket bat, and so on – outweigh the similarities (indeed, one can kill a person with just about any solid, heavy object, such as a television; no one would argue that since we don't ban those potential murder weapons, we shouldn't ban firearms!).

The fact that arguments from analogy frequently turn out to be fallacious does not mean that they are universally unsuccessful. For an analogy to be effective in giving us a reason to accept a conclusion, an arguer must first present an argument for the claim that the objects that are allegedly analogous (in the case under consideration, cricket bats and guns) are sufficiently similar in the relevant respect. Once established, this conclusion would become the first premise of a subsequent argument for the claim that gun ownership should not be banned.

Rhetorical ploys and fallacies

Causal fallacies

These fallacies are committed when we make mistaken inferences about the cause(s) of something. Three types can be distinguished:

◆ *Post hoc ergo propter hoc* (from the Latin meaning 'after that, this, therefore this because of that);
◆ fallacy of mistaking correlation for cause; and
◆ inversion of cause and effect.

Post hoc ergo propter hoc

This fallacy occurs when we mistakenly infer that an event X caused an event Y merely on the basis that Y occurred after X. In the following example, the fallacy is committed twice. We can tell that the argument includes a causal claim because the proposition expressed by the conclusion states that one event – making a will – *makes* another event happen – living longer:

> Making a will makes you live longer. That's the conclusion reached by legacy specialists Live and Let Live, who compared the mortality figures of those who had made a will with those who hadn't. The average age of death for people who hadn't made a will was 72 years, 6 months. However, bequeath your personal possessions on paper and life expectancy shoots up to 80 years, 5 months. Want to live even longer? Leave some cash to charity: generous donors lasted until the ripe old age of 83 years.

As with previous cases, we can first represent causal fallacies as neither invalid nor inductively forceful arguments, then to demonstrate that they are driven by a false assumption, make the argument valid or inductively forceful but unsound:

P1) Generally people who make a will live longer than those who don't.

_____

C) (Probably) Making a will causes people to live longer.

Adding the implicit general assumption and making explicit a further implicit premise needed to make the argument inductively forceful gives

the following, which is inductively forceful but unsound due to the falsity of P2:

P1) Generally people who make a will live longer than those who don't.
P2) Whenever one event (Y) occurs after another event (X), Y is caused by X
P3) Living longer (not dying) occurs after making a will.

C) (Probably) Making a will causes people to live longer.

The *post hoc ergo propter hoc* fallacy is frequently committed in public discourse when arguers are attempting to persuade people of the merits of a policy or piece of legislation. Thus tougher sentencing policies might be fallaciously inferred to be the cause of a drop in the crime rate or a teachers' pay increase to be the cause of better examination marks in schools. Where the second event is alleged to be causally linked to the first solely on the basis that it occurred after the first, the argument is almost definitely fallacious. This is not to say that tougher sentencing, say, couldn't be a cause of a drop in the crime rate, it's just that we need to be given stronger reasons to accept the causal claim than simply the fact that one event occurred before the other.

Fallacy of mistaking correlation for cause

Whereas the fallacy of *post hoc ergo propter hoc* occurs because the temporal priority of one event over another is taken as sufficient to establish a causal relationship between those events, this fallacy is committed when the fact that one type of event or state of affairs is always or usually found in conjunction with another type, is mistakenly taken to be sufficient to establish that events or states of affairs of the one type cause the other. In summary, the fallacy is committed when a statistical correlation is assumed, without any further justification, to establish a causal relation. So, for instance, someone who argues,

You only have to look at the statistics to see that poverty is the obvious cause of educational under-achievement. Eighty per cent of those who leave school with no qualifications come from homes whose income is at least 50 per cent below the average.

commits the fallacy of mistaking correlation for cause. Our point here is not that it is completely mistaken to think that there is a link between poverty and under-achievement in school. Indeed, there is plenty of respectable research establishing a relationship between them. And certainly a statistical correlation is necessary for a causal relation: if there is no statistical correlation, then there cannot be a causal relation. What is crucial, though, is that we resist making the mistaken inferential leap from the fact of their statistical correlation to the alleged fact of a causal relationship between them. There may well be such a relationship, but correlation alone is not sufficient to conclude that the relationship is causal, as we see when we reconstruct the argument in standard form:

P1) There is a statistical correlation between students' under-achievement at school and poverty at home.

C) Poverty at home causes students' under-achievement at school.

Making explicit the hidden assumption, we derive the following valid but unsound argument:

P1) There is a statistical correlation between students' under-achievement at school and poverty at home.
P2) Whenever two phenomena X and Y are correlated, X is the cause of Y.

C) Poverty at home causes students' under-achievement at school.

That P2 is false can be made obvious by considering certain counter-examples. Creatures that wear clothes, for example, tend to be more intelligent than those that do not. These things are correlated. But wearing clothes does not cause intelligence. In cases where there is correlation between two or more events or phenomena and they are, in fact, causally related, it is often difficult to distinguish which is cause and which is effect. Poverty and ill health tend to go together, but which causes which? It is probable that in many cases they are mutually causally efficacious. We discuss the relation between correlation and cause in more detail in Chapter 5.

Inversion of cause and effect

Here one mistakenly infers that if X causes Y, an absence of X will prevent Y:

> Research carried out by a team of researchers shows that vitamin E may be the key to the secret of everlasting youth. The team believe that the vitamin may be at work normally in humans to prevent ageing. Experiments carried out on laboratory rats have enabled the scientists to discover that animals deprived of vitamin E seem to age faster and even become senile. Their spokesperson, Dr Young, said that although the ageing process remains mysterious, these experiments have demonstrated an 'interesting causal relationship' between vitamin E and the ageing process.

Even if we grant that it is generally warranted to make inferences about human physiology from that of rats, it is *not* warranted to infer from the fact that vitamin E deficiency seems to hasten ageing to the claim that vitamin E is the secret of everlasting youth. The proposition that a lack of something (vitamin E) causes X (hastened ageing) does not entail the proposition that the presence of that thing (vitamin E) causes the opposite of X (slowed ageing). We can use a reconstruction to expose the mistaken causal assumption driving this fallacy (P2):

P1) Vitamin E deprivation causes hastened ageing.
P2) If a lack of something X is the cause of a phenomenon Y, the presence of X will cause the opposite of Y.

———————————————————————————

C) Vitamin E causes the ageing process to slow.

One sometimes encounters the inversion fallacy and the fallacy of mistaking correlation for cause committed simultaneously. The following reasoning fallaciously infers a causal relation between eating certain foods and eating disorders from a correlation between those things (in this case the correlation is anecdotal rather than statistical), but *also* commits the inversion fallacy by mistakenly inverting that causal relationship:

> Contrary to what the healthy eating lobby dictates, traditional British cooking is good for you. Remember the good old days when we ate bacon, eggs, sausages, fried tomatoes, potatoes, toast and marmalade? And that was just for breakfast. Well in

> those days, people didn't succumb to anorexia or bulimia and the incidence of obesity was much lower. Seems the 'experts' have got it wrong again.

This argument seems to conclude not only that eating traditional British cooking protects people from anorexia, bulimia and obesity, but that not eating it causes those things. To expose the fallacies, we can reconstruct the entire argument as an extended argument:

P1) When we ate traditional foods (X) there was a significantly lower incidence of anorexia, bulimia and obesity (lack of Y).

P2) Whenever two phenomena (X) and (Y) are correlated, X is the cause of Y.

C1) Eating traditional foods (X) causes a low incidence of anorexia, etc. (lack of Y).

P3) Now we don't eat those foods (lack of X) and the incidence of anorexia, etc. is much higher (Y).

P4) Whenever X causes a lack of Y, a lack of X causes Y.

C2) Not eating traditional foods (lack of X) causes a higher incidence of anorexia, bulimia and obesity (Y).

The fallacy of mistaking correlation for cause is signalled by P2. The mistaken inversion is signalled by P4, which is false. Some causal relationships can be inverted, but it is certainly not true of causal relationships in general that they can be. For example, drinking milk causes us not to be thirsty. But it is not true that not drinking milk causes us to be thirsty!

The next two fallacies we consider are committed when we make unwarranted inferences from what is known, believed or proven.

Appeal to ignorance

This is the fallacy of concluding either that because a claim has not been proven it must be false (the negative form), or that because it has not been disproved it must be true (the positive form). It is often used when defending a belief in something that remains unproved such as astrology or the existence of a deity. The following commits the **negative form** of the fallacy:

> As no one's proved that UFOs exist, it's reasonable to assume
> that they don't.

whereas this equally fallacious argument commits the **positive form**:

> No one's managed to prove that UFOs don't exist, so we can
> reasonably conclude that they do.

As we see when we reconstruct the arguments, each fallacious and
unsound argument is driven either by the false assumption that absence
of proof means a proposition is false or that absence of disproof means
a proposition is true:

P1) No one has proved that UFOs exist.
P2) All unproven propositions are false.

———————————————————————————————

C) It is reasonable to conclude that UFOs do not exist.

P1) No one has proved that UFOs don't exist.
P2) All propositions that have not been disproved are true.

———————————————————————————————

C) It is reasonable to conclude that UFOs do exist.

Of course, in cases where efforts to prove something have been suffi-
ciently strenuous, it may be reasonable to infer the falsity of the proposi-
tion. For example, repeated efforts have been made, using sophisticated
scientific equipment, to find the Loch Ness monster (to prove the propo-
sition that Nessie exists); but to no avail. It is reasonable on that basis to
conclude, alas, that Nessie doesn't exist. But that is because we know that
if Nessie did exist, then she probably would have been detected by those
efforts. The mere fact that a proposition hasn't been proven, just by itself,
is no reason to think it false. Likewise, the mere fact that a proposition
hasn't been disproven, just by itself, is no reason to think it true.

'Proof' connotes certainty, and part of what is going on with this
fallacy is that people sometimes think that if a claim is not certain, then
it can reasonably be denied. But that is not how things are, as should be
reasonably clear from Chapter 3. Where we know we have an inductively
very forceful argument with true premises, then despite not having
perfect certainty, it would be unreasonable to deny the conclusion (with

certain exceptions, as explained in Chapter 6). Some claims and theories provide the most probable explanations of the phenomena they concern even though they remain neither proved nor disproved. The theory of natural selection is one such example. The arguments in favour of it have a great deal of inductive force, but as yet no one has managed to prove it.[4] Reasons why it should be considered the most plausible explanation of the evolution of species ought to be incorporated into relevant arguments; whether or not it is proved is not essential to the question of whether we ought to believe it.

Epistemic fallacy

This fallacy (from the Greek *episteme*, meaning knowledge) arises because of the tricky nature of knowledge and belief, and the difficulty of discerning from the third-personal point of view what someone believes or knows. It is committed when we make a fallacious inference from the fact that if someone believes that P then they must also believe that Q on the grounds that P and Q are about the same thing or person even though the way in which they refer to that thing or person is different. The following provides a simple instance of the epistemic fallacy:

> Chris believes that Tony Blair enjoys sky-diving. Tony Blair is the Prime Minister, so Chris believes that the Prime Minister likes sky-diving.

A reconstruction gives us:

> P1) Chris believes that Tony Blair enjoys sky-diving.
> P2) Tony Blair is the Prime Minister.
> _____
> C) Chris believes that the Prime Minister enjoys sky-diving.

The inference is unwarranted and the argument invalid because the arguer has assumed that in addition to having beliefs about Tony Blair's preferred leisure pursuits, Chris also knows that Tony Blair is the Prime Minister. But the arguer has no grounds for this assumption. Chris may only have a belief about Tony Blair but not know that Tony Blair is the Prime

---

4 We should note that although a theory may be **in principle** provable it may remain neither proven nor unproven indefinitely because there is insufficient evidence to prove/disprove it conclusively.

Minister. If so then C might be false. Another way of putting this is to say that Chris may not know that 'the Prime Minister' and 'Tony Blair' refer to the same person. Thus, if Chris is indeed ignorant of Tony Blair's being Prime Minister, the following argument, though valid, would be unsound due to the falsity of P3 and of C:

P1) Chris believes that Tony Blair enjoys sky-diving.
P2) Tony Blair is the Prime Minister.
P3) Chris knows that Tony Blair is the Prime Minister.

C)   Chris believes that the Prime Minister enjoys sky-diving.

It is important to note that similar inferences made in different contexts are warranted and the arguments containing them valid. Consider the following:

P1) The Prime Minister is a world champion darts player.
P2) Tony Blair is Prime Minister.

C)   Tony Blair is a world champion darts player.

Inferences such as this are sanctioned by an apparently indubitable logical principle known as **Leibniz's Law** (after the seventeenth-century German philosopher and mathematician, Gottfried Leibniz). This law holds that if one thing is the same identical thing as another, then what is true of one must be true of the other. For example, if Superman has blond hair and Superman and Clark Kent are the same person, then Clark Kent must have blond hair. Sentences about people's beliefs or knowledge, such as our example about Tony Blair and sky-diving, are exceptions to Leibniz's Law. If Chris believes that X is thus-and-so, even if X and Y are the same thing, it does not follow that Chris believes Y to be thus-and-so because we do not know whether or not Chris knows that X and Y are the same thing. The inference would only be warranted if the arguer knew that Chris knew this. How to make sense of these sorts of cases is a famous philosophical puzzle, but we need not let that worry us here.[5]

---

5   The most famous discussion is 'On sense and reference', by Gottlob Frege, reprinted in *Meaning and Reference*, edited by A.W. Moore (Oxford: Oxford University Press, 1993).

Rhetorical ploys and fallacies

The epistemic fallacy is often used knowingly to discredit someone's opinion. For example:

> Mr Smith believes that the cultivation and use of cannabis should remain a criminal offence in this country. But cannabis is the most effective anti-nausea drug for chemotherapy patients. So Mr Smith believes that it should remain a criminal offence to produce or use the most effective anti-nausea drug for chemotherapy patients.

An initial reconstruction of the argument gives us the following valid argument:

P1) Mr Smith believes that it should be a criminal offence to produce or to use cannabis.
P2) Cannabis is the most effective anti-nausea drug for chemotherapy patients.

___

C) Mr Smith believes that it should be a criminal offence to produce or to use the most effective anti-nausea drug for chemotherapy patients.

If we add a hidden premise of the same form as P3 in the previous example, we see that we cannot conclude that the argument is sound because unless we have some grounds for saying that P3 (and hence C) are true:

P1) Mr Smith believes that it should be a criminal offence to produce or to use cannabis.
P2) Cannabis is the most effective anti-nausea drug for chemotherapy patients.
P3) Mr Smith knows that the most effective anti-nausea drug for chemotherapy patients is cannabis.

___

C) Mr Smith believes that it should be a criminal offence to produce or to use the most effective anti-nausea drug for chemotherapy patients.

The inference drawn here is unwarranted if we don't know that Mr Smith is aware of the anti-nausea properties of cannabis. It is possible that Smith is unaware that cannabis is the best anti-nausea drug for

chemotherapy patients. Indeed, it might well be that Smith believes that the best remedy for nausea ought to be made available to chemotherapy patients and so long as he is genuinely ignorant that cannabis *is* the best such remedy, he would not be inconsistent in making this claim. The danger of epistemic fallacies, then, is that they may attribute beliefs to persons that they do not really hold.

Notice that these examples turn on verbs such as 'knows', 'believes', 'wants'. Philosophers and linguists call verbs such as these **propositional attitude verbs**. If we reflect upon how these verbs are used, we see that we say someone **believes that . . .**, where the blank is filled by the expression of some proposition or other. A propositional attitude, then, expresses the fact that someone holds some attitude towards a specific proposition. Smith **believes** that such and such is the case; Jones **wants** such and such to happen; Brown **knows** that such and such is the case. Other examples are 'desires', 'hopes', 'prays', 'wishes'.

Further fallacies

We said early in our discussion of fallacies that almost all are either formal or substantive fallacies. All such fallacies make for unsound arguments; they are either irremediably invalid, or depend on some very general but false implicit assumption. We turn now, however, to a different group of fallacies. These are labelled as fallacies, but not every instance of them will be invalid, or inductively unforceful, or even unsound. However, they are all **poor techniques at argument**; they should be criticised when we analyse arguments, and avoided in our own attempts to persuade by argument. Many of them, however, are useful for non-rational persuasion: they are frequently used to avoid engagement with an opponent, or to trump an opponent in the knowledge that their premises do not actually give good reason to accept their conclusions. In many cases they do have persuasive power.

Although reconstruction will be helpful in analysing instances of these fallacies, they cannot be exposed by making explicit a false assumption that drives *all* instances of the fallacy in question. This is because there is no single false assumption (expressed as a generalisation or a conditional) that underlies *all* instances of each of these fallacies. So while we should, in order to expose their fallacious reasoning, continue the practice of reconstructing arguments that we suspect of committing these fallacies, it is not so easy to give a straightforward method for detecting these fallacies.

Rhetorical ploys and fallacies

Equivocation

The rhetorical ploy of **trading** on an equivocation is the ploy whereby we deliberately use a word or form of words with the intention to confuse the audience; one hopes that the audience will conflate the two or more possible interpretations. A single unsupported claim, rather than an argument, may be the instrument of the ploy. To fall prey to the fallacy of equivocation, by contrast, is to fail to notice an ambiguity, **thereby accepting the conclusion of an argument**, when one should not have. Silly but clear examples are easy to come by; for example: 'In the philosophy department, someone broke one of the chair's legs; therefore one of the philosophy department's professors has a broken leg' (equivocation on the word 'chair'). Such a case is simple and amusing but no one would actually be taken in by it. In the more interesting cases, explaining the fallacy can be a subtle conceptual task. For example:

> Some conservatives claim moral universal truths; they claim that throughout history, in all times and places, people fundamentally have the same rights. This displays a lamentable ignorance of history, and – characteristically of conservatives – of other cultures. It is a plain fact that at other times in history, and in other parts of the world today, human beings do not have the same rights. In some countries, for example, a man has the right forcibly to confine his wife to the home if he sees fit; not so in our culture. The conservative claim of universal rights is plainly false.

The arguer wishes to conclude that, contrary to certain conservatives who believe in universal moral truths, whether or not a human being possesses a given right depends on what culture they are in. Thus ignoring some irrelevant material, we may reconstruct it as a very simple argument:

P1) In some countries, men have the right to confine their wives forcibly; in other countries they do not.

C1) It is not the case that human beings have the same rights in all places and at all times.

C2) The conservative claim – that throughout history, in all times and places, people fundamentally have the same rights – is false.

The argument equivocates on the word 'right', however. Both senses are established items in our language, but they are close together in meaning. In one sense of the word, to possess a 'right' is to be allowed, by the culture or other social environment one is in (often, but not always, this is a system of laws), to perform a certain action. Call this the 'conventional' sense of the word. In the other sense, to possess a 'right' is to be such that one ought, whatever culture or other social environment one is in, to be allowed to perform a certain action – even if one is not in fact allowed to. Call this the 'philosophical' sense of the word. Thus one may possess rights in the philosophical sense that are not rights in the conventional sense. The trouble with the argument is that it uses both senses: if we keep to the conventional sense of the word, P1 is true, C1 is true, and the inference from P1 to C1 is valid. The inference to C2 would be valid if the conservative claim were intended in the conventional sense. But the conservative claim, no doubt, was that rights are invariant in the philosophical sense of the word. In that case C2 cannot be inferred from C1; the inference would be no better than the silly one about the broken chair.

Red herring

So named after the practice of dragging a smelly, salt-cured (and therefore reddish) herring across the trail of an animal tracked by dogs. The red herring fallacy is used as a technique to throw someone off the scent of one's argument by distracting them with an irrelevance. The rhetorical ploy of the smokescreen constitutes a similar tactic. However, where an irrelevant premise(s) is given as a *reason* for accepting the conclusion being advanced, the red herring fallacy is committed. For example:

> The judge should rule against the plaintiff's charge of
> sexual harassment against the president. The President is
> very popular, and presides over an extremely healthy
> economy.

The arguer seems to advance the President's political success as a reason to rule against the charge of sexual harassment. If we make the reasonable assumption that the judge should rule strictly on the basis of the president's guilt or innocence, then the President's political success is utterly irrelevant. Reconstructed, the argument looks like this:

Rhetorical ploys and fallacies

P1) The President is very popular, and presides over an extremely healthy economy.

P2) If the President is very popular, and presides over an extremely healthy economy, then the judge should rule against the plaintiff's charge of sexual harassment.

C) The court should rule against the plaintiff's charge of sexual harassment.

In general, the red herring fallacy is that of inferring a conclusion from a premise that is strictly irrelevant to it, but in a way that has the potential to fool the audience into accepting the inference. Normally this is accomplished by a premise that tends to instil some sort of positive attitude towards the conclusion. In this case, the premise is intended to make the audience feel supportive towards the President, thus unreceptive to the idea that he should be convicted of misconduct.

Note that, although red herring arguments can easily be represented as valid, red herring is not a substantive fallacy. P2 is obviously false, but our ability to recognise this depends on our knowledge of what is, and what is not, relevant to the establishment of guilt in a court of law. More generally, what is and what is not relevant to a conclusion will depend on the conclusion's particular subject-matter. So there is not going to be one characteristic premise that red herring fallacies assume, in the way that there is, for example, in the case of inverting cause and effect. So red herring, according to our categories, is not a substantive fallacy.

It is worth re-emphasising, finally, that to say that someone has been taken in by a red herring fallacy is to say that they have been fooled. Unlike most other fallacies, the ability to recognise red herring varies depending on our knowledge of the subject-matter of the argument. But if X honestly believes, for example, that cancer is always caused by thinking morally bad thoughts, then, although having developed cancer is irrelevant to the question of the moral character of their thoughts, X does not commit red herring if X infers, from the fact Y has cancer, that Y must have been thinking bad thoughts. X is just badly informed. The point of distinguishing red herring as a fallacy of irrelevance is to single out the cases where one is fooled by an irrelevance, where one ought to have known better. Every minimally educated person knows, for example that guilt or innocence in a court of law is properly established only by the preponderance of evidence; because of this, one who advances or accepts the argument given above has been fooled by an irrelevance, and has thus committed red herring.

Slippery slope

This fallacy occurs when an arguer wrongly assumes that to permit or forbid a course of action will inevitably lead to the occurrence of further related and undesirable events, without providing good reasons to suppose that the further events will indeed inevitably follow; and thus to allow the first is to tread on a slippery slope down which we will slide to the other events. Since its rhetorical power is derived from fear or dislike of the undesirable events, it is from a rhetorical point of view closely related to the appeal to fear. Slippery slope arguments are sometimes used to justify particularly harsh laws or penal sentences and occur frequently in debates about the liberalisation or toughening of laws or constraints on behaviour, as in the following example about the decriminalisation of cannabis use:

> The decriminalisation of cannabis would be just the start. It would lead to a downward spiral into widespread abuse of harder drugs like heroin and cocaine.

The implicit conclusion is that cannabis should not be decriminalised. The only explicit premise is that if cannabis were decriminalised then the use of hard drugs would increase. So an initial reconstruction represents the argument as invalid:

P1) If cannabis were to be decriminalised, the use of hard drugs would increase.

C) Cannabis should not be decriminalised.

Notice that as it stands the argument also commits the fallacy of deriving ought from is. To correct this, we need to add a premise to make good the connection between the non-prescriptive premise and the prescriptive conclusion, thus ending up with the following argument:

P1) If cannabis use were decriminalised, the use of hard drugs would increase.
P2) Anything that leads to increased use of harder drugs should be avoided.

C) Cannabis should not be decriminalised.

Rhetorical ploys and fallacies

The immediate problem is that we have not been given a reason to think that P1 is true; that is, no reason to think that decriminalisation of cannabis will unavoidably be the beginning of a slippery slope to an increase in the use of hard drugs. Of course, some slopes really are slippery; even in this case, it might be possible to give such reasons and they might form part of an extended argument for the same conclusion. But as it stands the argument remains fallacious, because the arguer has not given a reason for supposing that it is inevitable that allowing the first event will precipitate a slide into even worse events. (This form of argument is sometimes called **floodgates** – the arguer alleges without evidence that allowing X will inevitably open the floodgates to Y and Z.)

Straw man

This is the fallacy that occurs when an arguer ignores their opponent's real position on an issue and sets up a weaker version of that position by misrepresentation, exaggeration, distortion or simplification. This makes it easier to defeat; thereby creating the impression that the real argument has been refuted. The straw man argument, like the straw man himself, is easier to knock down than the real thing. Suppose that Jones is an advocate of the legalisation of voluntary euthanasia; that is Jones believes that terminally ill patients should have the legal right to choose to have their life ended if their suffering has greatly diminished their quality of life, and doctors agree that the patient's mental state is sufficiently sound to make the decision rationally. Smith, Jones' opponent, responds as follows:

> How can you support giving doctors the right to end a person's life just because *they* decide that the person's life is no longer worth living; no one should have that power over another person's life, and doctors should not kill patients.

According to Smith, Jones advocates that doctors should unilaterally have the power to end a patient's life, if they think that the patient's life is not worth living. That would be a very controversial position. But it is not Jones' position. Jones' position is that patients should have the choice of euthanasia, so long as that choice is approved by doctors. Of course the doctor administers the lethal drug, but only at the behest of the patient, as Jones envisages things. Smith thus fails to engage with Jones' **actual** position and instead misrepresents it as a more extreme and therefore weaker position that (as far as we know) Jones does not advocate.

Begging the question

An argument commits the fallacy of begging the question when the truth of its conclusion is assumed by one or more of its premises, and the truth of those premise depends for its justification on the truth of the conclusion. Thus the premises ask the audience to grant the conclusion even before the argument is given. Contrary to the way in which the phrase is sometimes used in ordinary discourse, 'begging the question' does not mean raising the question without offering an argument. By way of example, imagine the following scenario:

> Three thieves pull off a successful heist and steal four diamonds, but they can't decide how to divide their haul. Eventually the first thief says: 'I should get two diamonds, and you two should get one each, because I'm the leader.' The second thief says: 'Wait a minute, who says you're the leader?' The first thief replies: 'I must be the leader, I'm getting the largest share of the haul.'

A reconstruction enables us to see clearly the way in which the first thief's reasoning begs the question. Initially thief number one appears to argue from the proposition that he's the leader together with the implicit premise that leaders should always get the biggest haul to the conclusion that he should get the biggest share:

P1) I'm the leader of the gang.
P2) Gang leaders should always get the biggest share of their gang's haul.

C) I should get the biggest share of the haul.

But when requested by thief number two to justify P1 he argues as follows:

P1) I'm getting the biggest share of the haul.
P2) Whoever gets the biggest share of the haul is the leader.

C1) I must be the leader of the gang.

Rhetorical ploys and fallacies

Once we put the two steps of the argument together to form a complex argument, we see that it begs the question of who should get the biggest share:

P1) I'm getting the biggest share of the haul.
P2) Whoever receives the biggest share of the haul is the leader.

C1) I must be the leader of the gang.

P3) Gang leaders always receive the biggest share of their gang's haul.

C2) I'm getting the biggest share of the haul.

Thief number one is guilty of begging the question (in addition to armed robbery) because P1, which he uses to reach C2, expresses the same proposition as that expressed by C2. So the conclusion – that thief number one gets the biggest share – is already assumed by the premises. Notice that each inference stage passes our test for validity, if the premises were true, the conclusion would have to be true, but the argument is a clear instance of the fallacy of begging the question of who should get the biggest share of the proceeds of the robbery.

An argument's premise(s) and conclusion need not express a proposition in precisely the same way (as they do in the previous example) in order for it to count as an instance of begging the question. It is sufficient that the premise be a version of, or relies upon, the claim made by the conclusion. If someone were to argue from the premise that newspaper editors claim that their publications are better than any other medium for finding out about events overseas, to the conclusion that newspapers are the best source of international news, their argument would commit the fallacy of begging the question. For obvious reasons, the fallacy of begging the question is often referred to as 'circular reasoning'.

False dilemma

This is the fallacy of limiting consideration of positions on an issue to fewer alternatives than should be considered. Typically, the arguer pretends that there are only two options, when in fact there are more: the arguer sets up a dilemma where none really exists by misrepresenting the possible positions on an issue so that there appears to be a straight choice between their own and its opposite. The fallacy is committed by a politician who argues as follows:

> There is a tough choice facing the government and the nation:
> either we cut taxes and increase everyone's spending power,
> thereby providing much needed stimulation to the economy, or
> we increase spending on health and education. It is impossible
> to do both; and without a tax cut the economy will remain
> weak. So increased spending on health and education will have
> to wait.

As a reconstruction shows, the argument is valid, but the fallacy is driven
by the false assumption that cutting taxes and increasing public spending
on health and education preclude each other.

P1) We should stimulate the economy.
P2) The only way to stimulate the economy is by cutting taxes.

C1) We should cut taxes.

P3) We cannot both cut taxes and increase public spending
    on health and education.

C2) We should not increase public spending on health and
    education.

The assumption P3 is not true because (a) if a tax cut does improve
the economy, then even if the government takes a smaller proportion of
the Gross National Product it might take in more in absolute terms
(because the GNP will be larger); (b) expenditure on health and educa-
tion might be increased by diverting government funds from other areas
such as defence; (c) the government might be able to sustain a period of
decreased tax revenues and increased overall expenditure by increasing
its debt (or decreasing its surplus, as the case may be).

The fallacy of false dilemma is often used to make the false assump-
tion that if someone does not agree with X, they must be anti-X, whereas
they may hold some intermediate position or be undecided. For example,
suppose someone asks you if you are in favour of positive discrimina-
tion towards under-represented groups in the award of promotion to
higher grades of a profession (affirmative action). You reply that you are
not and, employing a fallacious inference, they accuse you of being against
affirmative action, when in fact you may just be undecided, neither for
nor against it. Or, worse, perhaps they retort that you support discrim-
ination against under-represented groups, and therefore are racist, ageist
and misogynist (among other things).

Rhetorical ploys and fallacies

CHAPTER SUMMARY

Rhetorical ploys and **fallacies** are both instances of **sham-reasoning**. While rhetorical ploys seek to persuade by non-argumentative means, fallacies are argumentative attempts to persuade that embody some characteristic type of confusion or mistaken assumption. Many fallacious arguments may function at the very same time as an effective rhetorical ploy, thus causing the audience not to notice the fallacy. Insofar as we aim to be rational, and to appeal to the rationality of others, we should avoid using rhetorical ploys and fallacies in our own attempts to persuade and should take care not to be persuaded by others' rhetorical or fallacious attempts to persuade us to do and believe things. The best way of doing so is to familiarise ourselves with various common rhetorical ploys and fallacious forms of argument.

Many rhetorical ploys are **appeals to specific feelings or desires**; these include the appeals to **novelty**, **popularity**, **compassion**, **pity**, **guilt**, **fear**, **cuteness**, **sexiness**, **hipness**, **coolness**, **wealth**, **power** and many others. Typically a position or consumer item is represented in association with some object of the specific feeling or desire, in order that the feeling or desire should be directed upon the position or consumer item. The **direct attack** involves the bald assertion of a position or command; the **hard sell** is the direct attack repeated. The use of **buzzwords** is the use of words with high emotive or otherwise rhetorical charge to manipulate the passions of an audience. **Scare quotes** are used mockingly to make an opponent's position or other phenomenon look ridiculous or dubious. **Trading on an equivocation** occurs when someone knowingly makes an ambiguous statement that may be true when interpreted a certain way, which, when interpreted another way, may be false, but is more favourable to the position being advanced or to the product being advertised. **Smokescreen** occurs when one talks about some highly controversial, compelling or otherwise arresting issue or object, in an effort to divert the audience momentarily from the issue under discussion. A successful smokescreen causes the audience to overlook the fact that the issue has not been addressed.

Fallacies can be grouped together according to certain shared features. **Formal** fallacies are simply mistaken inferences – inferences of certain characteristic kinds that are often mistakenly thought to be valid or inductively forceful. These include: the fallacies of **affirming the consequent** (of a conditional proposition) and **denying the antecedent** (of a conditional proposition);

the **fallacy of deriving ought from is**, which is committed by any argument that attempts to move from solely descriptive premises to a prescriptive conclusion, and **the base rate fallacy**.

Substantive fallacies are committed by arguments that tacitly assume some very general principle of a characteristic kind that it may be tempting to rely upon, but which is false, and which can easily be seen to be false the moment it is brought to light. The fallacies of **majority belief** and **common practice** make illegitimate inferences from the commonality of a belief or an action to its acceptability. The *ad hominem* fallacy and the *tu quoque* fallacy use facts about the person(s) putting forward a position as grounds for rejecting it. The **appeal to authority** makes a mistaken appeal to the opinion of someone who is not qualified (or under-qualified) on the matter in hand. The **perfectionist fallacy** is committed when excessive demands are placed on an idea or proposal; the fallacy of **conflating morality with legality** occurs when we mistakenly assume that anything that is legal must also be moral or that anything that is illegal must be immoral. The fallacy of **weak analogy** arises when an argument employs an unsustainable or an unjustified analogy. **Causal fallacies** are committed when we make mistaken inferences about the cause of a phenomenon or an event. There are **three types** of causal fallacy: *post hoc ergo propter hoc* – the fallacy of assuming that the temporal priority of X over Y makes X the cause of Y; **fallacy of mistaking correlation for cause** – falsely assuming that the simultaneous occurrence of X and Y makes one the cause of the other, and the **fallacy of causal inversion** – the mistaken inference that if X causes Y, an absence of X will prevent Y. **Epistemic fallacies** and the fallacy of **appeal to ignorance** occur because of unwarranted inferences from what is known, believed or proved to additional knowledge, beliefs or proof of which the arguer has no independent evidence.

Substantive fallacies can be exposed by careful argument reconstruction. An initial reconstruction will demonstrate that the argument, as stated, is either invalid or inductively unforceful. A second reconstruction can be used to reveal the false assumption that drives the fallacious reasoning, thereby demonstrating that the (amended) argument is unsound. For each substantive fallacy, all instances of the fallacy will be driven by the same or a very similar assumption.

Some fallacies are neither formal nor substantive. Not all instances of the fallacies considered in the final grouping of

Rhetorical ploys and fallacies

this chapter are invalid, non-inductively forceful or unsound even when carefully reconstructed. A deductively sound, hence valid argument may beg the question, for example. While the method of exposing hidden assumptions will prove helpful for some instances, there is no single false assumption that underlies all instances of each of these types of fallacies.

The **red herring** fallacy occurs when irrelevant premises are given as a reason for accepting a conclusion. A **slippery slope** fallacy is committed when an arguer assumes without justification that to permit or forbid a course of action will inevitably cause a chain of undesirable events. The **straw man** is deliberately set up as a target that will be easier to defeat than an opponent's real argument. An argument commits the fallacy of **begging the question** when the truth of its conclusion is assumed by its premise(s). The fallacy of **false dilemma** is committed when an argument limits consideration of positions on an issue to two mutually exclusive ones, thereby setting up an apparent dilemma, when there are other positions that could be considered.

EXERCISES   **1**   Name the following rhetorical ploys.

Example:
If we allow the development of genetically engineered foodstuffs, we are leaving our children and our children's children open to the threat of genetic mutation and environmental disaster.
Appeal to fear/scare tactics

a   The 'relationships' of homosexuals cannot be compared to that of marriage: marriage is something that involves a man and a woman united through the word of God.

b   Take a look at the latest hatchback from Fraud: the new *Ergo*.

c   Successfully applied knowledge management is now impacting on organisations in terms of both direct competitive advantage, and in building a learning culture that identifies itself through a shared vision and common purpose.

d   More pet owners feed their puppies *first choice* puppy food than any other. If you want to give the little fella a head start, shouldn't you, too?

e   These people say we have a moral duty to honour our United Nations commitments and provide refugees with a safe haven. I say to you: look at all these people arriving here and then claiming state benefits; they are sucking this country dry.

f   Fatbusters – the most successful diet programme on the market today!

g   Imagine yourself, or worse, your daughter, alone beside your broken down car on a remote country road in the middle of the night. Few people pass by and no one stops to help. Don't get caught like that – don't get caught without a Phonecom mobile phone.

h   By scheduling the broadcast at 1 a.m., the TV company is committing an act of censorship.

i   Hair unmanageable and dull? Try *Goldie Glow* from Bella. The latest in hair management technology.

j   **Smith**: I think it's extremely hypocritical of Mr Big to sit in judgement over those kids who were caught using drugs in school, when he himself has been caught breaking the law by drinking and driving. **Jones**: But Mr Big is a successful businessman who gives up his valuable time to serve our community.

---

2   Name the fallacy committed by the following arguments. If no fallacy is committed, write 'N/F'.

Example:
No one has ever been able to prove the existence of extrasensory perception. We must therefore conclude that extrasensory perception is a myth.
Appeal to ignorance

a   Our employees have asked us to provide lounge areas where they can spend their breaks. This request will have to be refused. If we give them lounge areas, next they'll be asking for a swimming pool and sauna. Then it will be tennis courts, football pitches, fitness centres . . .

b   A few minutes after the minister made his speech to the City of London, a devastating explosion occurred. For the safety of the people who live and work in the City, it is imperative that the minister makes no more speeches here.

Rhetorical ploys and fallacies

c Publishing these vile criminals' names and addresses in our paper was definitely the right thing to do. We have been inundated with calls and faxes from readers who support our stance.

d If she doesn't finish her assignment on time, she will fail the course. She has failed the course so obviously she didn't finish her assignment on time.

e We have a simple choice between developing genetically manipulated crops or continuing to stand by while thousands of people in the Third World die of starvation and malnutrition.

f As my client has pointed out, tax avoidance is not illegal, so she has done nothing immoral.

g The Japanese diet is low in dairy products and certain cancers have a very low incidence in Japan. So if you want to avoid cancer, give up drinking milk and eating milk products.

h Teachers say that their job is becoming harder and harder and that they deserve more pay. But the government should ignore them, they're just a bunch of whingeing liberals.

i It's all very well for the caring middle classes to say that we should provide more aid to the Third World, but remember that most of them are preaching at us from around their well-provided dining tables.

j Given the evidence that so many students are cheating in examinations, examinations should be abolished in favour of assessment by course-work.

---

3 (i) Name the fallacy committed by each of the following arguments and (ii) reconstruct the argument to demonstrate its fallaciousness.

Example:
Democracy is the best system of government. Most people in the world believe in the superiority of democratic systems.
Fallacy of majority belief

P1) Most people believe in democratic systems.
P2) Any belief held by the majority is true.

C) Democracy is the best system of government.

a   Our lecturers are always extolling the virtues of critical thinking, but they would say that wouldn't they? They only keep their jobs if they've got students to teach.

b   It's not illegal for me to exaggerate my skills on my curriculum vitae, so it can't be immoral either.

c   The opposition criticises the government for spending money on the arts, but when they were in govt. they increased the arts budget by two hundred per cent.

d   Smoking causes lung cancer so people who do not smoke will not suffer lung cancer.

e   Jo believes that Chancellor of the Exchequer is doing a poor job. Thus Jo believes that Gordon Brown is doing a poor job.

f   Whenever the cherry trees blossom, the weather begins to get warmer. So cherry blossom causes the weather to get warmer.

g   David Beckham is the best mid-fielder ever to play for England and he says that Brit-Burgers are his favourite. They must be the best burgers.

h   We shouldn't allow children to watch *Pokemon*, because pretty soon they will be collecting farm animals and staging contests to see who has the strongest pig. It's only a short step from there to complete disrespect for animal life.

i   There's nothing wrong with getting drunk once in a while, everyone I know is partial to a few beers now and again.

j   There is no point in implementing harsh penalties for drunk drivers because there will always be some people who are going to drive drunk no matter what penalties are put in place.

k   If a car breaks down on the road, no one thinks that a passing mechanic is obligated to render emergency road service. So why expect doctors to render emergency medical assistance to all and sundry?

---

4   Provide an example of each of the following fallacies. Try to make your example appreciably different from those already provided for you.

> **Example:**
> Denying the antecedent: If he took all of his medication, he should be feeling better by now. He flushed half of his medication down the lavatory so he can't be feeling any better.

Rhetorical ploys and fallacies

    a   Begging the question

    b   Red herring

    c   Straw man

    d   Mistaking correlation for cause

    e   Perfectionist fallacy

    f   *Ad hominem*

    g   Deriving ought from is

    h   Weak analogy

    i   Base rate fallacy

    j   *Tu quoque*

Chapter 5

The practice of argument reconstruction

The goal of argument reconstruction is to produce a clear and completely explicit statement of the argument the arguer had in mind. The desired clarity and explicitness is achieved by putting all the argument, and nothing but the argument, into standard form: this displays the argument's premises and conclusion, and indicates the inferences between them. The strength of the argument is understood in terms of the concepts discussed in Chapters 2 and 3 – validity, inductive force, deductive and inductive soundness. In this chapter, we begin to see, in more detail, how the practice of reconstruction goes; in particular, we learn some ways of coping with commonly encountered difficulties.

Extraneous material

The first step in analysing and reconstructing an argument is to identify its conclusion, then its premises. But much of what people say or write, when advancing an argument, plays no argumentative role. Much is there for emphasis, or is rhetorical, or plays some other role than that of expressing the propositions that properly constitute the argument. When reconstructing arguments, then, we have to hive off this extraneous material. Here's an example; we've given each sentence a number in order to facilitate a detailed discussion of it:

(1) Once again the problem of young people drinking in city centres and generally creating chaos rears its ugly head.

The practice of argument reconstruction

> (2) The recent trouble in York was some of the worst. (3) It happens over and over, so much that people just seem to shrug their shoulders, accepting it as a fact of life, or a law of nature. (4) So are we simply resigned to it? (5) Do we simply accept that our young people are going to waste the best years of their lives acting like hooligans? (or rather, *being* hooligans?). (6) Do we stand idly by? (7) I don't think so. (8) Not I, at any rate. (9) And there's a ready solution. (10) Let us turn to an old solution for a new problem: compulsory military service. (11) Because they would learn habits of discipline, and something about community spirit, it's pretty obvious that young people would be a lot less likely to cause trouble when finished.

The arguer's conclusion here seems to be that Britain should introduce compulsory military service for its youth. The most important premise is the conditional statement that if compulsory military service were introduced, then the problem of drunkenness and hooliganism amongst British youth would be curtailed. Sentence 11 seems to be provided as a premise in a sub-argument for that claim. But there is a lot of other material.

- The first two sentences function as stage-setting: they alert the reader to the problem being discussed, and perhaps serve to emphasise the immediacy and severity of the problem. Sentences 1–8, as a group, are intended to persuade the reader that the problem is serious enough that something ought to be done about it. That much is surely a premise of the main argument; but it should be evident that they do not provide an argument for that claim. They merely assert the claim in a rhetorically charged way.
- The function of sentence 9 is simply to announce that the author is now going to turn from stressing the gravity of the problem to suggesting a solution – in other words, that the author is now going to give the argument.
- Sentence 10 asserts the conclusion. But it does not do so in the most economical or straightforward way. The bit about the 'old solution for a new problem' is a rhetorical flourish that should be omitted from the reconstruction. Also, it seems clear that the arguer is saying that compulsory military service should be introduced, not just that it would solve the problem. The conclusion, then, should be rewritten as 'Compulsory military service should be introduced' (there will be more on this type of 'practical' conclusion later in the chapter).

◆ Sentence 11 includes the words 'it's pretty obvious that . . .'. Phrases of that kind – phrases that merely serve to emphasise the claim being made – should always be eliminated from reconstructed arguments.
◆ Sentence 11 exemplifies a use of the word 'because' that should always be eliminated from reconstructed arguments. Often we use 'because' when speaking of causes, as in 'The cake is dry because it was baked too long'. In sentence 11, however, the word does something else: it indicates a relation between premise and conclusion of a sub-argument to the main argument. In particular, the arguer is giving the following sub-argument:

P1) If British youth were to acquire habits of discipline and community spirit, then the problem of drunkenness and hooliganism among them would be reduced.
P2) If made to perform military service, British youth would acquire habits of discipline and community service.

C1) If British youth were made to perform military service, the problem of drunkenness and hooliganism among them would be reduced.

C1 serves as an intermediate conclusion in the arguer's overall argument. The word 'because', in this usage, is equivalent to the word 'since', which functions in the following way. If we say 'If there are no clouds, it isn't raining', then we have not asserted either that it isn't raining, nor that there aren't any clouds (review the section in Chapter 2 on conditionals if this is not clear). If we say, however, '*Since* there are no clouds, it isn't raining', then we've asserted both. So what the word 'since' does is to transform a conditional statement into a statement that asserts both the conditional and the antecedent of that conditional, and thereby asserts its consequent. So what it really does is to provide a compact way of expressing simple arguments of the form: 'If P, then Q; P, therefore Q'. 'Because', in the usage we are discussing, does exactly the same thing. And since our aim in reconstructing arguments is to lay out the arguments explicitly and clearly, we should eliminate such uses of 'because' and 'since', and unpack the arguments they serve to indicate. (The causal use of 'because' is another matter, however; it will be discussed later in this chapter; see also the discussion of 'because' in Chapter 1, pp. 16–17).

   Bearing in mind that the above argument laid out in standard form is a sub-argument for C1, a reconstruction of the main argument might go like this.

The practice of argument reconstruction

> C1) If British youth were made to perform military service, the problem of drunkenness and hooliganism among them would be reduced.
>
> P3) Something should be done to reduce the problem of drunkenness and hooliganism amongst British youth.
>
> ---
>
> C2) Britain should introduce compulsory military service for its youth.

We will be in a better position later in the chapter to say whether or not such an argument is valid. The main lesson here is that the first step in reconstructing an argument is to make a list of the argument's premises and conclusion that leaves out extraneous material, and that expresses the premises and conclusion as concisely and clearly as possible. Note however, that making such a list is only the first step towards a complete reconstruction. A complete reconstruction also includes premises and conclusions that are only implicit in the original argument (and which therefore do not appear on the initial list). A complete reconstruction also displays the structure of an extended argument by displaying intermediate conclusions.

Implicit and explicit

Not only do actual statements of arguments typically include a lot of material that is inessential to the argument, they often exclude some of what *is* essential to the argument: some essential propositions are left **implicit**, when our task in reconstruction is to make the argument fully **explicit**. To say that a proposition is implicit in an argument is to say that it is part of the argument intended by the arguer – either as a premise or as the conclusion – but that it has not actually been stated by the arguer. To **make the proposition explicit** is simply to state it – in particular, to include it in our reconstruction of the argument. So a large part of argument reconstruction is to make explicit what was merely implicit in the original statement of the argument.

Consider a very simple case:

> Does Mr Jenkins know how to read? Well of course. Didn't you know that he's a successful politician?

The arguer might be presumed to be arguing as follows:

> P1) Mr Jenkins is a successful politician.
> _____
>
> C)   Mr Jenkins knows how to read.

So rendered, the argument is invalid. Nor would it be inductively forceful, because in order to know whether or not the premise makes the conclusion probable, you have to know a certain fact which the argument does not tell us – namely that successful politicians, at least usually, know how to read. But it is clear that, in drawing this inference, the arguer is making the assumption that successful politicians do know how to read. We have to make this assumption explicit:

> P1) Mr Jenkins is a successful politician.
> P2) All successful politicians know how to read.
> _____
>
> C)   Mr Jenkins knows how to read.

We leave the reader to judge whether the argument is sound.

Let us now consider a more complicated and realistic example. The Moa was a very large, flightless bird which was native to New Zealand but which is now believed to be extinct. The Yeti is a large, white-furred ape which is probably only mythical, but which some people believe actually inhabits the Himalayas. Here is the argument.

> The Moa is thought to have been extinct for at least 100 years. So naturally a great deal of scepticism is met by Paddy Freaney's claim to have seen a Moa in 1993. But as Freaney himself has pointed out, Freaney has climbed Mount Everest twice, and has not claimed to have seen a Yeti. So we ought to believe Mr Freaney's claim.[1]

The first task of argument reconstruction should be to identify, and if necessary to restate, the conclusion of the argument. Now in this case it might seem obvious that the conclusion of this argument is stated explicitly in the last sentence: We ought to believe Mr Freaney's claim. But this is not the most straightforward and informative way to put it.

---

1   Sunday *Star-Times*, Auckland, New Zealand, 10 December 1995, p. C5. This argument was actually given by Freaney himself.

The practice of argument reconstruction

Ask yourself: exactly what proposition is the ultimate concern of the arguer here? The arguer attempts to persuade us to accept a certain claim of Paddy Freaney's; but this does not really tell us the ultimate concern of the arguer until we specify what claim that is. Freaney's claim, the one that ultimately concerns the arguer, is the claim that he saw a Moa. So what the arguer is really trying to persuade us of is that Paddy Freaney did see a Moa. Here then is a case in which the conclusion is at least partly implicit; none of the sentences directly says that Paddy Freaney saw a Moa. Indeed, it might plausibly be suggested that the conclusion ultimately at issue here is one that is inferred from the proposition that Freaney saw a Moa, namely, that the Moa is not extinct. In that case it would be even more obvious that the conclusion of the argument is only implicit in the original statement of the argument. But let us not push the matter quite as far as that. For simplicity we shall take the conclusion to be that Paddy Freaney saw a Moa.

Once we have identified the conclusion, we must identify the argument's premises. What are they? It is clear that the first two sentences in the example are not premises of the argument. The arguer is arguing that Paddy Freaney saw a Moa; this proposition is certainly not supported by the fact that the Moa is thought to be extinct, nor by the fact that many people doubt Freaney's claim to have seen one. If anything, those facts make it less credible that Freaney saw a Moa. So, by the **Principle of Charity**, we should not take the arguer to be citing those facts in support of the conclusion that Freaney saw a Moa. Rather, the function of these propositions is perhaps to point out that resistance to Freaney's claim is unsurprising. (This sort of thing is very common: very frequently, an arguer will begin by explaining why there is opposition to his or her conclusion.) The only premise explicitly given is that Freaney has twice climbed Mount Everest, and has not claimed to see a Yeti. So our first shot at a reconstruction might look like this:

P1) Paddy Freaney climbed Mount Everest twice, and did not claim to see a Yeti.

C) Paddy Freaney saw a Moa.

Now this is a good start, but clearly it does not do justice to the arguer's intentions. As it stands, the argument is not deductively valid. Nor would it be inductively forceful if we were to write 'probably' before the conclusion. For suppose the only information you are given is that a certain person climbed Mount Everest twice, and that that person did not

claim that he saw a Yeti. Does that give you any reason to infer that that person did see a Moa? Obviously not. P1 does not itself make the conclusion probable.

Yet the argument as originally intended by the arguer surely does have at least some inductive force. Therefore the arguer must be implicitly relying upon some further premise or premises, which we should try to make explicit. One proposition which is surely relied upon by the argument, but which is not explicitly stated in the argument, is simply that Paddy Freaney claimed he saw a Moa. The second sentence almost states this explicitly, but does not quite do so. So we should include this in our reconstruction:

P1) Paddy Freaney climbed Mount Everest twice, and did not claim to see a Yeti.
P2) Paddy Freaney has claimed he saw a Moa.

_____

C) Paddy Freaney saw a Moa.

This is better. We now have in the premises, it seems, a statement of all the particular facts about Paddy Freaney from which the arguer is inferring that Paddy Freaney saw a Moa. But still we have not made everything explicit. In order to see this, ask yourself, why does P1 – in view of the fact that Paddy Freaney claimed he saw a Moa – support the conclusion that he did see a Moa? Well, it is something like this: The arguer seems to be responding to the charge that Freaney is a liar, that he wishes fraudulently to claim credit for a certain sort of zoological discovery. As against this, the arguer points out that Freaney climbed Everest twice without claiming a Yeti-sighting. But surely, reasons the arguer, if a person has twice climbed Mount Everest without claiming to see a Yeti, then he is not the sort of person to go around making false claims about seeing strange creatures (for such a person, having visited Everest, surely would not be able to resist claiming to have seen the notorious Yeti). Freaney climbed Everest twice without making such a claim, so he must not be such a person. If he is not such a person, then he didn't lie about the Moa. So, if he said he saw a Moa, he did see a Moa.

What is missing from the preceding reconstruction, then, is an assumed generalisation: roughly, that a person who has visited Mount Everest without claiming to see a Yeti is not a person who would fabricate the sighting of an exotic creature. This is a generalisation because it is not something said on the basis of knowledge about Paddy Freaney in particular. It is plausible to say, as a general fact about human nature,

that a person who wishes to tell 'tall tales' about some particular kind of thing is unlikely to pass up a favourable opportunity to do so.

The argument implicitly depends upon this generalisation as a premise; hence in our reconstruction we should make it explicit:

> P1) Paddy Freaney climbed Mount Everest twice, and did not claim to see a Yeti.
> P2) Paddy Freaney has claimed he saw a Moa.
> P3) A person who lies about sighting unusual or non-existent creatures would claim to see a Yeti, if he or she were to climb Mount Everest.
> ───────────────────────────────
> C) Paddy Freaney saw a Moa.

Note that, although we could have produced a valid argument by writing 'Everything Paddy Freaney says is true' (or an inductively forceful argument by writing 'Most of what Paddy Freaney says is true'), doing so would have failed to take into account the relevance of P1. If we had written that premise in place of the actual P3, then the conclusion could have been inferred from that premise along with P2, and P1 would have played no role in the argument. This is important, because the arguer clearly means to give us a very specific reason to believe this particular claim of Freaney's; he or she is clearly not simply relying on Freaney's general honesty or trustworthiness.

This reconstruction of the argument now seems to include the relevant facts about Paddy Freaney upon which the arguer is relying, and it includes the generalisation upon which the arguer is (implicitly) relying. We could make its structure a bit more explicit by making an intermediate conclusion explicit (note that we have had to change the numbering of the premises):

> P1) Paddy Freaney climbed Mount Everest twice, and did not claim to see a Yeti.
> P2) A person who lies about sighting unusual or non-existent creatures would claim to see a Yeti, if he or she were to climb Mount Everest.
> ───────────────────────────────
> C1) Paddy Freaney does not lie about sighting unusual or non-existent creatures.
> P3) Paddy Freaney has claimed he saw a Moa.
> ───────────────────────────────
> C2) Paddy Freaney saw a Moa.

We will discuss the evaluation of arguments in more depth in the next chapter, but we might as well pause to evaluate this argument. The argument is not deductively valid. To see this, suppose that the premises P1–P3 of the argument are indeed true. In that case, Paddy Freaney was not lying, making up a story, in saying that he saw a Moa. But it does not follow that Paddy Freaney saw a Moa, for he could have been mistaken rather than lying. In that case, his claim to have seen one was not a lie, but nor was it true. The inference from C1 and P3 to C2, then, is invalid. Therefore the truth of the premises would not rule out the conclusion's being false, so the argument as a whole is not deductively valid.

Might the argument be made inductively forceful, then? If the premises including P3 are true, then Freaney was not making up a story; but again, all we can conclude is that either he saw a Moa or he was mistaken. So whether the premises make it probable that Freaney saw a Moa depends upon which of these alternatives is more likely. For that, we would need more information. We would need to know the circumstances under which the alleged sighting took place – whether the light was good, at what distance it was seen, and so on. We would also need to know whether Paddy Freaney knows what the Moa looked like, hence whether he would know one if he saw one.

Since we simply do not know these things, we cannot estimate the inductive force of the argument very highly. All we can say with any confidence is that if the premises are true then either Paddy Freaney saw a Moa, or he was mistaken. We could rewrite the conclusion to state exactly that. There may be better arguments for the conclusion that Paddy Freaney saw a Moa, but given only the material at hand, the reconstruction we have given does seem to capture what the person giving the argument must have had in mind.

Let us now look at an example from Australia, that shows how making a generalisation explicit can reveal the weakness of an argument.

> The suggestion that in order to protect children from sunburn, a rule should be instituted requiring all children at our school to wear a sunhat when they are outside after 11 a.m., is unacceptable. For clearly such a rule would be an infringement upon the freedom of the individual.

The conclusion, obviously, is that a sunhat rule at school would be unacceptable. The only explicit premise is that given in the second sentence. So we may begin with:

The practice of argument reconstruction

    P1) A sunhat rule at school would infringe upon the freedom
           of the individual.

    C)   A sunhat rule should not be instituted.

Obviously, a generalisation is being implicitly assumed. A first shot at making it explicit, and which might well seem to capture what the arguer must have been thinking, might be this: that no rule which infringes upon the freedom of the individual is acceptable. So we write:

    P1) A sunhat rule at school would infringe upon the freedom
           of the individual.
    P2) No rule which infringes upon the freedom of the individual
           is acceptable.

    C)   A sunhat rule at school is unacceptable.

Here is a deductively valid argument. If the premises are true, then certainly the conclusion is true. However, this does not establish that the conclusion is true. For that, we have to ask whether the premises are true – whether the argument is not only deductively valid but deductively sound.

    Surely P1 is true, in a sense: By definition, a rule of any kind restricts, hence infringes upon, the freedom of those to whom it applies – in this case the children of the school in question. So let us grant that P1 is true. But look at P2. We often hear this sort of statement, and we are often so impressed by such a phrase as 'the freedom of the individual' that we accept the statement as true. We have the feeling that 'the freedom of the individual' is something important and valuable, therefore that anything which takes it away must be a bad thing. But, as stated here, this proposition is absurd. For as we just said, all rules 'infringe upon the freedom of the individual'. So what P2 amounts to is the absurd proposition that all rules are unacceptable. Unless you are a kind of radical anarchist, you have to conclude that this argument is not deductively sound.

    Now there may be some less sweeping generalisation that the argument might employ instead of P2, which would be more plausible, yet sufficient to obtain the desired conclusion. But the arguer has not given any hint as to what generalisation this might be. So we cannot credit the

arguer with actually having supplied a helpful argument on the issue. They may have something more plausible in mind, but they have not conveyed it.

This argument, by the way, is a good illustration of the importance of distinguishing between argument and rhetoric. It was by thinking carefully about the precise literal meanings of the words expressing the argument that we came to see that the argument is unsound. In order to do this, however, we had to have some courage: a phrase such as 'the freedom of the individual' is rhetorically powerful; nevertheless, it is just not literally true that every rule that 'infringes upon the freedom of the individual' is unacceptable; on the contrary, it is absurd to suggest it.

Connecting premises

Look back at the argument about successful politicians. And now consider this argument:

> We can assume that inflation will increase, because we know that consumer confidence is increasing.

Reconstructing, we might write:

> P1) Consumer confidence is increasing.
> _____
> C)   Inflation will increase.

We can see that the arguer is implicitly assuming that if consumer confidence is increasing, then inflation will increase. So we write

> P1) Consumer confidence is increasing.
> P2) If consumer confidence is increasing, then inflation will increase.
> _____
> C)   Inflation will increase.

In each case – the inflation argument and the politician argument – the premise which had to be made explicit in order to make the argument

valid is what we call a **connecting premise**. We have had to add connecting premises in many of the arguments we have so far considered. Usually, when people give arguments, the premises they give explicitly will be only those which pertain to the particular facts or subject-matter they are talking about. For example, someone might say 'My cat won't have kittens; she's been spayed'. The arguer explicitly sets forth the relevant fact about her particular cat, but doesn't bother to state explicitly the generalisation she assumes, that spayed cats can't have kittens. Probably she assumes this to be common knowledge; if she did not, then she would have made the premise explicit. The point then is that arguers very often leave implicit the more general assumptions they make. That is what happened in the inflation case and the politician case discussed earlier.

You can see why we call these premises 'connecting premises'. Typically an arguer informs us of some proposition, as support for another proposition. The connecting premise is what connects them. It enables the arguer to get from one to the other. But the arguer doesn't bother to spell it out, because they assume that their audience will understand what assumption is being made.

Covering generalisations

In the politician case, the connecting premise was a **generalisation** (review the discussion of generalisations in Chapter 1 if you are not clear about what a generalisation is). Connecting premises are usually generalisations. But in the inflation case the connecting premise we used was a **conditional** (see Chapter 2). This is also common. However, there is an important relationship between conditionals and generalisations that must be appreciated. Consider the following propositions:

(a) If Betty is a Siamese cat, then she has blue eyes.

(b) All Siamese cats have blue eyes.

(a) is a conditional, (b) is a generalisation. The special relationship in which they stand is that (b) is a **covering generalisation** for (a). (We can also say that (a) is an **instance** of (b), as is common in logic.) Note that (a) may be inferred from (b); the argument from (b) to (a) would be valid. But covering generalisations need not be hard generalisations: 'If Jane's cat is Siamese then its eyes are blue' is an instance of 'All Siamese cats have blue eyes', but it is also an instance of '*Most* Siamese cats have blue eyes'.

There is one further aspect to the relationship between covering generalisations and their instances that should be appreciated. Another way to express the proposition expressed by the generalisation (b) is this:

(c) If something is a Siamese cat, then it has blue eyes.

Or we could say, 'Whatever X may be, if X is a Siamese cat, then X has blue eyes'. This is exactly like (a), except we use the indefinite pronoun instead of the name 'Betty'; this makes the conditional statement into a generalisation. In other words, generalisations of the 'All A are B' sort are themselves conditionals, except they are *generalised*. (b) and (c) are generalised forms of (a).

The same goes for generalisations of the form 'No A are B', as in 'No ungulates are carnivores'. This says, in effect, that all ungulates are non-carnivores. So it could be expressed as 'If something is an ungulate, then it is not a carnivore'.

Very often, when people assert conditionals, they do so on the basis of some covering generalisation. It is important to be aware of this when reconstructing arguments. Suppose you are given the inflation argument, but without a connecting premise:

P1) Consumer confidence is increasing.
_____

C) Inflation will increase.

Suppose you reply to the arguer by saying that you just have no idea whether or not P1 constitutes a reason to infer C, that inflation is going to increase. You point out that the argument, at any rate, is certainly not valid as it stands, and is not inductively forceful. Suppose that the arguer now tries to satisfy you with:

P1) Consumer confidence is increasing.
P2) If consumer confidence is increasing, then inflation will
     increase.
_____

C) Inflation will increase.

Does this really improve the argument? It makes it valid all right. But this doesn't really help you. All that P2 says is that if P1 is true then so is C. You have already said that you have no idea whether P1 constitutes a reason to infer C. The arguer is not going to convince you merely by asserting P2; to do so is merely to assert that P1 is a good

reason to infer C. So the new version is not really an advance on the first version; although you cannot deny that the new version of the argument is valid, you can't really say whether the argument is sound, even if you grant that the arguer does know that P1 is true.

Suppose, however, that you wished to find out whether or not to believe P2. P2 says that if consumer confidence is now increasing (which it is), then inflation will increase. So the only way to find out whether P2 is true would be to look into the future to see whether inflation does increase. But that you cannot do. So what do you do? What you need is to find out whether, in general, increases in consumer confidence bring about increases in inflation. And that you can find out, by doing some statistical research (consumer confidence is defined in terms of how much people buy, and how much they borrow in order to buy). So what you really need to know, and what it is possible to find out, is whether the following covering generalisation is true:

> Whenever consumer confidence increases, inflation increases.

Probably it isn't, but the corresponding soft generalisation might well be true:

> Usually, when consumer confidence increases, inflation
> increases.

This could be established by inductive inference from a survey of past cases in which consumer confidence increased.

So: the arguer giving the inflation argument gives a conditional premise rather than a generalisation, but if the arguer believes the conditional, he or she probably does so on the basis of believing the corresponding covering generalisation. Since that is what the arguer is assuming at bottom, that is what should be included in the reconstruction of the argument. If we do not do this, then our analysis of a given argument may be superficial. Indeed, it is easy to reconstruct an argument as valid in a completely superficial way, if we do not take pains to discover what connecting premises an arguer is relying on at bottom. For example, the following argument might have been given in 1914:

P1) If the Archduke Ferdinand has been assassinated, then a general war involving all the European powers will soon break out.

P2) The Archduke Ferdinand has been assassinated.

C) A general war involving all the European powers will soon break out.

Now in fact this argument would have been sound, but someone ignorant of the political situation in Europe at the time would have needed a lot of explanation in order to see why they should accept P1. P1 is a connecting premise which ensures validity, but it hardly begins to tell us why the arguer thinks that a single assassination will lead to a general war involving several countries. In reconstructing, we should try to bring out as much of this as we can.

Finally, you should take care to distinguish the relationship we have just been discussing between conditionals and generalisations from the following sorts of cases:

If all men are mortal, then all women are mortal.

If no men are mortal, then no women are mortal.

If all men are mortal, then Socrates is mortal.

If Socrates is mortal, then all men are mortal.

These are conditionals that have generalisations as antecedents, as consequents, or as both. They are quite common.

Relevance

This chapter began with a discussion of the necessity of eliminating extraneous material from argument reconstructions. We turn now to some further aspects of this, and to some cases where the matter is more complex. Consider this argument:

At *Les Champignons* they are usually fully booked (Professor Gilmour once took me to dine there). So we can safely expect *Les Champignons* to be fully booked tonight.

Without bothering to reconstruct, you can see that this would be an inductively forceful argument (of course it isn't deductively valid). You might try reconstructing the argument like this:

P1)   *Les Champignons* is usually fully booked.
P2)   Professor Gilmour once took me to dine at *Les Champignons.*

───────────────────────────────────

C)   (Probably) *Les Champignons* will be fully booked tonight.

The practice of argument reconstruction

But, if you think about it, there is no point in including P2 in the reconstruction of the argument. P1 alone is the basis for thinking that the restaurant will be fully booked; that the arguer once dined there with Professor Gilmour is irrelevant. When a proposition stated by the arguer is irrelevant to the reasoning which delivers the conclusion, that proposition should not be included in a reconstruction of the argument.

It might seem that the only reason not to include irrelevant material in an argument is that it would be distracting. But in fact it can affect our assessment of the argument in a more important way. For suppose P1 is true, but P2 is not; the arguer, wishing to boast of a greater intimacy with Professor Gilmour than they in fact enjoy, fibbed when telling us that they dined out with him. If we were to discover this, yet persist in the above reconstruction of the argument, we should have to pronounce the argument to be unsound. Indeed, upon discovering the lie and becoming annoyed with the arguer, we might well be eager to do just this. But this would be wrong, because P2 is irrelevant to the argument. The argument should be represented like this:

P1) *Les Champignons* is usually fully booked.

C) (Probably) *Les Champignons* will be fully booked tonight.

Since P1 is true, and the argument is inductively forceful, the argument is inductively sound. The arguer's snobbish mendacity is irrelevant.

This is an obvious point, but some cases are more subtle. In such cases it is a point that is easy to lose sight of, especially when we wish to refute the argument. Consider this example:

If the tuna industry is not regulated more stringently, it will collapse altogether, because tuna populations will vanish. The evidence is obvious: tuna catches have decreased significantly every year for the past nine years. Indeed, because of depleting stocks, the Mid-Pacific Tuna company went out of business.

Suppose we initially reconstruct the argument as follows:

P1) Tuna catches have decreased significantly every year for the past nine years.
P2) Last year, the Mid-Pacific Tuna company went out of business.

C) If the tuna industry is not regulated more stringently, it will collapse altogether.

For the moment, we leave out connecting premises. Here the arguer is claiming that P1 supports the conclusion by itself. But the arguer also seems to regard P2 as providing some additional evidence, in the form of an inductive inference from one case to the general proposition that tuna stocks have declined sufficiently to cause difficulty for the tuna industry as a whole. So surely P2 is not completely irrelevant to the argument. Both P1 and P2 support the conclusion (it might be plausible to say that they support it jointly rather than independently – review the section on evidence in Chapter 3 – but here we can ignore this point). Suppose now you discover that although P1 is true, P2 is false, because the Mid-Pacific Tuna company did not go out of business; rather, it was taken over by the larger Pan-Pacific Tuna company. You might then conclude that since P2 is false, the argument is unsound, end of story. You might, indeed, exploit this mistake in attempting to discredit the arguer. If you were a person who opposed greater regulation of tuna fishing, drawing attention to this factual error might be a good rhetorical strategy to use in persuading people not to listen to arguments for increased regulation. Nevertheless, it would be a mistake from the point of view of critical thinking. This is simply because P1, by itself, constitutes a good reason for accepting the conclusion. To fixate upon the falsity of P2 simply diverts attention from the arguer's having cited P1. According to the Principle of Charity, we should simply omit P2 in our final reconstruction of the argument. More exactly: both P1 and P2, quite independently of each other, would support the conclusion. So really we should regard the arguer as having given two arguments for the conclusion, one of which we know to be unsound. So according to the Principle of Charity, we should now focus on the other argument, the one utilising P1. Adding the needed connecting premises and an intermediate conclusion, we get:[2]

P1) Tuna catches have been decreasing significantly for the past nine years.

P2) If tuna catches have been decreasing significantly for the past nine years, then the tuna population will vanish unless the tuna industry is regulated more stringently.

———————————————————————————

C1) The tuna population will vanish unless the tuna industry is regulated more stringently.

P3) If the tuna population vanishes, the tuna industry will collapse altogether.

———————————————————————————

C2) If the tuna industry is not regulated more stringently, it will collapse altogether.

———

2   It might be more plausible to reconstruct this argument as inductive rather than deductive, but we have ignored this in order to maintain the focus on the point about relevance.

The practice of argument reconstruction

This reconstruction makes for a plausible argument as it stands, but there are two further points about it that it may be useful to make.

You might find it hard to see that the argument from C1 and P3 to C2 is valid. This is probably because of the word 'unless', which many people find confusing. But this problem is easy to get round. Remember from the discussion of conditionals that, in almost every sentence in which it appears, the word 'unless' can be removed, and replaced with 'if . . . not'. For example, the consequent of P2, and thus C1, can be rewritten as:

> The tuna population will vanish, if the tuna industry is not
> regulated more stringently.

This sentence is itself equivalent to this:

> If the tuna industry is not regulated more stringently, then the
> tuna population will vanish.

which you might find easier to work with. Now it is easier to see that the argument from C1 and P3 to C2 is valid. Suppose that the tuna industry is not regulated more stringently. Then according to C1 (as rewritten above), the tuna population will vanish. But if so, then according to P3, the tuna industry will collapse. So it follows from C1 and P3 that if the tuna industry is not regulated more stringently, then the tuna industry will collapse; which is exactly what C2 says. (This is an example of a 'chain argument', as discussed near the end of Chapter 2.)

Returning now to the main theme of this section. The basic moral of this section is that the truth-values of the premises actually advanced by an arguer can be more or less relevant to the soundness of the argument. Sometimes it is highly relevant that a given premise is false, sometimes it is much less so. It depends upon the nature of the mistake, and upon the role played in the argument by the premise. The degree of relevance must therefore be taken into account in the process of reconstruction.

Ambiguity and vagueness

If our explanations of ambiguity and vagueness are not clear to you, then you should review the discussions of those concepts in Chapter 1. In this section, we explain how to cope with vagueness and ambiguity when you encounter it in arguments.

Ambiguity

In reconstructing arguments, we have to eliminate any ambiguities in the original statement of the argument. If the original statement contains an ambiguous sentence, we have to decide which of the possible interpretations was most likely intended by the arguer, and, in our reconstruction of the argument, rewrite the sentence, choosing a form of words which conveys the intended meaning unambiguously.

Let us take an example. Suppose that Jane, a Londoner, decides to invest some money in the share market. She decides to look for an investment adviser, and finds an advertisement which runs as follows:

> Sharemasters: *London's leading personal investment advice service*!

Suppose that this expresses a true proposition, and Jane accepts that what it says is true; she accepts that Sharemasters is London's leading personal investment service. Assuming that it would be fraudulent to publish such a claim if it were not true, Jane reasons as follows (for simplicity, we will ignore the reference to London in the advertised claim; this will not affect the point we are making):

P1) Sharemasters is the leading personal investment service.
P3) If I employ a personal investment service, then I should employ the leading personal investment service.

C) If I employ a personal investment service, then I should employ Sharemasters.

This looks to be a deductively valid argument. But this is not quite clear, because P1 is ambiguous. What does 'leading' mean, in this context? It could mean (and this is no doubt what the advertisers hope that people like Jane will think it means) that Sharemasters is the most *effective* personal investment service – that it secures better returns for its clients than other personal investment services. But it could also mean that Sharemasters is the *biggest* such organisation – in the sense of having the most clients (this is the sense in which McDonald's is the world's leading restaurant); it might even mean that Sharemasters is the most *profitable* personal investment service – that is, that it makes the most

profits for the owners of Sharemasters itself, not necessarily for its clients (after all, this organisation is a business, and profitability can plausibly be said to be the measure of who is the 'leader' in a certain field of business). Suppose now that you wish to evaluate Jane's argument. You are aware of the ambiguity, so you rephrase P1 in each of the three ways, yielding three arguments:

A

P1 (a) Sharemasters secures higher returns for its clients than any other personal investment service.

P2 (a) If Jane employs a personal investment service, then she should employ the personal investment service that secures higher returns for its clients than any other personal investment service.

C) If Jane employs a personal investment service, she should employ Sharemasters.

B

P1 (b) Sharemasters is the most profitable personal investment service.

P2 (b) If Jane employs a personal investment service, then she should employ the most profitable personal investment service.

C) If Jane employs a personal investment service, she should employ Sharemasters.

C

P1 (c) Sharemasters is the biggest personal investment service.

P2 (c) If Jane employs a personal investment service, then she should employ the biggest personal investment service.

C) If Jane employs a personal investment service, she should employ Sharemasters.

Now all three arguments are deductively valid. But are any of them sound? First of all consider arguments B and C. Even if we assume that P1(b) and P1(c) are true, it does not seem that these are sound arguments, because it does not seem that either P2(b) or P2(c) are true. Perhaps they have some plausibility – since presumably if Sharemasters were highly *un*successful at making money for its clients, then it would not have grown so big or become so profitable. But these would not be very reliable assumptions; to reason in the manner of C, for example, would be like concluding that since McDonald's is the biggest restaurant in the world, it has the best food.

So you can conclude that Jane's argument has a good chance of being sound only if argument A is what she had in mind. Suppose you investigate, and find that it is true that Sharemasters is the biggest as well as the most profitable personal investment service in London. Suppose you find, however, that it is not the most effective personal investment service in London; its size and profitability are due to its high fees, its organisational efficiency, and effective advertising. Hence P1(b) and P1(c) are true, but P1(a) is false. You must therefore conclude that although it is valid, argument A is unsound. Since neither B nor C seems to be sound, you must conclude that Jane's original argument is unsound; on none of the three possible interpretations does it appear to be sound.

Eliminating the ambiguity of the original argument was crucial in discovering this, for one could very easily think that the original argument was sound. Advertisements of this kind may be said to exploit this sort of ambiguity. If Sharemasters is indeed the largest personal investment service in London, but not the most effective, then the slogan does express something true; yet the advertisers might hope that readers will interpret the claim along the lines of argument A (or possibly along the lines of B, making the questionable assumption that the biggest service of a given sort is likely to be the most effective).

Remember that a primary purpose of reconstruction is to represent the propositions which constitute an argument in the clearest possible way. Thus we should have no qualms about changing the language used to express those propositions; in changing it, we are only trying to gain a better grasp of what the arguer was thinking. There is no guarantee that we will not change or distort the arguer's thinking, but there is no point in allowing ambiguous language to remain unchanged. For we simply cannot evaluate an argument if we do not know exactly what argument we are evaluating. If we simply cannot decide between two

interpretations of an ambiguity, then we must give both interpretations of the argument, and evaluate the two arguments independently.

Vagueness

As noted in Chapter 1, many words and phrases are vague in the way that 'bald' is vague. Examples include 'tall', 'orange', 'heap of sand' (i.e. we cannot draw a precise boundary between tall and not-tall, between orange and red, and cannot say exactly how many grains of sand you need to have a heap). These words pertain to a certain one-dimensional quality, such as height, region of the colour spectrum (which proceeds gradually from red through orange, yellow, green, blue and violet), or amount of sand.

These sorts of words seldom present problems for actual argument reconstruction and assessment.[3] More important for critical reasoning are forms of vagueness which do not pertain simply to one-dimensional qualities like height. Indeed, many of the most rhetorically powerful or emotionally provocative words in public (and private) discourse are vague in this way. Consider:

• politics		• liberal
• elitism		• racism
• conservative		• love
• sovereignty		• radical
• rights		• responsibility

As pointed out in Chapter 1, we often have the feeling that these things are bad, or that they are good, without any precise idea of what we mean by them. What they signify is typically a whole group or cluster of things which are not unified in any exact way.

Let us take 'conservative' (in the political sense) as our example. There are various attributes often associated with someone's being 'conservative':

---

3   But they do cause trouble for logic. Consider the term 'sand dune'. There is no exact rule for how much sand is needed to constitute a sand dune, but here is one proposition about sand dunes that seems undeniable: If X is a sand dune and one grain of sand is removed from X, then X is still a sand dune. But now suppose that X is a sand dune. If we remove one grain of sand from X, then X is still a sand dune. But then if we remove one grain of sand from this, we still have a sand dune. And so on. But then eventually we run out of sand, so we have no sand dune. So at some point, X was a sand dune consisting of n grains of sand for some number n, such that X ceased to be a sand dune, upon the removal of one grain of sand. But that seems absurd; we cannot imagine a sand dune such that removing one grain of sand destroys its status as a sand dune. So we seem to have a contradiction. This paradox – *Paradox of the Heap* or *Sorites paradox* – was puzzled over by the ancient Greeks, and there is still no consensus as to how to avoid the apparent contradiction.

Conservative

► Believes in minimum government.

► Favours free market economics.

► Favours privatisation of government industries and services.

► Against the dole, welfare, etc.

► Supports 'traditional' values – pro-family, anti-gay, etc.

► Against government help for minorities, women, etc.

► Supports strong police, military, severe prison sentences.

► Supports business as against labour, environmental groups.

Someone could rightly be called 'conservative' without exhibiting all these attributes or tendencies. Exactly which attributes you must have, and which you must not have, to be a conservative, is simply not clear. So 'conservative' is vague.

Now, in reconstructing arguments, the best thing to do with vague words is simply to eliminate them. This can be seen from the following example:

> MP Jeremy Price has made his conservative stance clear by favouring 'traditional values' and stronger sentencing for criminals. Therefore we can certainly assume he will oppose any new laws protecting the environment.

It might occur to you to represent this as a two-stage argument, thus:

P1) Jeremy Price favours 'traditional values' and stronger sentencing for criminals.

C1) Jeremy Price is a conservative.

C2) Jeremy Price will oppose new laws protecting the environment.

The inferences from P1 to C1, and from C1 to C2, depend upon implicit generalisations, which, having learnt earlier lessons well, you know you would do well to make explicit:

The practice of argument reconstruction

> P1) Jeremy Price favours 'traditional values' and stronger sentencing for criminals.
> P2) All those who favour 'traditional values' and stronger sentencing for criminals are conservatives.
>
> ---
>
> C1) Jeremy Price is a conservative.
>
> P3) All conservatives oppose new laws protecting the environment.
>
> ---
>
> C2) Jeremy Price will oppose new laws protecting the environment.

The argument from C1 and P3 to C2 is deductively valid, but it might then occur to you that P3 is implausible. Surely not *every* conservative is so disdainful of environmental restriction. So you replace 'all' with 'most':

> P1) Jeremy Price favours 'traditional values' and stronger sentencing for criminals.
> P2) All those who favour 'traditional values' and stronger sentencing for criminals are conservatives.
>
> ---
>
> C1) Jeremy Price is a conservative.
>
> P3) Most conservatives oppose new laws protecting the environment.
>
> ---
>
> C2) (Probably) Jeremy Price will oppose new laws protecting the environment.

Now this argument might seem plausible to the average person; it might seem to be inductively sound. But the argument appears better than it is, because of the presence of the word 'conservative'. And in fact, all the actual logic of the original argument can be reproduced without making use of that word at all. It is simple: the actual evidence the arguer has for their conclusion that Jeremy Price will oppose new environmental laws is that Price favours traditional values and stronger sentencing for criminals. The arguer supplies no further evidence, and the word 'conservative' plays no role at all. Thus, adding the necessary generalisation, we might just as well represent the argument like this:

P1) Jeremy Price favours traditional values and stronger
sentencing for criminals.
P2) Most people who favour traditional values and stronger
sentences for criminals oppose new laws protecting the
environment.

C) (Probably) Jeremy Price opposes new laws protecting the
environment.

This remains inductively forceful, and it is no less inductively forceful
than the previous version. But the inductive soundness of this argument
is far from certain, because P2 is far from certain. This reconstruction
thus makes the weakness of the original argument perfectly clear. It does
so by removing the distraction created by the word 'conservative' (even
if you do think that P2 is true, you have to admit that this reconstruc-
tion centres our attention on the real issue).

Since many of the most rhetorically highly charged words in public
discourse are also vague, eliminating them from our argument-
reconstructions achieves two things: it clarifies the argument, and, by
eliminating emotionally provocative words, enables us to focus without
distraction upon the logic of the argument.

The best thing to do with ambiguous or vague language, then, is
simply to replace it with language which is not vague or ambiguous. The
aim is to employ language which will express the intended propositions
without ambiguity or vagueness. But this is not always possible. Where
a sentence is ambiguous, we cannot always tell which of the different
possible interpretations was intended by the arguer, even if we apply the
Principle of Charity. In such a case we can assess each of the possible
versions of the argument, but we may have to confess that we cannot
tell which version the arguer intended. And – especially where the
language used by the arguer is vague rather than ambiguous – we have
to admit that the arguer's *thinking* may simply have been vague or
confused, not just his or her language. Indeed, where vague words such
as those listed above play a role in an argument, this is very often the
case. In constructing your own arguments, such words are best simply
avoided. They tend to obscure the issue rather than clarify it. However,
not all vague language can be removed. We will return to this in the
next chapter.

The practice of argument reconstruction

More on generalisations

As we noted in Chapter 1, only hard generalisations can rightly be conveyed by using a quantifier-word like 'all', 'no', or 'every'. Indeed, that is the usual function of those words – to make it perfectly explicit that a hard generalisation is what is intended. For example someone at a meeting of Parliament might say, 'Every single MP in this chamber takes bribes', rather than 'The MPs in this chamber take bribes'. Soft generalisations, indeed, are very often expressed without any quantifier at all, as in 'Children like sweeties'. Other times, we add a quantifier such as 'most', or 'almost all' to a soft generalisation, in order to make it clear that a soft generalisation is what is intended, and also to make it clear just how soft (or how close to being hard) it is meant to be.

Since there is often confusion over the difference between hard and soft generalisations, we should, when reconstructing arguments, always make it clear whether a generalisation is hard or soft (the one exception to this is the case of causal statements; this will be discussed near the end of this chapter). The confusion most often arises when the generalisation, as stated, lacks an explicit quantifier. The intended quantifier is merely implicit, so there is room for misinterpretation. A generalisation in which the quantifier is merely implicit, then, is actually a kind of ambiguity. The way to alleviate the ambiguity is to add an explicit quantifier. All generalisations in reconstructed arguments should contain explicit quantifiers.

The scope of a generalisation

Consider the following hard generalisations:

1 All cows are herbivores.
2 All black cows are herbivores.

The **subjects** of these generalisations – what the generalisations are about – are cows and black cows, respectively. Both generalisations are true, and they stand in a special relationship. There are two aspects of this relationship. First, they attribute the same feature to their subjects (that of being herbivores). Second, the subject of the second is a subset of the first (all black cows are cows). Thus we say that the **scope** of (1) is **wider** than that of (2) and conversely that the **scope** of (2) is **narrower** than that of (1). Note that we can compare generalisations in this way only
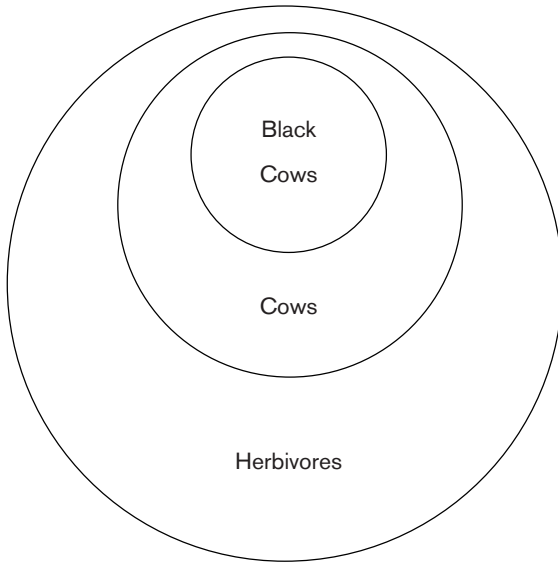
Figure 5.1 The scope of a generalisation

when the subject of one is a subset of the subject of the other. We cannot, for example, say that the generalisation 'All lions are carnivores' is narrower than, or wider than, either (1) or (2). Nor can we compare the scope of 'All sheep are herbivores' to that of either (1) or (2).

Figure 5.1 represents the situation expressed by (1) and (2). It can sometimes be important to adjust the scope of a generalisation, making it either narrower or wider. Usually, in reconstructing arguments, we have to narrow them; hardly ever do we have to widen them. Suppose, for example, you are a fervent environmentalist, and believe that radical measures must be taken immediately to protect the environment from air pollution. In particular, you believe that our reliance upon petroleum products – such as petrol for cars – must be halted as soon as possible. So you argue:

> We must find alternatives to machines that emit carbon
> monoxide; hence we must find alternatives to the automobile.

The reasoning here is pretty clear, but in reconstructing, some care must be taken with the generalisation implicit in it. The implicit generalisation, obviously, is that automobiles emit carbon monoxide. The argument might be reconstructed as follows:

The practice of argument reconstruction

P1) We must find alternatives to all vehicles that emit carbon monoxide.

P2) All automobiles emit carbon monoxide.

C) We must find alternatives to the automobile.

Note that in writing P1, we added the quantifier 'all'; this makes for a somewhat awkward sentence, but such awkwardness is a price we sometimes have to pay for the sake of explicitness. Now assume you really do believe P1. This argument is deductively valid. However, you must admit that it has a weak point. For P2, as written, is clearly untrue. Some automobiles, for example, are electric, and run on batteries. Since they do not burn petrol (or any other petroleum product), they do not emit carbon monoxide (indeed they do not pollute the air at all). So even if P1 is true, you have failed to give a sound argument.

Now you could make the argument inductively sound by changing the word 'All' in P2 to 'Most'. But you could do better, for the following two reasons:

First, it could become false that most automobiles emit carbon monoxide, without any reduction at all in the number of petrol-driven cars, in which case the pollution (and other problems) caused by the burning of petrol might not have been curtailed at all. For it would be possible for a large number of electric cars or other alternatives to be produced over a few years without any reduction in the number of petrol-driven cars. In that case, P2 would be false but you would still want to argue against petrol-driven cars, since they would still be emitting the same amount of pollution.

Second, changing 'All' to 'Most', although it renders the argument sound, does not clarify the issue as much as another alternative.

This other alternative is to change the scope of the generalisation expressed by P2 – in particular, to reduce its scope. Instead of 'All automobiles emit carbon monoxide', we can more accurately write 'All petrol-driven automobiles emit carbon monoxide', rewriting the argument thus:

P1) We must find alternatives to all vehicles that emit carbon monoxide.

P2) All petrol-driven automobiles emit carbon monoxide.

C) We must find alternatives to the petrol-driven automobile.

Assuming the truth of P1, this is a deductively sound argument. Even if you doubt P1, this is, at any rate, a much better argument than the earlier version; since the new P2 is a narrower generalisation than the old P2, its premises have a better chance of being true. Furthermore, by narrowing the generalisation, the issue is defined more exactly. We now have it explicitly before us that the point at issue is the use of petrol-driven automobiles, not automobiles in general.

Thus, when reconstructing arguments, we should take care not to employ a hard generalisation that is wider in scope than we need, if there is anything doubtful about the wider one that could be eliminated by employing a narrower one. If a narrower (but hard) generalisation will suffice for constructing an argument for the desired conclusion, then we should employ the narrower one. This is not to say we should always choose narrower generalisations whenever possible; for example, it would not improve the argument to rewrite P2 as 'All *green* petrol-driven auto-mobiles emit carbon monoxide'. For adding the word 'green', besides not making the generalisation any more likely to be true (or at least not appreciably so), renders the argument invalid. The point is that if one generalisation is narrower than another and more likely to be true, then, provided that it is sufficient for obtaining the desired conclusion, the narrower one should be employed.

Note that in some cases there is no natural word or phrase for the class of cases we wish to generalise about. In such cases we have to reduce the scope of a generalisation by explicitly excepting a certain class of what would otherwise be counter-examples. For example, consider the infer-ence 'That is a mammal; therefore it doesn't fly'. The generalisation needed is 'No mammals *except bats* fly', since 'No mammals fly' is false (and of course we would also need the premise that the creature in ques-tion is not a bat; we ignore flying squirrels).

Practical reasoning

The conclusion of the above argument about the automobile is one that has a practical conclusion: rather than saying that some proposition is true, it enjoins or commends a particular action. What the argument says, roughly, is that doing one thing (finding alternatives to the petrol-driven automobile) is necessary if a certain desirable outcome or end (finding alternatives to vehicles that emit carbon monoxide) is to be achieved. Other arguments with practical conclusions are those that say that a certain action would be sufficient to bring about the desired end, and for that reason ought to be performed. Still others say that a certain action would lead to a certain undesirable outcome, and therefore ought not to

be performed. Reasoning of this kind is often called practical reasoning, or means–end reasoning, and is normally based upon two sorts of considerations: First, an outcome is specified as being either desirable or undesirable (in this book we prefer to use 'must' or 'should' to specify the outcome, as illustrated below); second, there is some proposition put forward that says either (1) that if such-and-such action is performed, the outcome will result; or (2) the outcome will result *only* if such-and-such action is performed (that is, the outcome will not come about unless the action is performed).

In fact, there are no fewer than eight different types of argument that concern a relation between action and outcome. Just for the purpose of illustration, then, suppose that the outcome is that the amount of chocolate in the world be increased, and the action in question is X. Then, using 'should', the eight types of practical reasoning can be sketched like this:

1

P1) The amount of chocolate in the world should be increased.
P2) If action X is performed, then the amount of chocolate in the world will be increased.

C) Action X should be performed.

2

P1) The amount of chocolate in the world should be increased.
P2) If action X is not performed, then the amount of chocolate in the world will not be increased.

C) Action X should be performed.

3

P1) The amount of chocolate in the world should be increased.
P2) If action X is not performed, then the amount of chocolate in the world will be increased.

C) Action X should not be performed.

4

    P1) The amount of chocolate in the world should be
increased.

    P2) If action X is performed, then the amount of chocolate in
the world will not be increased.

    C)   Action X should not be performed.

5

    P1) The amount of chocolate in the world should not be
increased.

    P2) If action X is performed, then the amount of chocolate in
the world will be increased.

    C)   Action X should not be performed.

6

    P1) The amount of chocolate in the world should not be
increased.

    P2) If action X is not performed, then the amount of chocolate
in the world will not be increased.

    C)   Action X should not be performed.

7

    P1) The amount of chocolate in the world should not be
increased.

    P2) If action X is not performed, then the amount of chocolate
in the world will be increased.

    C)   Action X should be performed.

8

    P1) The amount of chocolate in the world should not be
increased.

    P2) If action X is performed, then the amount of chocolate in
the world will be not be increased

    C)   Action X should be performed.

The practice of argument reconstruction

There is, however, an important complication that tends to arise most often in connection with type (1). Suppose that we believe that increasing the number of doctors would improve the NHS. We might set out the argument like this:

P1) The NHS should be improved.
P2) If the number of doctors were increased, the NHS would improve.

C) The number of doctors should be increased.

This is a start, but as it stands it is really not adequate. For if you think about it, the argument is clearly not valid. It does not follow, from the fact that such-and-such would bring about some desirable result, that we should do such-and-such (also look again at the earlier argument about the sunhat rule). There are two reasons.

First, we need to know that the cost of the proposed action does not outweigh the benefit of the outcome. For example, it would certainly improve the NHS if the government were to increase its budget by ten-fold. But that would not be a good idea, as the cost would be far too great. This sort of thing – 'weighing the costs' – is clearest where money is at stake, but it is not limited to monetary considerations. If you want stronger muscles, for example, then you have to weigh the desirability for you of stronger muscles against the 'cost' of exercising (the time expended, the pain, etc.). Indeed, even in the argument about the NHS, we are not assuming that the value of an improved NHS can be assessed in monetary terms.

Second, we need to know that there is not some other means that would bring about the same benefit but at a lower cost. We need to know, that is, that the proposed action is the most *efficient* or *economical* way to bring about the desired outcome.

In reconstructing practical reasoning, then, we have to incorporate both of these points as premises. The argument concerning the NHS, then, would go something like this:

P1) The NHS should be improved.
P2) If the number of doctors were increased, the NHS would improve.
P3) The benefits of improving the NHS would outweigh the costs of increasing the number of doctors.
P4) Increasing the number of doctors would be the most efficient means of improving the NHS.

C) The number of doctors should be increased.

This argument is valid.[4] Strictly speaking, P1 and P2 are now redundant, since they can be inferred from P3 and P4; but it is harmless, and it makes it clearer, to leave them in. Of course, the argument is a bit vague: it does not say to what extent the NHS should be improved, nor by how many the number of doctors should be increased. Someone actually advancing this argument would want to fill in these details.

In considering this issue we focused on an argument of type (1), but the weighing of costs and benefits may enter into any of the eight types of practical argument. None of the eight patterns is valid as it stands. For example, consider type (7). It might be that the cost of not performing action X would be greater than the benefit of not increasing the amount of chocolate (that is to say, it might be that the benefit of performing it is greater than the cost of increasing the amount of chocolate).

Balancing costs, benefits and probabilities

Someone giving the above argument about the NHS might concede that there is some possibility, even if the number of doctors were increased, that the NHS would not improve. Nevertheless, so long as he or she is reasonably certain that it would, it would be simplest to leave the conditional P2 as it is, rather than inserting 'probably' before its consequent, thereby making the argument inductive rather than deductive.

In other cases, however, practical arguments are more plausibly reconstructed as inductive. For we sometimes have to balance costs and benefits in a more complicated way. For example, suppose you are repairing a window. Having removed it, you are invited to a party, which is taking place now. You know the party would be a lot of fun, but although you could very easily repair the window tomorrow, there is no time to replace the window before going to the party; either you leave the window off and go to the party, risking the possibility of rain getting in (assume the

---

4   In endorsing this style of argument as valid we are assuming that all practical considerations can be subsumed under the categories 'cost' and 'benefit'. We do so for simplicity; but it should be acknowledged that many philosophers hold that *rights* and *duties* are practically relevant considerations which cannot be explained in terms of cost and benefit (cannot be explained, that is, from a *utilitarian* or *consequentialist* point of view). More generally, the relation of rule-theoretic concepts such as rights and duties to value-theoretic concepts such as cost and benefit is a central concern of the philosophical study of ethics. A way to accommodate this without complicating our argument reconstructions is simply to think of rights and duties as entering into the calculation of cost and benefit in their own right, without any assumption that can be explained in terms of some independent conception of cost and benefit.

window is too high up for there to be a risk of burglary), or you continue working on the window and miss the party. What should you do? Obviously the benefit of going to the party is high, but so would be the cost of rain getting in. If we assume that these are roughly equal, then clearly your decision should rest on the probability of rain: if the probability of rain is less than 1/2, then you should go to the party; if it is higher than 1/2, then you should not. But suppose the cost of rain getting in would be much greater than the benefit of attending the party. In that case you should not risk going to the party even if the chance of rain is fairly low. We can represent this in a table:

	Rain	No rain
Go to party	−10	+5
Not go to party	0	0

Here we are assuming that the badness of rain, if you leave the window off, would be twice the goodness of going to the party, and that not going to the party would be neither bad nor good (it does not matter that we used −10 and +5 to represent this relationship; we could just as well have used −4 and +2, or −2 and +1). In this situation, then, you should stay home, unless the probability of its raining is less than half the probability of its not raining. That is, you should stay home unless the probability of rain is less than one third.

What we are leading up to is the concept of **expected value**:

Let o_1, o_2 ... be the possible outcomes of an action A; let V(o) be the value (cost or benefit) of each outcome o, and let P(o) be the probability of each outcome (given that action A was performed). Then the *expected value* of an action A is:

$$[P(o_1) \times V(o_1)] + [P(o_2) \times V(o_2)] + \ldots$$

That is, for each possible outcome of the action, you multiply the probability of the outcome by its value (its cost or benefit, as the case may be). Then you add these figures together to get the expected value of the action. The idea behind this is that, given a range of possible actions, one should do whatever maximises expected value. If one's possible actions are A_1, A_2, ..., and one of these – say A_k – has the highest expected value, then

one should perform A_k. In the case of the window and the party, the possibilities are: stay and repair the window, or go to the party. If we assume that the probability of rain is 0.5, then the expected value of going to the party, according to the values given in the chart above, would be

$$0.5 \times (-10) + 0.5 \times 5 = (-5) + 2.5 = -2.5.$$

The expected value of not going to the party is

$$0.5 \times 0 + 0.5 \times 0 = 0.$$

Since the expected value of staying to repair the window (0) is greater than that of going to the party (–2.5), you should stay to repair the window.

It is very helpful to have a firm grasp of the concept of expected value, because it is one of the areas in which people most frequently make mistakes in reasoning. For example, we sometimes see arguments like this:

> The bottom line is this: no matter how safe the government says it is, they cannot rule out the possibility of a catastrophic accident. We should decommission all nuclear power plants as soon as possible.

If the arguer is right, then the expected value of decommissioning all nuclear plants should be greater than that of not doing so. But, although this may be true, the arguer has given us no reason to believe it. In order to reach that conclusion, we would need to know the actual probability of an accident; we would also need to know how the cost of an accident would compare with the benefits and other costs of continuing to use nuclear power, as well as the benefits and costs of relying on other sources of energy. The arguer has provided none of this, and the mere fact that an accident is not impossible is not disputed by anyone. Nor does anyone dispute that a nuclear accident would be a very bad thing. But the mere fact that something bad – no matter how bad – is a *possible* outcome of some action certainly does not establish that the action should not be performed.

Explanations as conclusions

In the first chapter, we took care to distinguish **arguments** from **explanations**. An argument supplies reasons why we should believe a certain proposition whose truth-value is in question. By contrast, an explanation tells us why it is that a certain proposition is the case, when the truth-value of that proposition is not in question. It is especially important to observe this distinction when dealing with arguments whose conclusions

are themselves explanations. The aim of this sort of argument is to persuade the audience that such-and-such is the actual cause of a fact or event. Such arguments are very common. For example:

> We must not have fertilised that plant. It hasn't got a disease. We've watered it correctly, and it gets enough sun, but it still isn't growing well.

This passage does involve an explanation; but it is, nevertheless, an argument. The conclusion is that the cause of the plant's poor growth is that it wasn't fertilised. So we might reconstruct it like this:

P1) If a plant is not growing well, then either it has a disease, it is getting the wrong amount of water, it isn't getting enough sun, or it has not got enough fertiliser.
P2) This plant is not growing well.
P3) This plant has not got a disease, it gets the right amount of water, and it gets enough sun.

C)   This plant has not got enough fertiliser.

This argument is valid. But it is not quite right, because the conclusion does not say that the plant's poor growth is caused by its not having enough fertiliser. It is a perfectly good argument, and indeed it is part of what the arguer is trying to establish. But it is not all. In order to display the full force of the argument, we need to use the word 'cause' in both the premises and the conclusion:

P1) If a plant is not growing well, then the cause is either that it has a disease, that it is getting the wrong amount of water, that it isn't getting enough sun, or that it has not got enough fertiliser.
P2) This plant is not growing well.
P3) This plant has not got a disease, it gets the right amount of water, and it gets enough sun.

C)   The cause of this plant's not growing well is that it has not got enough fertiliser.

Or, as we might more naturally put it, this plant is not growing well *because* it has not got enough fertiliser – that would be the causal use of 'because' that was discussed in the first section of this chapter. Note

that in order to obtain the desired conclusion, it was necessary to include the word 'cause' (or the word 'because') in P1. If we had left P1 as it is in the first version, then the inference to the conclusion in the second version would have been invalid. For it is certainly not true that if A is true (the plant hasn't enough fertiliser), and B is true (the plant isn't growing well), then A is the cause of B (though of course if A is the cause of B, then A and B are both true).

In fact this is not quite the whole argument, but only a sub-argument. In order to get the final conclusion, 'This plant is not growing well because we did not fertilise it', we'd have to add another premise after the given conclusion, 'If we had fertilised this plant, then it would have had enough fertiliser'.

Note that the correct reconstruction is not like this:

P1) If a plant has got a disease, it is getting the wrong amount of water, it isn't getting enough sun, or it has not got enough fertiliser, then it doesn't grow well.
P2) This plant is not getting enough fertiliser.

C) This plant doesn't grow well.

In the passage, the arguer assumes that the audience agrees that the plant isn't growing well. This fact is not in question. The aim of the argument is to persuade the audience as to the cause, that is, the explanation, of this fact.

All such arguments begin with an accepted fact – some proposition that the arguer expects his or her audience already accepts. The arguer's aim is to answer the question, 'What is the correct explanation of this fact?', by giving an argument. In the above example, the arguer recognises that there are several possible causes of the agreed fact, and, by eliminating all but one, advances the remaining one as the correct explanation. The pattern can be represented like this:

P1) (The agreed fact).
P2) (The agreed fact), was caused by either A or B (or C, or . . .).
P3) B is not the case (nor is C, nor . . .).

C) (The agreed fact) was caused by A.[5]

---

5   A more complex kind of case is that where we know that the cause was A or B, and that A and B were both true, yet we still wish to determine which was the actual cause (e.g. we might know that the patient's death was caused by either heart disease or liver disease, and that both his heart and liver were diseased). But we shall not discuss these sorts of cases here.

The practice of argument reconstruction

Normally P2 will be backed by a covering generalisation, of the form 'Whenever such-and-such happens, it is caused by . . .'. But not always; sometimes we know what the possible causes of a given fact or event are, but are unable to articulate generalisations from which they may be inferred. In cases where we do know the appropriate generalisations, however, we should, as always, include them in the reconstruction. That is what we did in the plant example. Of course the list of possible causes listed in P2 may be of any length.[6]

The generalisations appealed to in arguments of this kind are often soft rather than hard. For example, a sudden increase in the temperature of an automobile engine's cooling system is usually, but not always, caused either by a leak in the cooling system or a broken fan belt. In such cases, we must insert 'usually', or 'probably', or some suitable variant, into the appropriate place in P2. Likewise, we may not be able to rule out with certainty the possible explanations listed in P3. Again, in such cases we should have to insert 'probably', 'usually', or some suitable variant, into the appropriate place in P3 (or perhaps 'probably not', etc.). In these cases the arguments will be inductive rather than deductive. Such an argument might look like this:

P1) The temperature of the engine's cooling system has suddenly increased.
P2) Almost always, a sudden increase of the temperature of an engine's cooling system is caused by a broken fan belt, or a leak in the cooling system.
P3) It is very unlikely that the fan belt has broken.

C) Probably, the increase of the temperature of the engine's cooling system was caused by a leak in the cooling system.

In this case, P3 might be the conclusion of another argument; for example it may have been inferred from the fact that the fan belt is new.

Causal generalisations

In the preceding examples we were concerned with causal relationships between particular events or states of affairs. We were concerned, for

---

6   It may, indeed, contain only one item, in which case there is no need for P3. For example, smoke is always caused by fire; therefore if we see smoke, then since we know that it is always caused by fire, and never by anything else, we can conclude immediately that the smoke we see is caused by a fire. Such cases of course are not very interesting.

example, with whether or not the particular event in a car's cooling system was caused by a particular fault. However causal statements often appear as generalisations about types of events or states of affairs, as in:

Powerful electric shocks cause the muscles to contract.

This sort of statement is relatively unproblematic. What it means, roughly, is that whenever powerful electric charges are applied to a living person's body, the muscles contract (unless there is something wrong with their nervous system). So it is a hard generalisation about events. However, the word 'cause' does not always, or even typically, indicate a hard generalisation of this kind. Consider:

Smoking causes cancer.

This is not a hard generalisation about events, persons or states of affairs. It does not mean, for example, that every act of smoking causes an outbreak of cancer, or that everyone who smokes gets cancer. It does not even mean that smoking usually causes cancer (since it is not true that most smokers develop cancer). The following argument, then, is not inductively forceful:

P1) Smoking causes cancer.

C) If you smoke, then, probably, you will get cancer.

What, then, does 'Smoking causes cancer' mean? Roughly, what it means is that each person in a given situation is more likely to develop cancer if they smoke, than that person would be, in that very same situation, if they did not smoke. Smoking raises the probability of getting cancer. So this argument is deductively valid:

P1) Smoking causes cancer.

C) Therefore, if you smoke, you will be more likely to get cancer than if you don't.

The practice of argument reconstruction

But note that the statement about smoking does not tell us to what degree smoking raises the probability of getting cancer. Because of this, it usually does not help to try to add an explicit quantifier to a causal statement.

This connects with a further issue regarding generalisations. Consider:

A recent study of primary school children has discovered a strong correlation between diet and school performance: better diets were strongly linked to better test marks. One of the simplest things we can do to improve performance in primary schools, then, is simply to improve the foods on offer at school refectories.

To discover a 'correlation' between X and Y is to discover that the proportion of things that exhibit the feature Y is higher amongst things that exhibit feature X. The arguer here seems to infer a causal relationship from the correlation of better test marks with better diets. As noted in Chapter 4, this is the fallacy of confusing correlation with cause. The inference goes like this:

P1) X is strongly correlated with Y.

C)   X causes Y.

But this inference is unwarranted. This can be seen easily by considering the following example of the same pattern:

P1) High concentrations of bicycles are strongly correlated with air pollution.

C)   High concentrations of bicycles cause air pollution.

P1 is actually true, since high concentrations of bicycles and high levels of air pollution are typically found in the same place, namely large cities. But obviously C is false. This argument-pattern is evidently not valid, and not inductively forceful either.

Usually, when there is a correlation between X and Y without causation, what is going on is that some underlying factor causes both X and Y. In the case of bicycles and pollution, the common underlying cause is

the presence of a large and concentrated population of human beings. As for the case of children and test scores, it is certainly plausible to suppose that diet influences school performance, but there may be a common underlying cause in that case as well. For example, it might be that children from better-educated households tend to be more intelligent (if only because intelligent people are more likely to complete a university education), and that better educated households tend to eat better diets (perhaps because they tend to have more money, and hence can afford a better diet). Or it might be that children from poor households tend to eat poorer diets, and that poor households tend to suffer from other sorts of family problems that tend to impair school performance.

Under what circumstances, then, can we legitimately infer the presence of a causal relationship? The answer should be evident from the rough-and-ready definition of causation given above: in order to infer a causal relationship from a correlation between X and Y, we need to know that the correlation holds, or would hold, even when other possible causes of Y are or would be absent. In other words, we need to know that X makes Y more likely regardless of the circumstances in which we find X. We need to rule out other possible causes. So what we would need to know in the case of the children, for example, is whether or not children from well-off households do worse at school if their diets are poor, and likewise for children from well-educated households.

It is important to be aware of these issues, and especially to be able to point it out when a causal relationship is wrongly inferred from a correlation. But it would take us too far into the subject to give a general recipe for validly inferring causal relationships from correlations. Suffice it to say that a causal relationship entails a correlation, but a correlation does not entail a causal relationship.

A shortcut

Where an argument contains a conditional amongst its premises, we have, in order to infer the consequent of the conditional, to write down its antecedent as a separate premise. This means that we have to write a certain sentence down twice. This can be cumbersome, and it often makes the structure of an argument more difficult to see than it need be. Consider again the first sub-argument in the argument about the tuna industry:

The practice of argument reconstruction

     P1) Tuna catches have been decreasing significantly for the
          past nine years.
     P2) If tuna catches have been decreasing significantly for the
          past nine years, then the tuna population will vanish
          unless the tuna industry is regulated more stringently.

     C1) The tuna population will vanish unless the tuna industry is
          regulated more stringently.

P2 is a conditional whose antecedent is P1. Instead of rewriting P1 out in full, we may abbreviate it simply as 'P1'. So we may rewrite the argument as follows:

     P1) Tuna catches have been decreasing significantly for the
          past nine years.
     P2) If (P1), then the tuna population will vanish unless the
          tuna industry is regulated more stringently.

     C1) The tuna population will vanish unless the tuna industry is
          regulated more stringently.

This is easier to read, and it saves you having to write things out unnecessarily. Henceforth, whenever an argument draws an inference from a conditional premise, feel free to abbreviate in this way.

CHAPTER SUMMARY

This chapter was concerned to address some of the main logical problems encountered in the **reconstruction** of arguments.

    In many cases, one or more premises upon which a conclusion depends is left **implicit** by the arguer. A necessary part of reconstruction is to make such premises **explicit**. Implicit premises are usually, though not always, **connecting premises**: these are either **conditionals** or **generalisations**. Where an argument contains a conditional amongst its premises, the conditional is often regarded by the arguer as being supported by a **covering generalisation**. In reconstruction, we should take care to include a suitable covering generalisation where appropriate. Covering generalisations may be either hard or soft generalisations.

Not everything explicitly stated by an arguer is **relevant** to the argument. A proposition asserted by an arguer may be completely irrelevant; that is, the arguer does not advance that proposition as support for the conclusion. Such propositions should not be included in the reconstruction. In other cases, a proposition explicitly stated does provide some support for the conclusion, but is such that if it were ignored, the arguer would still, by virtue of other propositions put forward, have given an argument for the intended conclusion. If that proposition is false, then it should be removed from the argument.

Where possible, **ambiguity** and **vagueness** should be coped with by removing them. If it is possible to do so, the ambiguous or vague language should be replaced by language that is not ambiguous or vague. If a word in a premise or conclusion is ambiguous (can be read as expressing either of two meanings), and it is not clear which meaning is intended, then two reconstructions of the argument should be given, each reflecting one of the two meanings.

All **generalisations** in reconstructed arguments, whether hard or soft, should have explicit quantifiers. Where a premise is a generalisation, the **scope** of a generalisation should be as narrow as possible, but not so narrow that the inference to the conclusion is no longer valid (or inductively forceful, as the case may be).

Practical reasoning, or **means-end reasoning**, is embodied in arguments that specify an **outcome** as being either **desirable** or **undesirable**, along with an **action** said to be either necessary or sufficient for bringing about that outcome. Such arguments can take any of eight basic forms. In order to reconstruct such arguments as valid or inductively forceful, we often have to add a premise stating that the proposed action is the most efficient means of bringing about the outcome, and a premise stating that the benefit of the outcome outweighs the cost of the action (or, in the negative case, that the cost of an action would outweigh the benefit of the outcome). But sometimes the outcomes of actions are only probable. This may require us to calculate the **expected value** of an action.

Arguments are distinct from **explanations**, but many arguments have **explanations as conclusions**: such arguments attempt to establish which, of various possible causes of a given fact or event, is the actual cause. Arguments that attempt to

The practice of argument reconstruction

establish **causal generalisations** present more difficulty. It is common to see causal generalisations fallaciously inferred from mere **correlations**. Causal generalisations are not as informative as one might be tempted to think: what they tell us is that one type of event or state of affairs increases the probability of another type of event or state of affairs.

EXERCISES

1   As given, the arguments below are invalid. But each can plausibly be thought to contain either a generalisation or conditional as an implicit premise – a proposition that the arguer is assuming but has not explicitly stated. The arguments can easily be made deductively valid by making the implicit premise explicit. Do so: identify the conclusion and explicit premises, then add the premise needed to make the argument valid. Reconstruct the argument in standard form. Use a hard generalisation if you can think of one that is plausible; otherwise settle for a conditional. You should write all conditionals in the 'if–then' form.

a   Mr Bean is an idiot. You shouldn't marry him.

b   All men are idiots. You shouldn't marry one.

c   Johnny will like the Blandings Castle books; he likes the Jeeves books.

d   Pavarotti is going to sit in that chair! It's going to break!

e   Prices are not going to rise, because the savings rate is not decreasing.

f   Hakkinen won't win unless Schumacher's car breaks down. So Hakkinen won't win unless Schumacher's mechanic is inept.

g   Wagner was greater than Verdi, so obviously Wagner was greater than Puccini.

h   Since only democracies are just, no socialist countries are just.

i   If he doesn't accept the offer, then we will either withdraw the offer or raise it. If we raise it, then we will incur financial risk. Therefore, if he doesn't accept the offer, then either we give up trying to acquire another company, or we will incur financial risk.

j   Unless John brought wine, there isn't any. John didn't bring wine. So everyone will drink water.

k   Cigarette advertisements don't encourage people to smoke more?
     Ha. I think it safe to say that ads for chocolates encourage people
     to eat chocolate.

---

2   Find appropriate covering generalisations for the following
conditionals. Choose a hard or soft generalisation as seems
appropriate.

a   If John is under 18, then he cannot legally purchase alcohol in the
     UK.

b   If that is a scorpion, then it is poisonous.

c   If that wine is not French, then it is probably not overpriced.

d   That picture is not likely to be an oil painting if it is not Dutch or
     Flemish.

e   He is good at analysing arguments if he is a successful lawyer.

f   If the patient is now haemorrhaging, then his blood pressure is
     decreasing.

3

a   Suppose you know that Doctor Bowes does not own or drive a
     new Mercedes-Benz. Reconstruct the argument in accordance with
     the Principle of Charity.

         Doctor Bowes, the candidate, drives a new Mercedes-
         Benz. Furthermore, he has over £6 million in a personal
         Swiss bank account. Of course he is wealthy.

b   The following passage includes the assertion that the defendant
     has been addicted to cocaine for two years. Suppose you know
     this assertion to be entirely false (that is, the defendant has never
     used any illicit drug). Reconstruct the argument in accordance with
     the Principle of Charity. In order to reconstruct it, you will have to
     rephrase some of the sentences, ignore some material, and also to
     make implicit premises explicit. The conclusion is also implicit.

         The evidence is very compelling. The defendant has been
         addicted to cocaine for over two years. Her blood and
         fingerprints were found on the murder weapon. She was
         seen emerging from the victim's flat not more than half an
         hour after the murder took place. She boarded a flight for

Greece only twelve hours after the murder. And finally, the victim had recently ended the sexual relationship between himself and the accused.

---

4   We very frequently hear of claims or actions being criticised on the grounds that they are *politically motivated*. Sometimes this is put by saying that someone is 'playing politics', or using something as a 'political football'. What does this mean? Why is it a criticism? Find examples in the print media, and try to explain what the point is, in such cases, in calling something 'political'.

---

5   Reconstruct the following arguments, taking care to eliminate vague or ambiguous terms. You will have to rewrite sentences, ignore some material, and make implicit propositions explicit.

a   Ms Jones has demonstrated her commitment to feminism by supporting an across-the-board pay rise for female academics. If she is elected MP, then no doubt she'll support a reduction of the evidential standard for rape convictions.

b   Make no mistake: whatever their keepers say, these so-called 'domesticated' wolves are *wild*. Wild animals are too dangerous to be kept as pets.

c   Researchers have found that heroin use in teenagers is linked to parents with histories of depression. Parents under psychiatric care for depression should therefore be told the symptoms of heroin use.

---

6   For this argument, reconstruct it twice: once retaining the vague term 'political correctness', and once eliminating it.

When the Bunbury Women's Group proposes that the City Council bar men from the 'Women's Safe House', we realise that we must not allow them to have influence on the council. For this makes it obvious that the virus of 'political correctness' has infected the women's group, and we know what 'political correctness' stands for – they would promote such horrors as homosexuality being taught in schools and legal prohibitions against language deemed 'incorrect'.

7   Reconstruct the following arguments, making implicit propositions explicit as needed. If the argument is not deductively valid, increase its inductive force by decreasing the scope of the generalisation. If the argument is deductively valid, increase its chances of being sound by decreasing the scope of the generalisation.

a   Of course your yearling can be trained without much difficulty. Most horses can be.

b   Most men marry. Therefore most men have, have had or will have a mother-in-law.

c   No primates can learn to talk. Bobo is a chimpanzee, hence a primate; therefore he cannot learn to talk.

d   All countries can be attacked by sea. If so, then all countries require naval defence.

8   Reconstruct the argument about compulsory military service discussed at the beginning of this chapter (pp. 155–156). Take into account the advice given in the section on practical reasoning.

9   Look at argument-pattern (1) in the section on practical reasoning. In order for such arguments to be valid, we need to add two premises: first, that the benefit of increasing the amount of chocolate outweighs the cost of performing action X; second, that action X is the most efficient way of increasing the amount of chocolate. What similar premises are needed to make argument-patterns (2)–(8) valid?

10 The following argument contains a fair amount of rhetoric, stage-setting and explanation, but also a practical argument. Decide which of the eight patterns of practical reasoning the central argument of the passage fits, and reconstruct it. Don't worry too much about including every point that might be relevant; your main task should be to get the central argument laid out before you (then you can add sub-arguments, if you find any). Part of the first paragraph is concerned to reply to a point made by the correspondent referred to; think carefully about whether, and in what way, this contributes to the writer's intended argument.

The practice of argument reconstruction

> I despair of your correspondent who responded to the article on otters and mink (Letters, 9 July). Mink are certainly not playful and delightful animals. She must be confusing them with ferrets. While I agree it is no fault of their own that they are with us, it is not a question of survival. Mink are in paradise here in the wild: no harsh winter; no natural predators, a vast range of prey.
>
> In the Mustelid family they have found a niche halfway between otters and pine martens with the aquatic ability of the former and the natural aggression of the latter. Mink are consequently able to reach nesting colonies of ducks and seabirds, formerly secure from predation, wiping them out and rendering the island sanctuaries untenable by birds. I would invite your correspondent to come to Argyll and see for herself. Mink are a scourge on our environment and should be eliminated at every opportunity.[7]

11 Another practical argument, this time giving you more practice identifying implicit premises and conclusion. Reconstruct according to the instructions on practical reasoning. The letter is humorous, but contains a serious (though implicit) argument.

> Sir, Having read your report on how the police plan to test drivers for drugs, I and both my children in their early twenties attempted the tests described, without having taken any illegal substances.
>
> We consistently failed the second test (on one leg, head tilted back, eyes closed, other leg off the ground, arms extended, touch nose with each index finger). Swaying and giggling helplessly at each other's attempts, we would have given every appearance of intoxication.
>
> We have therefore decided that, if we are stopped and invited to perform these tests at the roadside, we will first ask the officer for a demonstration.[8]

12 A friend offers you a wager: if you draw an ace out of a (normal) pack of playing cards, he'll give you £10. If you don't, you pay him £1. What is the expected value of accepting the wager? Should you do it?

---

7   Michael Murray, *The Independent*, 16 July 2000.

8   John Tayler, *The Times*, 4 August 2000.

13 There is a sweepstakes draw: all you have to do is post a letter with your name and address to the address given, and your name will go into the draw. The lucky winner, drawn from a barrel containing 400,000 names, will receive £100,000! Assuming that a second-class stamp costs 26p, what is the expected value of posting your letter, thereby entering the sweepstakes?

14 The following are not simply explanations, but rather arguments that have explanations as conclusions. Reconstruct them, eliminating extraneous material and making implicit material explicit as needed.

a   At your last visit I said that the pain in your abdomen is caused either by a kidney infection, a musculoskeletal injury, or cancer in the pancreas or liver. But no indicators for a kidney infection showed up in the urinalysis, and if it were musculoskeletal, the pain would have subsided, not grown, by now. I'm afraid you have cancer.

b   The stress of modern life is not because we work more than our parents did – we don't – but because we no longer go to church. Instead, we shop: the world's most stressful activity.

c   Nietzsche's going irrevocably mad was caused either by his tortured intellectual life – as romantically-minded people would like to believe – or by the syphilis he contracted as a young man. But if it were the former, then we should expect a great many other notable intellectuals to have gone irrevocably mad. The truth – again contrary to romantic fantasies – is that very few have.

d   We read with horror that the age of the onset of puberty in girls is getting younger. We don't believe this has to do with artificial hormones in the milk or anything like that – surely in these more environmentally aware times, there are fewer artificial hormones and suchlike in our food than there were say, twenty years ago. The real reason is the unprecedented onslaught of sexualised images to which children are now subjected, especially on television.

e   Mr Jenkins blames the high unemployment rate on Britain's high interest rates: these strengthen the pound, making British exports uncompetitive, which in turn forces British manufacturers to cut costs by means of redundancies. We blame it on the ease with which people out of work can go on the dole. Jenkins' argument assumes that the non-manufacturing sector of the economy is not growing as fast as the manufacturing side shrinks. Manifestly, it is.

The practice of argument reconstruction

f   The president says that greenhouse gases are mostly produced by trees, not by industry. That makes sense. Global warming is on the rise, and we all see how, all over the world, people are busy planting trees and closing down factories, getting rid of their cars, and so on. The world of our pre-industrial ancestors must have been a real hothouse.

---

15 Review the section on causal generalisations. Every example below is a case either of (1) assuming that a causal relation entails a hard generalisation or a generalisation of the 'most' type; or (2) too readily inferring a causal relationship from a statistical correlation. In a short paragraph, criticise the following arguments: say whether each argument is a case of either (1) or (2) and explain your answer.

a   It is just not true that exposure to the sun causes skin cancer. If it did, then everyone who has ever had a suntan or a sunburn would get skin cancer.

b   I am so tired of hearing that social and economic deprivation 'cause' teenage criminality. Plenty of teenagers suffer from social and economic deprivation without turning to crime.

c   Taking lots of vitamin C makes my colds go away. I always take vitamin C when I get a cold, and it always goes away.

d   A causal link between violent video games and juvenile violence has been demonstrated. A study has shown that the incidence of violence is much higher amongst teenage boys who play violent video games regularly than it is amongst those who do not.

e   Oxbridge representatives claim that their admission procedures are not biased against students from poor backgrounds. How can they maintain this obvious falsehood, when *their own figures* show that the proportion of Oxbridge undergraduates from impoverished backgrounds is far lower than the proportion of applicants from impoverished backgrounds? Oxbridge continues to exclude the poor because they are poor.

---

16 Reconstruct the following arguments, and say whether they are valid, inductively forceful, or neither.

a   The sharp increase in the death rate in ancient Antioch, in the year AD 364, has been ascribed to many factors: to famine, to an influx of disease brought by soldiers returning from the Persian

campaign, and to problems with the city's water supply. It was probably the water supply: if it were food shortages, it is unlikely *both* that no contemporary historian mentions it, and that there should be no record of similar sufferings in the lesser towns of the area or the surrounding countryside. And if the soldiers had contracted disease in Mesopotamia, then, since they were seriously weakened by hunger and the strain of a long and arduous campaign, they would have died of disease in significant numbers. And then it is certain that Ammanius, an eyewitness, would have recorded those fatalities – so cruel after the humiliations and losses they had endured in battle – in his history. But he is silent on the point.

b Jepsen's Neural Syndrome – JNS – seems to be caused in dolphins by an excessive intake of heavy metals. A genetic cause is ruled out, because if the cause were genetic, then the correlation discovered between JNS and the presence of abnormally high levels of heavy metals in dolphins would be extremely unlikely. True, the correlation between JNS and abnormally high levels of heavy metals is no greater than that of JNS with the consumption of large amounts of squid. But squid, as we know, tend to retain, over time, large concentrations of the trace heavy metals that pass through their bodies. Dolphins tend to eat the larger, therefore older squid. It is very probable that if squid did not retain heavy metals, then the correlation between JNS and squid consumption would not exist. Since there are no other significant correlations with JNS, the heavy metals diagnosis looks highly probable.

c If alcoholic *behaviour* in parents – as opposed to genetic predispositions to alcoholism – caused alcoholism in children, then we would find that children whose parents are not alcoholic, and who are adopted at a very early age by alcoholic parents, were significantly more likely to become alcoholics. But our study shows that they are not. Therefore mere heavy drinking in parents does not, despite the well-known correlation, cause children to grow up to be alcoholic.

Chapter 6

Issues in argument assessment

Rational persuasiveness

It is plain common sense that the role of an argument is to give us reasons for accepting its conclusion as true. But we have not quite defined what it is, exactly, for an argument to do this. You might think that we have done this with the notions of deductive and inductive soundness, but that is not quite right. The reason is that even if we have reconstructed an argument perfectly, we cannot always tell whether or not the argument is sound. And that is because a sound argument must have true premises. A deductively sound argument is one that has true premises and which is deductively valid; an inductively sound argument is an inductively forceful argument with true premises. Since we do not always know which propositions are true and which false, we cannot always tell whether an argument is sound or not.

What we want, rather, is to characterise the property of 'giving us good reasons to accept the conclusion' in such a way that when an argument has that property, we are always able to recognise it – just by virtue of whatever knowledge we do possess, along with our ability to assess arguments logically. This property cannot be the property of validity or inductive force, since the fact that an argument is valid or forceful does not give us a reason to accept its conclusion unless we already have reason to accept the premises. But neither can this property be the property of soundness, since being able to recognise that normally requires that we have factual knowledge that goes far beyond the ability to analyse or reconstruct arguments.

Here is a very simple illustration. Suppose you are wondering whether or not you should expect that interest rates will increase in the next year, and someone presents you with the following argument:

P1) Interest rates will neither decrease nor remain unchanged
       in the next year.

C)   Interest rates will increase in the next year.

Clearly this argument is of no use to you. It is logically correct – it is deductively valid – but such an argument could not persuade anyone of the conclusion who did not already accept it. For if you did not know whether or not to expect interest rates to increase, then of course you didn't know whether to expect them to decrease, or to remain unchanged. If you had had some reason to accept P, then you could not have been wondering whether to accept C.

Here is another illustration. Suppose Bert drops a coin into a cup of coffee, and it settles to the bottom 'heads up'. Of course neither Bert nor Barney can see it, so neither knows that it landed heads up. Then Bert and Barney consider the following argument:

P1) If the coin landed tails down, then it landed heads up.
P2) The coin landed tails down.

C)   The coin landed heads up.

In the envisaged situation, the argument would be deductively sound, since it is deductively valid and the premises are true. But, although P2 is true, neither Bert nor Barney has any reason to think so. So, although the argument is deductively sound, and Bert and Barney can see that it is deductively valid, it gets them no closer to knowing the truth-value of the conclusion.

In such a situation, we say that the argument is **rationally unpersuasive**. More exactly we must say it is rationally unpersuasive for Barney, and also for Bert. Why we must put it this way – why we must *relativise* the notion of the rational persuasiveness of an argument, can be seen from a variation of the story.

Suppose that Barney, but not Bert, knows that the coin is weighted in such a way that it virtually always falls tails down. Then Barney, but

not Bert, has a good reason to accept P2. Therefore Barney has a good reason to accept that the argument is sound, but not Bert. Therefore the argument is rationally persuasive for Barney, but not for Bert.

Before we give our official definition of rational persuasiveness, there is one last complication to be discussed. Looking over the example just given, you might think that an argument is rationally persuasive (for a person) if (i) it is inductively forceful or deductively valid and (ii) the person reasonably believes the premises. But this is not quite right.

For consider now the following example.

P1) Almost all residents of Inverness own at least one item of woollen clothing.
P2) Jane is a resident of Inverness.

C) Probably, Jane owns at least one item of woollen clothing.

This is certainly an inductively forceful argument. Suppose in fact you have good reason to accept both P1 and P2. Indeed, suppose you know with certainty that they are both true. Does it follow that you should accept C? Can you imagine a situation in which you know that P1 and P2 are true, but in which you could reasonably reject C? Yes. For suppose that, in addition to your knowledge of P1 and P2, you know that Jane is violently allergic to wool. In that case, you could quite reasonably expect C to be false. Someone who knew the truth of P1 and P2, but nothing else about Jane, could reasonably expect C to be true, but not you. The argument would be rationally persuasive for them, but not for you.

What we say, in such a case, is that the argument is **defeated for you** by other evidence that you have; you have, in effect, a more compelling argument for the falsity of the first argument's conclusion (the concept of evidence is discussed further in Chapter 7).

In particular:

> To say that an inductively forceful argument is *defeated* for a person is to say: the person reasonably believes the premises, but, nevertheless, reasonably rejects the conclusion.

An inductively forceful argument whose premises you have reason to accept is rationally persuasive only if your total evidence does not

defeat the argument for you. Note that the definition pertains only to inductive arguments. The reason for this will be clear in a moment. Thus our definition:

> **To say that an argument is *rationally persuasive* for a person is to say:**
> **(i)** **the argument is either deductively valid or inductively forceful;**
> **(ii)** **the person reasonably believes the argument's premises; and**
> **(iii)** **it is not an inductively forceful argument that is defeated for that person.**

It can even happen that an argument is rationally persuasive for you at one time, but not at a later time, because you have acquired new evidence which defeats the argument (by suggesting that the conclusion is false, more strongly than the argument suggests its truth). To be strictly accurate, then, we should speak of rational persuasiveness for a person at a particular time. For the sake of simplicity, however, we have omitted this complication from what is already a rather complex definition.

There are six further points to bear in mind as regards rational persuasiveness.

1 It is not possible for the conclusion of a *deductively valid* argument to be *defeated* by a person's total evidence – this is only possible for inductively forceful arguments. Thus condition (iii) of the definition of rational persuasiveness only applies to inductively forceful arguments. The reasons for this are:

a If you accept with good reason the premises of an argument which you recognise to be deductively valid, then you must accept the conclusion as well; for you know that if the premises are true, then the conclusion must necessarily be true also. If a person has good reason to accept the premises of a deductively valid argument, then we may conclude that it is rationally persuasive for that person without further ado. Any further consideration which, for that person, casts doubt on the conclusion, must equally cast doubt on the premises.

  Why is this so? Well, the answer is because it is part of the definition of a valid argument that if its premises were true, then,

necessarily, its conclusion is true. Remember from Chapter 2: validity has nothing to do with the actual truth of the premises – so even if you don't have reason(s) to accept one or more of the premises, the argument remains valid (this is just to say that valid arguments can have false premises). But *if the premises are true and the argument is valid – the conclusion must necessarily (no exceptions) be true*. So if you have reason to accept the premises then those very same reasons are equally strong reasons to accept the conclusion. The only way that a **deductively valid** argument can fail to be rationally persuasive is if a person is without reason to accept the truth of one or more of the premises (remember that since the argument is deductively valid, condition (i) has already been fulfilled).

b The adverb 'probably' (or similar term) before the conclusion of an inductively forceful argument allows that it is possible that the premises are true and the conclusion false. When you have an argument that is inductively forceful, the third condition of the definition of rational persuasiveness becomes relevant; for the evidence that defeats the conclusion is precisely the evidence that shifts the balance of probability. Because inductively forceful arguments only claim that if the premises are true the conclusion is likely to be true, there is room for situations in which the conclusion is actually false.

Condition three of the definition of rational persuasiveness means that there are two ways in which an inductively forceful argument can fail to be rationally persuasive. The first is when you have no good reason to accept one or more of the premises (i.e. it fails condition (ii)). The second is when the argument is defeated for you: you have evidence that means you can't accept the conclusion even though you accept that the premises are true and the argument is inductively forceful.

2 It is not part of the definition that the argument be *sound* (either deductively or inductively). The reason is simple: the notion of rational persuasiveness is intended to capture what it is about an argument that makes it acceptable to a person. It is the notion of an argument's giving a person good reason to accept its conclusion. Indeed, an argument can have a false premise, and hence be unsound (neither deductively nor inductively sound), yet be rationally persuasive for a person. To illustrate, take the case of the weighted coin again. Suppose that on this one rare occasion, the coin settles to the bottom of the cup heads down, so P2 is false. Still, in that case, Barney would still be quite reasonable in thinking that P2 is true (for he knows that the coin is weighted). So the

argument would still be rationally persuasive for Barney, despite its having a false premise, and thereby being unsound. He would be quite right – in the sense of being rationally justified – to accept the conclusion of the argument, even though, in fact, it is false.

This point about rational persuasiveness illustrates an important fact to which we will return later in much more detail: a person may reasonably believe a proposition which is, as it happens, false. In other words, there is such a thing as a reasonable mistake. This is an elementary point, but it is an easy one to forget – and hence, for example, inappropriately to blame people for their mistakes. Indeed, a useful way to put the point is to say that the word 'mistake' is ambiguous. To say that someone is 'mistaken' could mean either (i) that they have accepted a false conclusion, or (ii) that they have been persuaded by bad reasons – by an argument which is not in fact rationally persuasive for them (or have failed to be persuaded by good reasons – by an argument that *is* rationally persuasive for them). Clearly we are responsible for our mistakes of type (ii); we ought to be persuaded by good reasons, and not by bad reasons. If we fail in this, then, typically, we are blameworthy. But it is much less clear that we are responsible for mistakes of type (i). If someone believes a proposition on the basis of good reasons – on the basis of arguments that are rationally persuasive for him or her – then it might just be bad luck if the proposition turns out to be false. If so, then the person need not be to blame for having made a 'mistake'.

Suppose, for example, that you are a doctor. Suppose there is a drug X which you know to have cured a certain dangerous disease every single time it has been used, in over a million cases; furthermore, you know that it has never had a negative side effect, and contains no substances known to be dangerous in any way. A patient has the disease and you prescribe X. Unfortunately, instead of being cured by X, the patient is made ill by it. Were you mistaken in believing that X would safely cure the disease? Not if by 'mistaken' we mean a mistake of type (i), and hence not in the sense required if you are to be blameworthy. Indeed, you would have been open to criticism if you had *not* prescribed the drug, since you would have been going against a massive body of evidence.

3   It should be appreciated why we have named rational persuasiveness as we have. Remember that at the beginning, we said that there are various kinds of attempts at persuasion. This book is about trying to distinguish argumentative from non-argumentative – chiefly rhetorical – attempts at persuasion, and learning to evaluate them. An attempt at persuasion by argument is an attempt at rational persuasion, as opposed to other kinds of persuasion, which do not appeal to your reason (but rather to your emotions or prejudices). Directed at you, it is an attempt

at providing you with a rationally persuasive argument. Note, again, that this does not mean that the argument must be a sound argument. The attempt depends upon what you reasonably believe regardless of whether or not those reasonable beliefs are in fact true.

4 Rational persuasiveness is a matter of degree; it is not all or nothing. Of two arguments for the same conclusion, one may be more rationally persuasive for you than the other, even if both are rationally persuasive. For example, suppose you have good reason to believe that if the barometric pressure is low, then rain is likely; you also have a slight reason to believe that if the wind is from the west, then rain is likely. Suppose you know that the barometric pressure is low, and that the wind is from the west. So you have two arguments for the same conclusion. But the one based upon the barometric pressure is more inductively forceful. In that case that argument is more rationally persuasive for you than the other.

5 'Rationally persuasive' does not mean merely 'persuasive' or 'convincing'. A rationally persuasive argument may fail to persuade anyone. Whether or not an argument is rationally persuasive for you does not depend upon whether you think it is. The crux of the matter is to understand this: *an argument may be rationally persuasive for you even though you are not persuaded by it*. This should not be regarded as paradoxical. All it means is that there are cases where you ought to be persuaded by an argument, but you are not. Likewise, there are cases where you are persuaded or convinced by an argument, but where you should not be, because the argument is not actually rationally persuasive for you. It is the task of rhetoric to cause people to overestimate the rational persuasiveness of an argument – to convince or persuade people without actually giving them good reasons.

To understand why this is so, examine Figure 6.1. You can see from this figure that when we examine the relationships between sound arguments, rationally persuasive arguments and the arguments that actually persuade people, it is possible for a given argument to be one, two or all three of them. The important point you need to note is that rational persuasiveness and soundness are properties that arguments can have independently of whether an individual or group actually finds them persuasive. Human beings are often irrational and part of the function of the concept of rational persuasiveness as we have defined it is to acknowledge this. If rational persuasiveness were defined so as to make all rationally persuasive arguments ones that actually persuaded people, we would have no way of accounting for the fact that sometimes people fail to be persuaded by arguments that they really should be persuaded
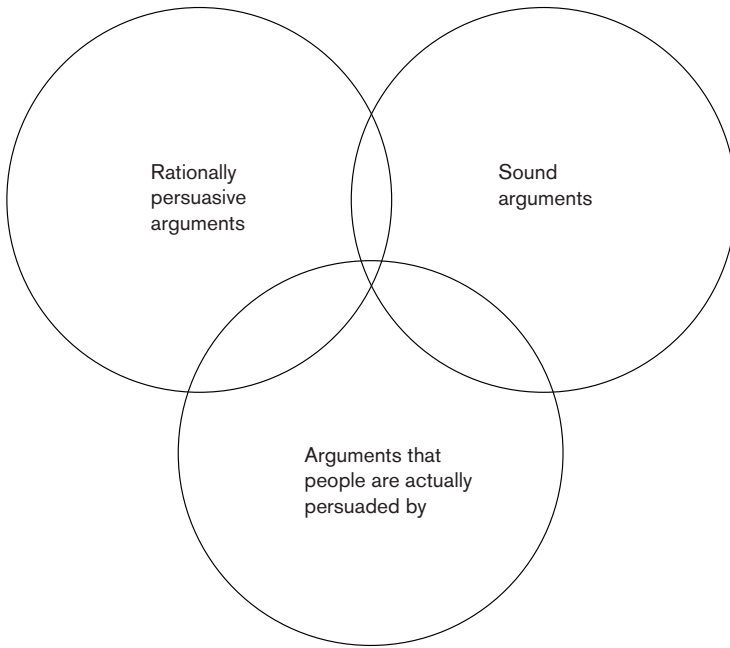
Figure 6.1 Soundness, persuasiveness and rational persuasiveness

by, and are sometimes persuaded by arguments that they should not be persuaded by.

There are three ways in which one can be mistaken about the rational persuasiveness of an argument. Study them carefully, especially if the preceding paragraphs are mysterious to you.

First, we can make mistakes concerning whether or not an argument is valid or inductively forceful. Especially in the case of a logically complex argument, we may for example think the argument deductively valid or inductively forceful when it is neither (for example when the argument contains a fallacy which is also an effective rhetorical ploy; see Chapter 4). If so, then even if we accept the truth of the premises, and are perfectly justified or reasonable in doing so, we may think that the argument is rationally persuasive for us when it is not. Equally, if we think an argument invalid when it is valid, or inductively unforceful when it is inductively forceful, it is possible to think an argument rationally *un*persuasive for us when in fact it is rationally persuasive for us.

Second, we can be mistaken in thinking we have good reason to accept a premise as true. For example, consider the following argument:

Issues in argument assessment

P1) Next summer will be a hot one.
P2) Hot weather is good for Cabernet Sauvignon grapes.

C) Probably, next year's Cabernet Sauvignon harvest will be good.

With suitable connecting premises added (for example: 'If next summer's weather is good for Cabernet Sauvignon grapes, then, probably, next year's Cabernet Sauvignon harvest will be good'), this becomes an inductively forceful argument, and P2 is true. Now suppose Jane knows that P2 is true. She also thinks that P1 is true, but her reason is that she has the superstitious belief that whenever it rains on Winter solstice, the following summer will be hot (assume it really is unreasonable for her to accept P1; as we will see in more detail later, what are often called 'superstitious' beliefs are not always unreasonable). In that case, assuming it did rain on Winter solstice, she may think her belief in P1 is reasonable or well-founded, but it is not. So in that case, having noticed its inductive forcefulness, she may well think that this argument is rationally persuasive for her, but she would be wrong. She may be persuaded, but she is not rationally persuaded.

Third, you can be mistaken about whether or not an argument is defeated for you. On the one hand, you might accept an argument for a given conclusion – thereby accepting the argument as rationally persuasive – without realizing that you have information sufficient to construct another argument that defeats the conclusion of the first argument. You might, for example, accept an argument for the conclusion that Mr Jones is at home, momentarily forgetting that you know that Mr Jones boarded a flight for Singapore yesterday. On the other hand, you might think that an argument is defeated for you when it isn't. You might, that is, mistakenly think you have good reasons to reject the conclusion of an argument that is inductively forceful and whose premises you accept.

Obviously one can give arguments for various purposes, including deceitful ones. But the rational, non-deceitful motivation for giving an argument is surely this: *It is to give a sound argument which is rationally persuasive for oneself, and for its intended audience.* Different arguments may be required for different audiences. We cannot always know with certainty whether an argument is sound, but that is the human predicament. It is simply a consequence of the fact that we do not always know with certainty which propositions are true and which false. If we did, we would never need arguments.

6   In saying that an argument is rationally persuasive for a person only if the person reasonably believes the premises, we are not requiring that the person have, at his or her disposal, further arguments with those premises as conclusions. What we are requiring is that the person be *justified* in accepting the premises. Justification is a wider concept than rational persuasiveness: if one has a rationally persuasive argument for a proposition then one is justified in accepting it, but one may be justified in accepting it by means other than argument. In particular, some beliefs, especially perceptual beliefs such as 'I see a dog in front of me', are often justified – reasonable – even though they are not inferred from other beliefs. We will return to this in the next chapter.

Some strategies for logical assessment

Once we have an argument represented in standard form, we have to pronounce whether or not the argument is valid, and if not, whether or not it is inductively forceful. To do this, the basic technique is simple: ask yourself, can I imagine or conceive of a situation in which the premises are true, but the conclusion false? If under no conceivable situation could that be, then the argument is valid.

   If you can think of ways in which the premises would be true but the conclusion false, then you must determine to what degree, if any, the argument is inductively forceful. What you do here is to imagine various situations in which all the premises are true. In these situations, is the conclusion more likely to be true than it is to be false? If the situations in which the conclusion is true would be more likely than those in which it is false, then the argument is inductively forceful; if not, not. If it is forceful, it remains only to specify the degree to which it is so.

   Here is an illustrative example:

P1)   The boy next door's pet is a mammal.
P2)   The only mammal that lays eggs is the platypus.

C)   The boy next door's pet does not lay eggs.

This is not deductively valid, since if the boy's pet is a platypus, then if it is female, it might lay eggs; that is certainly conceivable. Is it inductively forceful? Not as it stands, since the argument says nothing about the likelihood of the pet's being a platypus. If the premises are true, then

either the animal is a mammal but not a platypus, or it is a platypus. In the first case the conclusion is true; in the second, it would be false if the platypus were female. But of course it is highly unlikely that the boy next door keeps a platypus, since very few if any are kept as pets. So the first situation is highly probable if the premises are true. Therefore if we add the premise that it is highly unlikely that the animal is a platypus, the argument is inductively forceful (to a high degree). Making this explicit, we get:

P1)   Extremely few platypuses are kept as pets.

C1)   (Very probably) the boy next door's pet is not a
        platypus.

P2)   The boy next door's pet is a mammal.
P3)   The only mammal that lays eggs is the platypus.

C2)   (Very probably) the boy next door's pet does not lay eggs.

Whenever we find that an argument is not valid, we should always ask whether there are premises that (1) the arguer could reasonably be expected to know, or which we know to be true, and (2) would make the argument inductively forceful, if added.

Arguments with conditionals or generalisations as conclusions

The process described above, where we ask ourselves whether there is a possible situation in which an argument's premises are true but its conclusion false, can be mentally taxing, where the argument's conclusion is a conditional or a generalisation. Recall the argument from Chapter 5 concerning the tuna industry:

P1)   Tuna catches have been decreasing significantly for the
        past nine years.
P2)   If tuna catches have been decreasing significantly for the
        past nine years, then the tuna population will vanish
        unless the tuna industry is regulated more stringently.

C1) The tuna population will vanish unless the tuna industry is regulated more stringently.

P3) If the tuna population vanishes, the tuna industry will collapse altogether.

C2) If the tuna industry is not regulated more stringently, it will collapse altogether.

Concentrate on the argument from C1 and P3 to C2. C1, you may recall, is equivalent to

C1) If the tuna industry is not regulated more stringently, the tuna population will vanish.

It might be a bit difficult to imagine C1 and P3 true, and then ask whether C2 could be false. But in fact, when we considered this argument earlier, we used a method which makes it easier. We shall now make this more explicit, as it is very important when considering arguments that have as much logical complexity as this argument (or more). Note that C2 is a conditional. What a conditional asserts, roughly, is a certain relation between the antecedent and consequent – that if the antecedent is true then so is the consequent. So the question we want to answer is this: *if* the premises of the argument were true, then would this purported relationship hold? To answer this question we suppose, not only that the premises of the argument are true, but that the antecedent of the argument's conclusion is also true. And then what we want to know is whether, under those suppositions, the consequent of the argument's conclusion would also have to be true.

So here goes. Suppose C1 and P3 are true, and that the tuna industry is not (going to be) regulated more stringently (this is the antecedent of C2). We can now reason as follows (using our easier-to-understand, 'if–then' version of C1): we have assumed that the tuna industry is not regulated more stringently; therefore, according to C1, the tuna population will vanish. But in that case, according to P3, the tuna industry will collapse altogether. But that is precisely what the consequent of the C2 says. So we see that if C1 and P3 are true, then it does follow, from the antecedent of C2's being true, that the consequent of C2 is true also. So the argument is valid.

The same sort of technique is useful where the conclusion of an argument is a generalisation. Consider this argument.

Issues in argument assessment

P1) Every striker on the Italian national side is a good defender.
P2) No one is a good defender who does not tackle well.

C) Every striker on the Italian national side tackles well.

Here, as before, we begin by supposing the premises true. But in this case the conclusion is a generalisation. Now recall from Chapter 5 that generalisations can typically be regarded as generalised conditionals. So the conclusion C can (somewhat awkwardly) be reworded as: 'If someone is a striker on the Italian national side, *then* that person tackles well'. Thus suppose that P1 and P2 are true. What we do in this case is to suppose that someone is a striker on the Italian national side; that is, we consider any striker on the Italian national side. What we now want to know is whether P1 and P2 force us to conclude that that person tackles well. They do: according to P1 this person is a good defender; but then according to P2, this person must tackle well. But this was just *any* Italian striker, not any one in particular. So we can conclude that every Italian striker tackles well (if P1 and P2 are true). Now consider a slightly different argument:

P1) Every striker on the Italian national side is a good defender.
P2) Every player who tackles well is a good defender.

C) Every striker on the Italian national side tackles well.

If we imagine our Italian striker again, P1 entails that he is a good defender. But P2 does not entail that he tackles well. It says that every player who tackles well is a good defender, but it does not tell us that if he is a good defender, then he tackles well. So P1 and P2, in this argument, do not allow us to conclude that our Italian striker tackles well. The argument is invalid.

Supposing the conclusion false

Another way to assess the validity of an argument is to suppose the premises are true but the conclusion false. If we can see that this is impossible, then, according to the definition of validity, the argument is valid;

if we can see that this is possible, then we know that the argument is invalid. So consider the first argument about the Italian footballers. To suppose the conclusion false is to suppose that there is at least one Italian striker who does not tackle well. If he does not, then by P2 he is not a good defender. But then he is an Italian striker who is not a good defender, which contradicts P1. So if the conclusion is false, it is impossible for the premises to be true. So the argument is valid.

Do the same with the second argument about the Italian footballers. We now imagine an Italian striker who does not tackle well. According to P1, he must be a good defender. But P2 does not force us, or enable us, to draw any further conclusion about this Italian striker. It tells us that if he tackles well, then he is a good defender, but it does not tell us anything about the case where he does not tackle well: it does not say that if he does not tackle well then he is not a good defender. So the falsity of the conclusion is perfectly consistent with the two premises being true. For all the premises say, our striker might be a good defender for other reasons – he might be very quick and energetic, good at clearances, and so on. So the argument is invalid.

This method can be used on any argument, not just those whose conclusions are conditionals or generalisations. But it is especially helpful in those cases.

Refutation by counter-example

We move now to an important technique for showing that an argument is invalid or inductively unforceful. Normally, you would use this when you have already determined that the argument is not valid or inductively forceful, and, in an argument commentary, you wish to explain why the argument is not valid or inductively forceful. It is grasped most easily by looking at an example. Consider this argument:

P1) Almost all heroin addicts were marijuana smokers before becoming heroin addicts.

C) Marijuana smokers tend to become heroin addicts.

This argument is sometimes given as a reason not to legalise marijuana (or cannabis). But it is definitely a bad argument. This can be seen very readily by comparing it with this argument:

Issues in argument assessment

P1) Almost all heroin addicts were milk drinkers before becoming heroin addicts.

C) Milk drinkers tend to become heroin addicts.

Obviously this argument would not give us a reason to outlaw the drinking of milk. Yet it has a true premise, just as the first argument, and more importantly, it embodies exactly the same reasoning as the first argument. In particular, both arguments assume that if almost everyone who does X did Y beforehand, then having done Y makes them more likely to do X (or, those who do Y tend to become people who do X). As the second argument above illustrates, that is clearly wrong.

If you were given the first argument, you could show that it is a bad argument by presenting the second as an example of the same reasoning. That is what we mean by refuting an argument by means of a counter-example. The second example is a counter-example to the belief in this form of reasoning.

Now in this case we have been able to lay bare the mistaken assumption upon which the mistaken reasoning rested. This can often be very useful, and according to our general policy of making everything explicit, we should make the assumption explicit, rendering the two arguments thus:

P1) Almost all heroin addicts were marijuana smokers before becoming heroin addicts.
P2) If almost everyone who does X did Y beforehand, then those who do Y tend to become people who do X.

C) Marijuana smokers tend to become heroin addicts.

And:

P1) Almost all heroin addicts were milk drinkers before becoming heroin addicts.
P2) If almost everyone who does X did Y beforehand, then those who do Y tend to become people who do X.

C) Milk drinkers tend to become heroin addicts.

In making this generalisation explicit, we make it explicit exactly what is wrong with the original argument. For we have now represented the argument as a deductively valid one. But in this way we represent the argument as very obviously unsound, owing to the falsity of P2 (P1 in each case, is true). The second argument shows this, because it is valid, yet P1 is true and C is false. According to the definition of validity, it follows that P2 is false.

The technique just discussed is this: to show that an argument is unsound, make explicit the assumed generalisation upon which it is based, in such a way that the argument becomes deductively valid (or inductively forceful). Then find a true premise and false conclusion which are suitably analogous to the premise and conclusion of the original argument, and substitute them (as we did with P1 and C of the above argument).

Sometimes an arguer's assumed, faulty generalisation is perfectly obvious, but they seem not to notice that it is faulty. Here is an example:

> The continuing carnage on our roads must be stopped. Since it will certainly reduce the number of fatal road accidents, I propose that we immediately reduce the speed limit on all roads (perhaps except motorways) to 60 mph.

We can reconstruct this, making the generalisation explicit, as follows:

P1) Reducing the speed limit to 60 mph would reduce the number of fatal road accidents.

P2) Anything which will reduce the number of fatal road accidents should be done immediately.

C) The speed limit should immediately be reduced to 60 mph.

This is deductively valid, but it can easily be refuted by counter-example:

P1) Reducing the speed limit to 5 mph would reduce the number of fatal road accidents.

P2) Anything which will reduce the number of fatal road accidents should be done immediately.

C) The speed limit should be reduced to 5 mph.

Issues in argument assessment

Since P1 is true and C false, and the argument is valid, P2 must be false. So the original argument, since its P2 is the same, is unsound. Of course, this does not mean that nothing should be done to reduce road fatalities. It means only that the fact that a given remedy would reduce them is not sufficient for carrying it out. As explained in the Chapter 5 section on practical arguments, we must, in such cases, show that the envisaged remedy would not cause worse problems, and also that no other remedy would be more practical and effective.

Engaging with the argument: avoiding the 'Who is to say?' criticism

Sometimes an argument will contain a premise which no one would say can be known with certainty. Sometimes these will have to do with what a particular person was thinking, or with what motives people in certain circumstances are likely to have, or with the future, such as whether unemployment will increase or the Labour Party will win the next election. Consider this argument:

> If they close down the factory – making over 500 workers redundant – then unemployment in our town will immediately double, at the very least. This will surely lead to increases in family tensions, hence in domestic violence. Studies have found this happening in every case where there is such an abrupt and dramatic increase in unemployment.

The conclusion is that if the factory closes, domestic violence will increase. Without bothering to reconstruct further, you can well imagine someone reacting to this argument with, 'Well, who is to say what will happen? You don't *know* that increased domestic violence would result; you can't just assume that the people in our town are as bad as that'. But this is an empty criticism of the argument. The critic has given no reason to suppose that the people of the town in question are different from other people. The critic is just ignoring the argument, not analysing it. Assuming the premises to be true, the evidence that an increase in domestic violence tends to follow large increases in unemployment is very strong. Perhaps this sort of thing cannot be absolutely certain until it happens, but that is the nature of inductive inferences. Where an argument is inductively forceful, the person who says 'who is to say?' that the conclusion is true is either repeating what nobody doubts – that the argument is not deductively valid – or expressing a seemingly unreasonable scepticism, like a person who refuses to believe that past observation supports the hypothesis that spring will follow winter.

Another variety of the 'who is to say?' response is illustrated in the following example:

P1)   Our finest works of art are irreplaceable assets.
P2)   It is the government's responsibility to protect our irreplaceable assets.

C)   It is the government's responsibility to protect our finest works of art.

Now the phrase 'finest work of art' is a vague term; no doubt the boundary between the finest works and the not-so-fine works is not clear. It is also what is sometimes called a 'value-laden' term, since what makes something a fine work of art depends at least partly upon whether, in fact, people do value it (more generally, a value-laden term is one whose application to things depends in some way upon our attitudes towards those things; for example, the word 'weed' is value-laden, since whether or not something is a weed depends on whether we like or tolerate it in gardens or amongst crops). So now suppose some opponent of the conclusion responds to this argument with, 'Well, who is to say what are our finest works of art?; it's a matter of opinion'. But this is not an effective criticism of the argument. It may well be that there is some difficulty in achieving unanimity on such a question. But in order effectively to criticise an argument – in order to engage with it – one must either (i) show that the argument is a bad one – neither valid nor inductively forceful – or (ii) show or argue that one or more of the premises is false; or (iii) show, if it is an inductive argument, that it is defeated by some other argument. Merely pointing out that a term occurring in the argument is vague or value-laden is not sufficient. Certainly remarking the presence of vagueness is not sufficient. For example, 'bald' is a vague term, yet it is inductively forceful to argue 'Joe is not over 25, therefore Joe is not bald'. Nor is remarking the presence of a value-laden term sufficient as a criticism.

To make this vivid, consider this argument:

P1)   Cold-blooded murder of innocent children is immoral.
P2)   One should not do what is immoral.

C)   One should not commit cold-blooded murder of innocent children.

Issues in argument assessment

This argument contains the words 'immoral', and 'innocent', which are certainly value-laden. It would certainly be possible for someone to say, 'Who is to say what is immoral, or what is innocent?'. But the fact that these terms can present problems of application does not detract from the evident soundness of this argument.

Of course, it might be that if the government policy suggested by the argument about works of art were enacted, it would be very difficult to apply it, owing to the difficulty of identifying 'our finest' works of art. That might even constitute an argument against such a policy. But in order actually to engage with the argument, the critic must actually produce such arguments; merely to point out that an argument contains a problematic or vague term like 'fine work of art' is not enough, as shown by the murder example.

Argument commentary

We have now completed our survey of our basic concepts and procedures of argument analysis. By 'argument analysis' we mean a two-stage process; comprising first the reconstruction, then the assessment of the argument. At this point, we need to say a bit more about the final product of this process. When the analysis of an argument is undertaken, you may sometimes want to produce a piece of written work which summarises the analysis you have made. This should consist of three parts:

1 The argument as originally expressed.
2 The argument expressed in standard form.
3 A *commentary* on the argument, written in ordinary prose.

Generating 2 out of 1 is the stage of argument reconstruction; most of what was said in the previous chapter of this book is addressed to various sorts of difficulties commonly encountered in making that step. The concepts discussed in Chapter 1 also bear upon the step from 1 to 2, because, as explained earlier, the Principle of Charity enjoins us to reconstruct the argument in the most favourable way, and we need those concepts in order to determine exactly what that means.

What we have not yet discussed is 3. What this comes to is simply a written piece of work that covers the following points:

(i) A discussion of how and why the standard-form reconstruction was derived as it was, focusing especially on any problems encountered in the process. If any implicit premises (or any implicit conclusions) have been added, it should be explained why. If the conclusion or any premises have

been reworded, it should be explained why (for example, it should be explained why an ambiguous sentence has been rewritten). In general, this section should include everything necessary to justify the given reconstruction. You have to explain why you have reconstructed the argument in the way that you have (but there is no need to state points which are simply obvious). Frequently, this will have to make mention of the likely intentions of the author of the original argument, and the context in which the argument was given.

(ii) A discussion of the validity or degree of inductive force of the argument. You have first to pronounce whether or not the argument is deductively valid. If it is not, you should explain why it is not (here, for example, you might use the method of refutation by counter-example). And if it is not, you should pronounce and explain to what degree, if any, the argument is inductively forceful. If the argument commits a fallacy, then you may identify it at this point, especially if the fallacy is a formal one.

(iii) A discussion and verdict concerning the truth-values of the premises. If the argument is deductively valid or inductively forceful, this will amount to a verdict regarding the soundness of the argument. It should be explained in detail which premises are most debatable, and why. Except where it is more or less obvious, these explanations must be **substantive**; actual reasons for accepting or doubting particular premises must be given. If the argument commits a substantial fallacy, then you would explain this here.

(iv) In the case of an inductively sound argument, you should also say whether or not the argument is defeated for you. An argument's being defeated for you is a fact about your relation to the argument, not a fact directly about the argument itself. By contrast, an argument's validity or inductive force (or lack of it), and its soundness (or lack of it), are matters concerning which there is a single right answer, a fact about the argument which is independent of the state of knowledge of particular people. Nevertheless, when our ultimate concern is with the truth-value of the conclusion of an argument, it is obviously relevant, when the argument is defeated for you, to point this out in your argument commentary. In other cases, however, we may be interested, not so much in the truth of the conclusion, but in the merits of the argument as given by the arguer. For example, Napoleon may have advanced an inductively sound argument for the conclusion that his army would prevail at Waterloo. That argument would be defeated for us, since we know that his army did not win. But if our concern were to assess Napoleon's reasoning, then still

we could not do so on the grounds that his argument is defeated for us. We would have to show that the argument was not rationally persuasive for *him* – or at least not sufficiently rationally persuasive to justify risking the battle.

Note that these instructions do not tell you to specify whether or not the argument is rationally persuasive for you. The reason is this. What you are trying to assess is the argument – the argument as given by the original arguer. When you assert that an argument is sound – deductively or inductively – the only way that you could make that assertion but not find the argument rationally persuasive is where an inductively sound argument is defeated for you. If it is not defeated for you, then if you confidently assert that an argument is sound – that is, if you know that it is sound and undefeated for you – then you are already showing that you find it rationally persuasive for you to a high degree. So it would be redundant for you to say that the argument is rationally persuasive for you.

However, in some cases your belief in the premises may not be strong enough for you to pronounce confidently that the argument is sound. Or you may have some suspicion that the argument is defeated for you, but not so much that you do not think the conclusion probable. In either of those cases, you would say that you think that the argument is probably sound, and believe it not to be defeated for you (in the case of an inductive argument). In these cases it would be appropriate for you to say that the argument is rationally persuasive for you, and to explain why, despite your lack of complete confidence, you hold the argument to be sound and not defeated.

In some cases we fail to have a belief either way with respect to an argument's premises. What we should do in such a case, rather than arbitrarily committing ourselves to a soundness verdict, is to say that we cannot determine it because we are ignorant of the truth-values of the premises (again, one is, in such a case, showing that one finds the argument rationally unpersuasive). We should also explain, if we can, what we would need to find out in order to alleviate that ignorance.

Ideally, it is helpful to keep the discussions (i)–(iv) separate. However, this is not always practical. For example, we may justify a given reconstruction on the grounds that it is inductively more forceful than another possible reconstruction (following the Principle of Charity). So we would be jumping ahead to (iii) in the midst of (i). That is all right. The main thing is to ensure that all three tasks are completed, with as much clarity as possible.

Argument trees

An argument tree is a device that can be used for representing arguments in the form of a diagram. They are helpful when we are reconstructing arguments, particularly complex ones, because they provide a means of showing the ways in which the different parts of an argument are related to each other. They show how the premises support the conclusion. It is great discipline to construct argument trees and you will find it helpful to use argument trees in your own analyses of complex real-world arguments.

Note: the process of constructing an argument tree is especially useful *before* you have supplied missing premises, and before you have settled upon a reconstruction of the argument in standard form. In fact, it can be useful to construct an argument tree at almost any stage in the reconstruction, including when the reconstruction is complete, simply as a way of illustrating the structure of the argument, of making it clear to oneself and to others.

Consider the following simple arguments:

A

P1) Susan is a marathon runner.
P2) Susan eats well and sleeps well.

C)   Susan is healthy.

B

P1) Willy is in the music club.
P2) No member of the music club plays jazz.

C)   Willy does not play jazz.

These arguments both infer a conclusion from two premises, but there is an important difference.

In argument A, each premise supports the conclusion individually. That is, P1 is cited as a reason for C, and P2 is cited as another reason for C. One could argue from P1 alone to C; one could also argue from P2 alone to C. By contrast, in argument B, neither premise supports C by itself. Neither P1 nor P2 would, by itself, be a reason to accept C. Rather, they work together to support C.

Issues in argument assessment
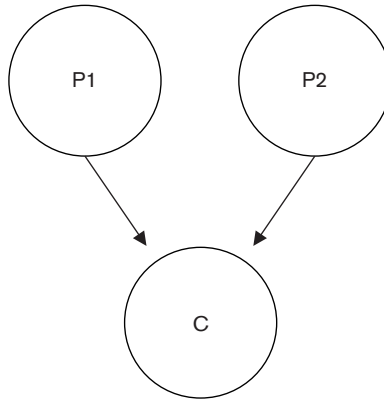
Figure 6.2 Two premises supporting a conclusion individually
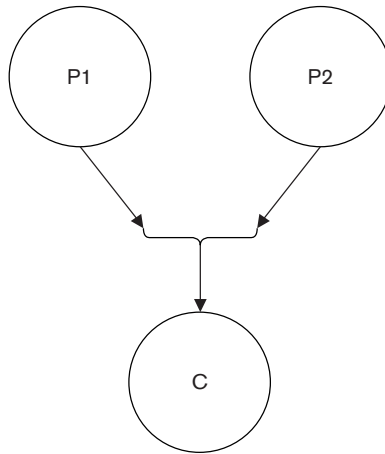
Figure 6.3 Two premises supporting a conclusion jointly

Argument A is represented as in Figure 6.2 and argument B is represented as in Figure 6.3.

Now these examples are very simple. The real value of argument trees emerges when they are used to represent more complex arguments with intermediate conclusions. For example:

P1)   Consumption is increasing.
P2)   The pound is weakening against other currencies.

C1)   Inflation will increase.

P3)   Whenever inflation increases, mortgage rates rise.

C2)   Mortgage rates will rise.

P4)   Whenever mortgage rates rise, the building trade suffers.

C3)   The building trade will suffer.
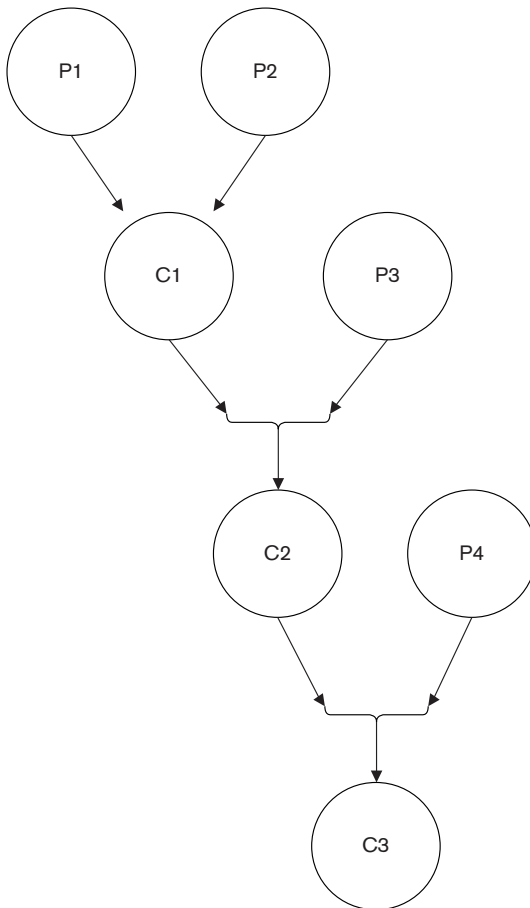
Figure 6.4 Extended argument

Issues in argument assessment

The correct argument tree for this would be as shown in Figure 6.4. Note that this is not a completely reconstructed argument. P1 and P2 do support C1, but we have not included the connecting premise or premises. The most plausible connecting premise would be of the form, 'If P1 and P2, then C1'. We had the same sort of situation in the argument about Susan the marathon runner. The connecting premise was left implicit. As plausible as any connecting premise, in this case, would be something like

P3) No marathon runner who eats and sleeps well is not healthy.

If we do it this way, then the three premises work together to support the conclusion; an instance of this generalisation would be: 'If P1 and P2, then Susan is healthy'. So the argument would be represented as shown in Figure 6.5.
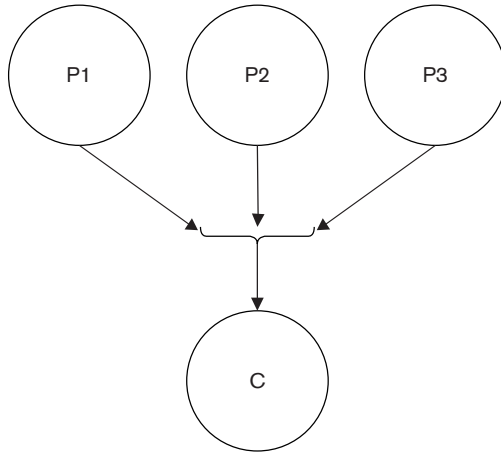
Figure 6.5 Three premises supporting a conclusion jointly

A deductively valid argument is **rationally persuasive** for you if you have good reason to accept its premises. An inductively force-ful argument is rationally persuasive for you provided that you have good reason to accept the premises, and it is not **defeated**: you do not have some other argument for rejecting the conclusion that is more rationally persuasive for you than the one in question.

Unlike validity, inductive force and soundness, rational persuasiveness is not a feature of arguments in themselves. It is also a matter of a given person's relationship to an argument. An argument may be rationally persuasive for one person but not for another. An argument may be sound but not rationally persuasive for you (for the premises might be true even though you lack good reasons to accept them). Moreover, an argument may be rationally persuasive for you but not sound (for you may have good reasons for accepting a set of premises, even when one of the premises is false). But the question of whether or not an argument is rationally persuasive for you is not simply the ques-tion of whether or not you find it persuasive, or whether you are in fact persuaded by it. For unlike the question of whether we are actually persuaded by an argument, we can be mistaken about rational persuasiveness: an argument may be rationally persuasive for you even when you think it isn't, and fail to be when you think it is. The importance of the concept of rational persuasiveness is that it captures what arguments are intended to do: The distinctive aim of persuasion by argument is to persuade people rationally, that is, by actually giving them good reasons to accept a given conclusion, and not just seeming to.

There are various informal methods of **logical assessment**, and of demonstrating that an argument is not valid or inductively forceful. The principle task is simply to apply the definitions of deductive validity and inductive force, but there are some strate-gies to facilitate this in more difficult cases. One strategy that is almost always pertinent is to suppose the conclusion of the argu-ment to be false, and then to ask whether it would still be possi-ble for all the premises to be true. Another strategy, where the conclusion of an argument is a conditional, is to suppose that the antecedent of that conditional is true, and then to determine whether the remaining premises force one to conclude that the consequent would also be true. If the conclusion of an argument is a generalisation, an analogous strategy is to suppose that the antecedent of an arbitrary instance of the generalisation is true. Finally, an effective way of demonstrating that an argument is

invalid, or inductively unforceful, is the method of **refutation by counter-example**. This involves giving an argument that employs the same type of inference as the one in question, but which is obviously invalid or unforceful.

It is important to avoid the **who is to say?** criticism. One form this takes is to complain that the conclusion of an inductively forceful argument has not been 'proved', that it 'might' be false. Another is to object to an argument on the grounds that it contains a value-laden term. It should be clear that these are not genuine criticisms from a rational point of view.

EXERCISES **1** For questions a–i, consider this argument:

P1) If Rangers won the match then the pub will sell pints for £1 tonight.
P2) Whenever Rangers are ahead at halftime of a Scottish Premiership match, it is highly probable that they win the match.
P3) Rangers were ahead at halftime, and this is a Scottish Premiership match.

C) (Probably) The pub will sell pints for £1 tonight.

Suppose that the premises are true. Suppose that both Andrew and James know that the premises are true (they have both seen the sign in the pub window about the £1 pints, they both saw the halftime score of 1–0 in favour of Rangers announced on television, and both know that Rangers have won every Scottish Premiership match in the past three years when leading at halftime). Andrew has no other information relevant to the truth of the conclusion; but unlike Andrew, James saw on television that the opposing side scored twice in the second half to win the match 2–1.

a Is the argument deductively valid?

b Is the argument deductively sound?

c Is the argument inductively forceful?

d   Is the argument inductively sound?

e   Is the argument rationally persuasive for Andrew?

f   Is the argument rationally persuasive for James?

g   Should Andrew be persuaded by the argument?

h   Should James be persuaded by the argument?

i   If the argument is not rationally persuasive for either Andrew or James, explain why.

---

2   Now consider this argument:

P1)   The majority of well-educated Germans speak English.
P2)   Jacob is a well-educated German.

C)   (Probably) Jacob speaks English.

Assume this time that P1 is true and that the conclusion is true; but assume that P2 is false: Jacob is Austrian, and not very well educated either (still, he did learn English quite well in secondary school). Assume that Catherine, Jane, Mary and Anna all know that P1 is true. Catherine believes P2 is true because her friend David told her it is; David is a reliable person who knows Jacob, and Catherine has no reason to doubt him. In fact David's mistake was quite reasonable: Jacob lied to him, telling him he completed a university degree in history; further, it was reasonable for David to infer, from Jacob's accent, that Jacob is German. For, although German is the national language of Austria, the vast majority of native German speakers are German. Catherine knows nothing else about Jacob, and accepts C. Jane also believes P2, but for different reasons: she fancies Jacob, having seen him, at a distance, at a party, and thinks he's well educated because she saw Jacob wearing glasses and a tie. This belief is clearly not well supported. Jane also accepts C, but has no other information relevant to the truth of C. Mary believes both P1 and P2, for the same reasons as Catherine; but she does not accept C, because she heard Jacob speaking German to David at the party, and inferred, quite reasonably, that Jacob does not speak English. She has no other information relevant to the truth of C. Anna does not believe P2. She heard what David said to Catherine, and believes that David is always sincere and well informed; but Jacob is athletic and handsome

and she has a stupid irrational prejudice according to which athletic, handsome men are almost always stupid, and therefore not well educated. She has no other information relevant to the truth of C.

a   Is the argument inductively forceful?

b   Is the argument inductively sound?

c   Is the argument rationally persuasive for Catherine? Why or why not?

d   Is the argument rationally persuasive for Jane? Why or why not?

e   Is the argument rationally persuasive for Mary? Why or why not?

f   Is the argument rationally persuasive for Anna? Why or why not?

---

3   Reconstruct the following arguments. Then, using the techniques discussed earlier, explain why the argument is, or is not, valid.

a   If this man is not a spy, then he is a detective. But all detectives wear trench coats, and he is not wearing one. If he is a spy, then either he is Russian or American. But no American spy knows how to order dessert wines, unless he is from New York. But all spies from New York wear trench coats. Therefore, if this man knows how to order dessert wine, then he is a Russian spy.

b   If the ancient Allemani had been both vicious and loyal, then the Romans would have purchased their allegiance. If they had done so, then the Franks would never have challenged the Allemani. But they certainly did. Therefore, if the Allemani were vicious, they were not loyal.

c   Some ancient Visigoths were literate. But all Roman citizens were literate. Therefore some ancient Visigoths were not Roman citizens.

d   If most educated Romans learned Greek, then most read Homer. But anyone who reads Homer knows the story of Achilles. Therefore most Romans knew the story of Achilles.

e   If all ancient Visigoths were pagans, then so were all ancient Picts. Some ancient Picts were non-pagans, or some ancient Saxons were. But some ancient Saxons were Christian, and no Christian is a pagan. Therefore some ancient Visigoths were non-pagan.

4   The conclusions of the first two arguments in exercise 3 are conditionals. Split them into two statements, antecedent and consequent. Now add the antecedent to the premises, and regard the consequent as the conclusion of the argument. Using the method of supposing the conclusion false, explain why the resulting argument is valid or not valid.

5   Reconstruct the following arguments, then refute them by counter-example – that is, by giving arguments that embody the same pattern, but which have true premises and false conclusions.

a   A significant increase in the rabbit population would bring about an increase in the number of foxes. And sure enough, the number of foxes has increased lately. This must, therefore, be due to an increase in the rabbit population.

b   If free will is impossible, then the concept of responsibility is nonsense. If so, then the justice system is based upon a confused idea. Therefore if free will is possible, then the justice system is not based upon a confused idea.

c   You and others have tried for years to prove that Nessie does not exist. But you have failed. You should by now admit it: Nessie is real.

d   You admit that the Loch Ness monster has not been seen in recent years. If so, then it must be very secretive. But if it is secretive, then, obviously, it exists. Therefore you have admitted that Nessie exists.

e   Why don't we do what we know will stop these evils? We should give drug-dealers and paedophiles life prison sentences.

6   Go back through the arguments in question 3, and draw tree diagrams for them.

7   List the *explicit* premises and the conclusions of the arguments in questions 13, 14 and 15 of Chapter 5. Draw tree diagrams for these, including only the explicit premises. Now add the needed implicit premises to the list, and draw new tree diagrams for the complete arguments.

Issues in argument assessment

8 Could one give an argument that (i) one knows to be sound (ii) is rationally persuasive for its audience, and (iii) is not rationally persuasive for oneself? If not, why not? If so, would that be deceitful, or otherwise naughty in any way? Why or why not? Make up an example.

Chapter 7

Truth, knowledge and belief

In this final chapter we will delve deeper into some philosophical issues underlying the principles of good critical thinking. We begin the chapter by extending the discussion of truth that appears in Chapters 1 and 2 and attempting to dispel what in our experience is a deep seated myth – that what is true depends upon nothing more than personal opinion or taste. As we will explain, this is the myth that **all truth is relative**. Later we examine the relationship between **believing** that something is the case, being **justified** in believing that it is the case and **knowing** that it is the case. Issues associated with these relationships are important for critical thinking because they are concerned with the **adequacy of evidence** for beliefs. Thus they are at the root of our determinations of whether or not premises are true, whether or not the conclusion of an inductively forceful argument is defeated for a person, and whether or not an argument is rationally persuasive for a person. Further, the concepts of truth, knowledge and evidence are such frequent sources of confusion in argumentative contexts that clarity about them is extremely valuable in its own right.

Truth and relativity

When we say something, claim something such as 'The kettle has boiled', we **assert** something; we express a **belief**. A belief is an attitude we take towards a **proposition**: to believe a proposition is to accept it as **true**. Assertion is a truth-claim, and belief is a truth-attitude. Assertion, belief, and truth are internally related in this way. From the outset we have been

working with an intuitive understanding of truth such that to say that a claim is true is simply to say that things are as the claim says they are: *To assert that a proposition is true is equivalent to asserting that very same proposition.* What this means is that a pair of sentences such as the following:

▶ Venus Williams was the Wimbledon Women's Singles Champion in 2000.

▶ It is true that Venus Williams was the Wimbledon Women's Singles Champion in 2000.

must have the same truth-value: if one of them is true, then so is the other. This necessary equivalence is the fundamental fact about the ordinary meaning of the word 'true'. Suppose, then, that Julie says that Venus Williams was the Wimbledon Women's Singles Champion in 2000. To say that Julie's claim is true, at bottom, is just to say that Venus Williams did win the Women's Singles at Wimbledon in 2000. Thus, although truth is a feature of claims that people make (of some claims, of course, not of all of them), whether or not a claim is true has nothing at all to do with the person who makes it; nor with that person's beliefs, culture or language (except when the proposition is explicitly about those things). Whether or not Julie's claim about Venus Williams is true depends only on whether or not Venus Williams won the Women's Singles at Wimbledon in 2000, and does not depend in any way on anything about Julie. In particular, the mere fact that Julie believes, and has claimed, that Venus Williams won Wimbledon, has nothing to do with whether or not her belief or claim is true.

Notice also the following consequence of the equivalence noted above. If John responds to Julie's claim by saying 'That's true', then what he does, in effect, is to assert the very same thing that Julie did. He agrees with her. Her claim is true if Venus Williams won Wimbledon and false otherwise; likewise with John's.

These points are quite straightforward, but they can be easy to lose sight of in other contexts, and these contexts can create confusion about truth generally. In order to dispel the myth that truth is relative, we first explain the concept of **implicit speaker-relativity**. We do so by considering some different types of **disagreement**. Consider the following claims:

▶ Chocolate ice cream tastes better than vanilla ice cream.

▶ David Blunkett is Home Secretary.

▶ Water is H_2O.

▶ Neptune is larger than Venus.

▶ George Clooney is more handsome than Brad Pitt.

▶ *Mission Impossible II* is a highly entertaining film.

▶ La Paz is the capital of Bolivia.

These are all declarative sentences. The claims are expressed using the same **assertoric** form. Assertoric form is generally used to express a belief that such-and-such is true (except when someone is lying or play-acting); but the fact that someone has made an assertion does not establish that the assertion is true, only that the speaker believes it to be so. Thus, if someone asserts 'Neptune is larger than Venus', they express their belief that one planet is larger than another, that it is a fact that the one is larger than the other. But whether or not the assertion is true depends only on whether or not Neptune really is larger than Venus. Similarly with 'David Blunkett is home secretary', and 'La Paz is the capital of Bolivia'. Note that although someone making one of these statements asserts their own belief, they are certainly not talking about themselves; someone who asserts that La Paz is the capital of Bolivia asserts their own belief, but the statement, and their belief, is only about La Paz and Bolivia.

Now look at the other claims which, as we have noted already, are expressed in the same assertoric form. Consider the first claim

Chocolate ice cream tastes better than vanilla.

It appears to assert a fact about the relative merits of chocolate and vanilla ice cream. But, unlike the cases considered above, the fact expressed is about the person who makes the assertion. This claim is, first of all, **implicitly relative**. A claim is implicitly relative when it involves a comparison or other relation to something it doesn't explicitly mention (see Chapter 1, pp. 28–29, for a review of this concept). For example, said of an adult man, 'John is tall' states a comparison, a relation, between John and other men; it really says that John is taller than the average man.[1] But furthermore, the claim about ice cream is implicitly **speaker-relative**. If Julie says 'Chocolate ice cream tastes better than vanilla', then what she is really saying is that *she* prefers chocolate to vanilla. Similarly, if John asserts this very same sentence, then he is saying that chocolate ice cream tastes better to *him* than vanilla does. So Julie and John are saying different things, despite the fact that they use the same sentence

---

1   Of course there is no such thing as 'the average man', as if, in addition to Tom, Dick, Harry and the rest, there were another chap, the average man. To say that John is taller than the average man is to say that if you take the average (the mean) of the heights of all men, then John's height exceeds that figure.

to say it. These two speakers' assertions are statements of fact about their respective preferences, not statements of fact about the superiority of chocolate over vanilla ice cream independent of anyone's preference.

This means that the sentence by means of which we express such a proposition – 'Chocolate ice cream tastes better than vanilla' – is an incomplete expression of the proposition we express by means of it. The statement actually expresses the preference of the person making it, but the sentence does not explicitly mention this; that is why the statement is **implicitly speaker-relative**.[2]

The importance of this emerges when we consider what happens when people appear to disagree over claims that are implicitly speaker-relative in this way, and compare this with genuine factual disagreement. Suppose now that Julie and John disagree about the capital of Bolivia. Julie says

La Paz is the capital of Bolivia.

But John denies it. He says: 'La Paz is not the capital of Bolivia' (perhaps he thinks it's the capital of Colombia). In this case, there is exactly one proposition – that La Paz is the capital of Bolivia – such that Julie asserts it and John denies it. That is what genuine factual disagreement is: genuine disagreement is when there is one proposition that is asserted by one person but denied by another. If Julie and John value the truth, they will want to know whose claim is true.

But suppose now that Julie and John also disagree about the relative merits of chocolate and vanilla ice cream: Julie contends that chocolate tastes better, while John is on the side of vanilla. As we have seen, in order to make her meaning perfectly explicit, Julie would have to say 'Chocolate ice cream tastes better than vanilla to me'; John would have to say the same in order to be explicit. So what Julie is really saying is that chocolate tastes better *to her*, John that vanilla tastes better *to him*.

The proposition that Julie expresses can equally well be expressed as:

(1) Chocolate ice cream tastes better than vanilla to Julie.

Whereas John's proposition is:

(2) Vanilla ice cream tastes better than chocolate to John.

These are two different propositions. There is certainly no logical conflict between them: they could both be true. But in that case, Julie and John

---

2   There is one slight complication, however: statements of this kind may mean that something is preferred or liked by most people. For example, this is plausibly what someone means who says: 'Soured milk does not taste good'. Nevertheless, these statements are still implicitly relative, because they still depend for their truth on reference to people's preferences. In such a case, the statement is a generalisation about people's actual preferences rather than a statement of one person's preference.

do not really disagree: there is not one proposition here that either Julie or John asserts, and that the other denies. They are not really disagreeing about the truth-value of the same proposition. That is, they do not dispute the facts of the matter; their claims are simply expressions of different preferences. To continue to dispute such an issue would be a waste of time. Indeed, notice that (1) and (2) are no longer implicitly speaker-relative; they are *explicitly* speaker-relative. So there should be no temptation to say that the truth of either (1) or (2) depends on who is making the claim. If you say 'Chocolate ice cream tastes better than vanilla', then you are implicitly talking about yourself, and the truth of what you say depends on facts about you (your preferences). If you assert (1), however, the truth of your assertion depends only on facts about Julie, not on facts about you.

It should not be supposed that implicit speaker-relativity always has to do with preferences and the like. Suppose, for example, that John has very fair skin, and says 'The sun is too strong' during a walk on a Mediterranean beach. He refers to the danger of sunburn. Julie has darker skin that is less sensitive to the sun; she says 'No, it isn't'. In such a case, it may well be that John is really saying that the sun is too strong for him, and Julie is saying that it isn't too strong for her. So there is no immediate factual disagreement. It is a case of implicit speaker-relativity, but it does not concern preferences, as in the case of ice cream.

Having discussed the concept of implicit speaker-relativity, we are now in a position to confront the myth that truth is relative.

True for me, true for you

Often, people who have succumbed to the myth that 'the truth is always relative' respond to a disagreement about the facts by saying something like, 'Well, that may be true for you, but it's not true for me.' In doing so, they use a common ploy to avoid proper engagement with the argument. Unless the matter under discussion is one that is actually implicitly speaker-relative, as it is in the ice cream example, this is not a legitimate move to make within an attempt to persuade rationally. It is a refusal to argue any further.[3] A similar refusal to engage in debate occurs when someone responds to others' claims by saying, 'that's just your opinion' as though expressing one's mere opinion is not an attempt to make a true claim about a matter but, rather, tantamount to expressing a preference

---

3   When you think about it, it is very hard to see exactly what this claim amounts to. It seems to say that truth is relative to persons; yet it is odd that the only way that this can be the case is for that statement itself to be true in just the way that the statement itself denies – thus it seems that relativism about truth may be **contradictory**.

for chocolate rather than vanilla ice cream. But when we express our opinion on a matter – the best way of reducing crime rates, say – we are expressing our beliefs about the truth of a matter. It is really a kind of self-deception not to face up to this: that when we express our opinion, we are making a claim to truth. So criticising someone's contribution to a conversation by saying, 'that's just a matter of opinion' is another attempt to hinder rational persuasion or debate, and unjustifiably denies that there is any such thing as disagreement.

The 'true-for-me' phraseology, however, is not just a device for evading arguments. It is a characteristic way of expressing the relativity-myth. Thus we shall try to dispel the myth by considering it in more detail. Consider again the sentence

Chocolate ice cream tastes better than vanilla.

The implicit speaker-relativity of a sentence like this might be described by saying that the sentence is true for Julie, and not true for John. Upon hearing Julie assert this sentence, John might say, 'Well, that may be true for you, but it isn't true for me'. John might, in this case, simply be making the point about implicit speaker-relativity. If so, then that is all right; he is quite right to do so. However, phrases such as 'true for me' are sometimes used in what appear to be factual contexts where implicit speaker-relativity is not in play. For example, suppose that Julie believes in astrology, and says:

Scorpios tend to be luckier than Libras.

John does not believe in astrology, and therefore does not believe that one's fortune depends on the part of the year during which one is born. So he thinks this proposition is false. But John, wishing to avoid a painful disagreement, expresses himself by saying, 'Well, that may be true for you, but it isn't true for me'.

As we have just seen, where implicit speaker-relativity is involved, the use of such phrases as 'true for you' is perfectly legitimate. But in the astrology case, the use of this phrase is at best misleading. Is this sentence implicitly speaker-relative? Is the sentence about a preference, belief, or other attitude? It certainly does not seem to be. It is like the sentence about La Paz: it purports to state a fact about the respective fortunes of Scorpios and Libras. Someone asserting the sentence about La Paz expresses their *belief* concerning La Paz, but they are not talking *about* themselves; they are talking only about La Paz and Bolivia. This is shown by the fact that the truth of what they say depends only on how things are with La Paz and Bolivia; its truth does not depend in any way on the beliefs of the

speaker. Likewise, someone asserting the astrology sentence is not saying anything about their own attitudes towards Scorpios and Libras. So when John says, of the astrology sentence, that it might be true for Julie but not for him, he cannot be saying that the sentence is implicitly speaker-relative, and could therefore be true when Julie says it but false when he says it.

What, then, could John reasonably mean by 'true for you' in this context? As we noted, when someone sincerely asserts a declarative sentence, they express a belief. By 'true for you', then, John could simply mean that according to Julie, the sentence is true. That is to say, he could mean simply that Julie believes the proposition expressed by the sentence: that she believes that Scorpios tend to be luckier than Libras. John, of course, denies this very same proposition. So this is a straightforward case of genuine factual disagreement. In suggesting that the sentence is true for Julie and not for himself, all that John is doing is pointing out the fact that he and Julie do disagree: Julie believes, and John disbelieves, the very same proposition. Unfortunately, by using the phrase 'true for you', he makes it *sound* as if it's a case of implicit speaker-relativity, in which case there is no actual disagreement. Since that is the sort of case where 'true for me' has a legitimate point to it, he makes it sound as if there is no actual disagreement, thus smoothing over his difference with Julie. This is perhaps polite of him, but it is really just an evasion.

With these points in mind concerning 'true-for-me', we can now dispel the myth that all truth is relative. The myth is often expressed by saying that we cannot legitimately speak simply of what is true, but only of what is true for me, or true for you, or more generally true-for-X, where X is some person (or perhaps culture). If this is explained as the claim that all statements are really implicitly speaker-relative, then this seems not to be so, as we have seen: a statement like the one about La Paz just isn't speaker-relative. However, the claim that all truth is truth-for-X might also be understood in accordance with the way that John employed the phrase in his dispute with Julie about astrology. According to this interpretation of 'true-for-X', to say that a proposition is true for X is to say that X believes it. Is it conceivable that all truth is really truth-for-X?

Let us work out the implications of supposing that it is. Consider these two sentences:

(3) Scorpios tend to be luckier than Libras.

(4) It is true that Scorpios tend to be luckier than Libras.

According to what we said about the word 'true' at the beginning of this chapter, (3) and (4) are necessarily equivalent: it is impossible for one of

them to be true and the other false. That is why we can always register our agreement with a claim simply by saying, 'that's true'. And that, we might say, is the point of having the word 'true': that is how the word is used. However, according to the version of the myth we are considering, if Julie asserts (4), then she would be speaking more accurately if she were to say 'It is true for me that Scorpios tend to be luckier than Libras'. Thus, as we ourselves would express it, what she says by means of (4), according to the myth, is really:

(5) It is true for Julie that Scorpios tend to be luckier than Libras.

And this, as we said, is more accurately expressed by

(6) Julie believes that Scorpios tend to be luckier than Libras.

So according to this version of the relativity-myth, Julie's utterance of (4) is equivalent to (6). But (6) is certainly not equivalent to (3). (3) makes no reference to Julie; it would be possible for (3) to be true but (6) false, or the other way round (in fact, if (6) were true, then since presumably (3) is false, they would differ in truth-value). Thus, according to this version of the myth, we would have to say that (4) is not equivalent to (3). But that cannot be right, for that is simply not how the word 'true' is used. This version of the relativity-myth violates the actual, ordinary day-to-day meaning of the word 'true' according to which (3) and (4) are equivalent. According to that meaning, Julie is right when she says (3) if, and only if, she is right when she says (4). So (3) and (4), whoever utters them, cannot differ in truth-value.

There is no getting around it, then. There is no way to make satisfactory sense of the relativity-myth. So truth is not relative. It is **objective**, and the truth of a proposition is **independent** of our desiring or believing it to be true. Just as thinking or desiring cannot make the moon be made of green cheese, thinking or desiring cannot make it true that the moon is made of green cheese. To believe is to believe something to be true, but truth is not the same thing as belief. This means that truth is independent of all of us; it does not mean that one powerful person or being could hold the key to all that is true about the world. Thus in saying that truth is objective, we are certainly not taking any kind of political stance, saying that certain cultures or institutions have or might have a monopoly on truth. The aim of good reasoning and argument is to get at the truth, at the way the world is, irrespective of how people think or feel it to be. Rationality is a great leveller. In the pursuit of truth, we are all equally placed before the world, and no amount of political power can provide a real advantage.

Truth, value and morality

So far we have been tackling the myth that all **facts** are relative. Many people are also tempted to think that **values** like those central to moral issues are relative to personal or cultural preferences. This view is much harder for the critical thinker to dispel. By way of example, we'll take the claim

> Abortion is immoral.

According to the relativist view, when anti-abortionists say abortion is morally unacceptable and their opponents contradict them and say that it is morally acceptable, there is no real disagreement; rather the two sides do not share the same moral preferences. Thus, for the relativist, value-statements are always speaker-relative, whether implicitly or explicitly. An apparent disagreement over a value is in this respect like that between Julie and John concerning chocolate versus vanilla ice cream. One reason why this relativistic view of moral issues is so tempting is that we feel uncomfortable about being seen to dictate morals to other people because we (rightly) value tolerance of different opinions.[4]

We cannot prove, and would not try to prove, that moral relativism is false. It is conceivable that all claims about values are implicitly speaker-relative, or perhaps implicitly relative in some other way. We will, however, try to explain why there is good reason to resist moral relativism. Whilst moral issues are almost always complicated and can be difficult to agree upon, the problem with this simplistic relativist approach is that it leaves no room for genuine disagreement about moral issues. It enjoins us to say, 'You feel OK about abortion, I don't, end of story' without attempting rationally to persuade one another of the truth of our beliefs. The poverty of such a view of morals is illustrated by the following case: suppose there emerged a terrible fascist regime that murdered millions of people on grounds of race, religion or political beliefs. Simplistic relativists who desire to remain consistent with their relativist commitment would be unable to hope that the fascists are wrong, and that, therefore, others could rationally be persuaded that the fascists are wrong; for the fascists' views about the moral status of their victims would be nothing more than preferences that just happen to be different from most other people's. In the same way that most people prefer ice cream to carrots, this regime prefers murdering people who are different

---

4   This repudiation of moral relativism is aimed only at the very naive type of relativism that holds that moral questions are simply a matter of personal or cultural preference and therefore not the subject of genuine disagreements. There are various more sophisticated versions of moral relativism. It is not our intention to repudiate such theories.

from them to living tolerantly alongside them. This is a very extreme example, but the point is well made. To adopt naive relativism about moral matters, to deny there can be a truth of the matter; and to say that claims such as the claim that torture is wrong have a similar status to the claim that chocolate ice cream tastes best, is to deny ourselves the opportunity to attempt rationally to persuade others that their moral beliefs are false and to persuade them not to follow courses of action that would be harmful to others. It places fundamental moral issues outside of the ambit of critical thinking.

Yet the critical thinker does not have to give up hope. While people may continue to believe that moral relativism is true, they must also be consistent in their judgements. For example: it is irrational to hold that murderers are bad people, and at the same time, in full knowledge of his crimes, hold that Jack the Ripper was not a bad person. Thus, the critical thinker can at least demand logical consistency of the relativist: they can show that there is a valid argument from a premise accepted by the relativist to the conclusion that Jack the Ripper is bad. This means that it is possible to refute the relativist's moral views if it can be established that they are incoherent. In moral arguments a good way to do this is to find a general principle that the relativist accepts and go on to show how it is inconsistent with the belief that you wish to challenge. For example, if a relativist is pro-abortion and you discover that they are also against any form of killing, then you could force them to revise their beliefs if you could establish by argument that abortion was a type of killing.[5]

Belief, justification and truth

It is not the case that for any given proposition, we either believe it to be true (believe it), or believe it to be false (disbelieve it). We may simply have no attitude towards it, either because it has never come to our attention, or because we choose not to consider the question. Or we may have an attitude that is midway between belief and disbelief: we may suspend judgement, because we find upon reflection that we lack sufficient evidence to make the judgement. This happens when we lack an argument either for or against that proposition that is rationally persuasive for us. The issue of the existence of a single deity who is said to have created the world provides a good example with which to illustrate the point. Consider the claim,

---

5   See Louis Pojman's essay 'Ethical Relativism versus Ethical Objectivism' in his, *Introduction to Philosophy: Classical and Contemporary Readings* (Belmont: Wadsworth, 1991) for an excellent critical analysis of moral relativism and a convincing case for a version of moral objectivism. This very accessible paper is aimed at beginning Philosophy students.

The world was created by an omnipotent and omniscient being.

Some people, let's call them 'theists', believe the proposition expressed here. Others, atheists, disbelieve it. Others, agnostics, are not sure whether or not to believe it so they have suspended judgement until such time as they acquire sufficient evidence (or faith) to support a belief either way. These three positions probably cover most adults in our society. But consider the position of most pre-school children. They do not believe that a deity made the world. They do not believe that a deity didn't make the world. However, it would not be accurate to say that they have suspended judgement on the issue. For they have not considered the issue, indeed they are probably too young to understand the claim that is the subject of debate. Of course, when assessing an argument's soundness and rational persuasiveness, the stance of having no opinion about the truth of the premises is not an option for critical thinkers. Either we consider ourselves to have sufficient evidence to believe or disbelieve the proposition expressed by the premise(s) or we recognise that we lack such evidence, and suspend judgement until such evidence is available. And even if we believe the argument to be inductively sound, we may find that having accepted its premises, the argument is not rationally persuasive for us because its conclusion is defeated by some further evidence that we have.

It is crucial to bear in mind that in saying that someone does not hold a certain belief, we are not saying or implying that they hold the opposite belief. If someone tells you that they don't believe that the Prime Minister is a bad person, they are not thereby saying or implying that they believe that the Prime Minister is not a bad person, let alone that the Prime Minister is a good person. They may not know the Prime Minister very well and may want to learn more about the Prime Minister before making a judgement, or they may simply not care and have no intention of ever forming a belief about the Prime Minister's moral standing. Or perhaps they have never even heard of the Prime Minister. As we saw in Chapter 4, an argument that assumes that someone who does not believe a proposition believes its negation commits a version of the **epistemic fallacy**.

Of the four stances we can take towards a proposition – **believing it**, **not believing it**, **suspending judgement**, **not engaging with it** – two, believing and not believing, admit of degrees. Smith and Jones may both believe that the Conservative Party will form the government after the next general election, but they may not hold the belief with the same degree of confidence. Smith may be nearly certain, whilst Jones, an arch-sceptic when it comes to predicting voting behaviour, still believes it but to a lesser degree. Similarly Jack and Jill may both disbelieve the proposition that the Conservative Party will form the next government, but

the strength of Jack's disbelief may be such that he'll bet two weeks' wages at odds of 20 to 1 on a Conservative defeat. Whereas Jill, normally a keen betting woman, would not risk nearly as much money (if offered the same odds).

Given that someone holds a certain belief, we can ask whether or not the belief is **justified**, and whether or not it is **true**. Suppose Smith works for a national polling organisation and has seen the results of several methodologically sound top-secret polls that predict the Conservative Party will gain 80 per cent of the vote on election day. In the absence of any evidence that goes against that prediction, Smith is in possession of an argument that is highly rationally persuasive for him. He would thus be well-justified (other factors such as the reliability of the data notwithstanding) in believing the conclusion of the argument with a strong degree of certainty. As critical thinkers we may conclude that Smith's belief is reasonable or **justified** given what we know about the evidence available to him. Whenever someone's belief is backed by an argument that is rationally persuasive for them, it follows that they are justified in believing the conclusion, precisely to the extent that the argument is rationally persuasive for them.[6]

However, the fact that Smith seems perfectly rational and justified in holding this belief does not establish that the belief is true. That depends upon whether the Conservative Party do in fact form the next government. A belief's being true is a matter of its fitting the facts, not of there being good reasons to think that such-and-such is the case. We see this clearly when we consider the possibility of being justified in holding a belief that is actually false. Suppose I form the belief that the government has just been deposed in a military coup. I am led to form this belief because, while writing this chapter, I have been distracted into reading electronic newspapers and I find the same story appearing on several previously reliable and reputable internet sites. I track down one or two colleagues, they look at their own favourite sites and, sure enough, find the same story. Given the widespread coverage of the coup and the fact that the sources are reliable and reputable, I would be justified in believing that the government had been deposed by a military coup. But in fact, despite my good reasons for having formed this belief, it turns out to be false. A clever computer hacker has managed to hack into several UK-based internet news services and post this false story. So I have a

---

6   This is not to say that a belief is justified only if one has a rationally persuasive argument for it. According to many theorists, some beliefs, especially those that arise in perception, are justified but not 'inferentially' – not, that is, by means of arguments or reasoning. Others deny this, holding that even the justification of perceptual beliefs depends upon underlying premises or assumptions. Since this is a fundamental open question in the theory of knowledge (epistemology), we have avoided taking a stand on it here. For further discussion see Robert Audi, *Epistemology A Contemporary Introduction to the Theory of Knowledge* (London: Routledge, 1997).

justified but false belief. I may be in possession of an argument that is rationally persuasive for me, but my belief in its conclusion still need not be true.

Knowledge

Truth and knowledge are intimately linked. Our desire to accumulate knowledge stems from a desire to get at the truth. And certainly if one knows something, then the proposition known must be true: one cannot rightly be said to know that the cat is on the mat if the cat is not on the mat. The truth of a belief is certainly one necessary condition of that belief's being knowledge. So it may seem reasonable to suppose that if Jones believes that the government has been deposed in a military coup, and the government has been deposed in a military coup, then Jones knows that the government has been deposed in a military coup. This is to say that knowledge that such-and-such is the case is simply **having a true belief** that such-and-such is the case.

However, this is not correct. Having a belief that is true is necessary, but not sufficient for knowledge. Suppose that Smith has been given LSD. He is hallucinating like mad, and his reason has gone completely haywire. The conviction suddenly impresses itself upon his unhinged mind that there are leprechauns in the next room; and that because of this, his mother's house in Sydney, 15,000 miles away, is on fire. In fact, Smith takes LSD every day, and believes, every time, that his mother's house is on fire (though for a different reason each time). But this time, it just so happens that his mother's house is on fire. Just by chance, the belief, this time, was true. No one would say that Smith, in this situation, knows that his mother's house is on fire. It was true belief, but not knowledge. He is like a person who shoots at a target blindfolded, dozens of times, and finally hits it, despite not having had any idea where the target was.

So knowledge cannot be identified simply with having a true belief. A further ingredient must be added to true belief, if it is to be knowledge. In particular, a true belief counts as knowledge only if we arrive at that true belief via the **right route**: we have knowledge only if we have **good reasons** for a holding belief that turns out to be true. We have to be **justified**, have solid **evidential support**. We must, if you like, earn the right to be sure. Lucky guesses do not count as knowledge.

This third requirement is incorporated into the traditional philosophical account of knowledge, which is called the **tripartite account**:

Truth, knowledge and belief

> **For any subject (S) and any proposition (P),**
> **S *knows* that P if and only if:**
> **(i)   S believes that P**
> **(ii) P is true**
> **(iii) S is justified in believing that P.**

According to this account, to know that P is to have a justified true belief that P. It amounts to saying, then, that **knowledge is identical with justified true belief**. If I have a justified true belief that Tony Blair is UK Prime Minister then I *know* that Tony Blair is UK Prime Minister.

Justification failure

There are two ways in which belief can lack the justification, the evidential support, required for knowledge: it can either be genuine but insufficient, or simply not genuine, that is, mistaken.

Insufficiency

Consider Mrs Green, the greengrocer, who discovers that a customer has left their shopping behind. She has served five customers so far that day – Mr Red, Mrs Pink, Mr Orange, Mr Yellow and Mrs Blue. She has no reason to suppose that any one of them is any more forgetful than any of the others. None of them is in the habit of forgetting their shopping, for instance. Mrs Green concludes that one of the men left the shopping behind. She does have some justification for forming this belief: of the five customers three were male, so on the balance of probability the forgetful shopper is more likely to be a man. That is to say, she does have an argument for that conclusion which is somewhat **rationally persuasive** for her. But, intuitively, we may feel that this probability is somewhat slim justification for her belief. Even if it turns out that the shopping belongs to Mr Orange, we don't feel inclined to say that she arrived at her true belief via the right route; that is, that her true belief is really **justified**. So even if she's right, we don't feel that she *knows* that the forgetful shopper is a man. This is not to say that probability as a justification for true belief should be ruled out per se. The problem here is that the probability is not strong enough. If there had been four male shoppers and one woman, we would be more inclined to allow that Mrs Green's belief is justified. If there had been ninety-nine men and one woman, we would certainly be so inclined.

There are occasions when there is far more at stake in having the proper justification for forming a true belief than simply making sure that someone gets their groceries. An obvious case is where principles of justice are involved, where evidence is required to demonstrate beyond reasonable doubt that a defendant is guilty of an alleged crime. Suppose Jones is accused of the murder of Brown. In fact, Jones is guilty: he strangled Brown to death. Smith is called as a prosecution witness to Jones' crime. He is called as a witness because he has given a police statement in which he claims to have seen Jones emerging from Brown's house carrying a knife. If this is all the evidence the jury has, is it, by itself, sufficient reason for jury members to form the true belief that Jones is guilty of Brown's murder? Clearly not, even though Jones is in fact guilty of the murder. Brown died of asphyxia, not from stab wounds, and probably the evidence would have been too weak even if Jones had stabbed Brown rather than strangling Brown. So although the jury members may have reached their true belief about Jones' guilt on the basis of Smith's evidence, they do not have a **justified** belief about Jones' guilt. If this is the only evidence they have for reaching their verdict, then they should not convict Jones of Brown's murder. Note that this holds even though, as it happens, Jones is guilty.

How strong must our evidential support be, if a belief is to be justified? How strong must it be if a true belief is to qualify as knowledge? Unfortunately, there is no precise answer to this question. The terms 'justified' and 'knowledge', like 'bald', are somewhat vague: just as there is no exact quantity or proportion of cranial hairlessness required for baldness, there is no exact strength of evidence for a belief short of which it is not knowledge, and beyond which it is knowledge. But this should not trouble us too much. In practice, we are adept at recognising the cases where our degree of justification is borderline. This does not change the nature of our task as critical thinkers: we know that our task is to fashion sound arguments that are rationally persuasive for ourselves or for their intended audiences. By paying attention to the reasons why we are entitled to form true beliefs, we are able to achieve the firmest possible grasp upon what we can legitimately claim to know.

Mistakes about justification

Sometimes mis-evaluation of evidence leads us to form **irrational beliefs**, even in cases that are not so extreme as Smith and his hallucinations. In such cases we over-estimate our evidential support: we think that the evidence supports a belief when it doesn't. This is an error that people who attempt to persuade by devious means can often exploit to their

advantage. For instance they might describe one or two vivid examples of some alleged phenomenon with the aim of getting us to form generalisations for which the examples do not provide adequate support. An example that should be familiar to readers and observers of the UK tabloid press is that of a vivid description of a recipient of state benefit or of an asylum seeker, which apparently supports the claim that they enjoy an undeservedly luxurious lifestyle. Such examples are rolled out in an attempt to persuade readers to believe that the majority of benefit recipients and asylum seekers are undeserving cheats. Often, the details of such examples and the manner of their presentation are so rhetorically powerful that they cause us to mis-evaluate the significance of the evidence presented to us. In particular, we make an unjustified **inductive inference** (one with very low force). When we are led to make such mistakes and form irrational beliefs, we allow ourselves to be distracted by factors other than principles of good reasoning.

There is another type of irrational belief. Sometimes we allow ourselves to accept false beliefs because we believe doing so will benefit us in some way. In such cases we not only lack evidence for the belief, but we also seem not to care that we have none. For instance, someone might believe that the predictions of their astrologer will come true; they might believe in faith-healing; ESP; life after death; that everything in the Bible is literally true; that their terminally-ill grandparent will make a recovery; that a certain alternative healing method will cure cancer. We often feel that it would be unfeeling to point out that such a person's beliefs are irrational. Indeed, we often admire their faith and sincerity. However, it would be a mistake to retreat into relativity and conclude that their belief is 'true for them'. The belief is false, that's that. Indeed, such people can sometimes paradoxically be described as believing that which they know to be false – self-deceived.

However, we should not be too cold-blooded. It's not usually the case that such people are intentionally deceiving themselves, but rather that they are trying to deal with their plight by having faith that their hopes (however unlikely) will be borne out.[7] So it's rather unfair and unfeeling to accuse such people of being irrational, even though strictly speaking they are. Whilst we should try to avoid irrationality, we must also accept that as human beings, it is sometimes psychologically better for us if our beliefs and behaviour fall short of rationality.

---

7   There are, of course, cases where people are guilty of self-deception. A classic case is the parent who knows (for they witnessed the act) that their child stole sweets from the supermarket, but refuses to believe that they did.

Knowledge and rational persuasiveness

Rational persuasiveness requires that you be justified in accepting the premises of an argument. If the argument is **deductively valid**, then if you are justified in accepting the premises, you are equally justified in accepting the conclusion. If the premises of such an argument are actually true, then since you know the premises to be true, you know the conclusion to be true.[8]

However, there are two complications. First, a rationally persuasive argument can be unsound: rational persuasiveness does not itself require that the premises of the argument actually be true. It requires that one be justified in accepting the premises; but, we have seen, it is perfectly possible to have good reasons, to be justified, in holding a false belief. Thus one can be rationally persuaded by an argument but not know the argument's premises to be true. In such a case, one can be rationally persuaded by an argument without thereby acquiring knowledge. **Knowledge** that an argument is sound, on the other hand, makes a stronger demand on our epistemological relationship with its premises. To know that an argument is sound, we have to know that its premises are true. To know that an argument is sound, we need justified true beliefs in the premises. Whereas for rational persuasiveness, we need only the justification; the belief need not be true.

Second, in the case of an inductively forceful argument that is not defeated, the degree of justification transmitted from known premises to the argument's conclusion may be insufficient to establish the conclusion as knowledge. For the degree of inductive force may be too weak. One might, for example, know with certainty that fifty-one of the hundred stones in a bag are black; one would thus have an inductively sound argument that is rationally persuasive for the conclusion that the next stone drawn at random from the bag will be black. But the argument is only slightly rationally persuasive. Even if that conclusion turns out to be true, the belief in it would be insufficiently justified by the argument, and so would not count as knowledge. More generally, the degree of justification one has for the conclusion of an argument that is inductively forceful but not deductively sound will be smaller than the degree one has for the premises (unless one has some other source of justification for the conclusion).

Let us now look at a case that displays what knowledge is, as opposed to what it isn't. The following scenario provides an example of the requirements placed on knowledge claims:

8   For simplicity we are assuming that if you know one proposition p, and the argument from p to another proposition q is deductively valid, then you know that q. In epistemology this is called the 'Closure Principle'. In more advanced philosophical treatments of these issues, this principle is called into question.

Truth, knowledge and belief

> Tabitha Tabloid, a journalist, is at a restaurant one evening and sees a prominent politician at an adjacent table with a person she knows not to be their spouse. 'Perhaps there's a story in this', she thinks. Later, as she walks to her car, she notices a couple locked in a passionate embrace in a parked car that just happens to be parked under a street lamp. On discreet but close inspection she sees that the couple is the politician and their companion. She forms the belief (which happens to be true) that the politician is having an illicit affair with this person. In fact, she believes that the evidence for her belief is sufficiently strong that she persuades her editor that her story should be the next morning's front-page lead.

Here we have the three components that are apparently required for legitimate knowledge claims: a true belief that the politician is having an affair and a rational justification for that belief based upon what Tabitha witnessed in the restaurant and subsequently in the parked car.

Gettier cases

We have said that knowledge is justified true belief. However, as you're no doubt aware by now, philosophers like to get to the bottom of things, and this can often mean throwing the spanner in the works of an apparently satisfactory theory. We are going to end this chapter by discussing some cases that cast some doubt on the plausibility of the intuitive theory that knowledge is justified true belief.

These sorts of cases, in which we have a true belief for which we have adequate rational justification yet can't really say that we have knowledge, could occur in a court case:

> Imagine that a jury in a murder case, having carefully considered the evidence, form the belief that the defendant is guilty. The strongest piece of evidence they have for their belief is the fact that the defendant's blood was found at the scene of the crime and on the gloves also found there. As it turns out, the defendant did indeed commit the murders of which they stood accused.

Thus, the jury have a belief, it's true and they have rational justification (in the form of the blood, gloves and other evidence) for it. However, suppose that the blood at the scene was planted by the investigating police officers in an (albeit unnecessary) attempt to incriminate the defendant.

Now, if the jury convict on the basis of the forensic evidence concerning the blood found at the crime scene, would you be comfortable with the conviction – would you feel that the jury are really justified in convicting?

Cases such as these should encourage us to reflect further upon the adequacy of the tripartite account of knowledge. Indeed, some philosophers are fascinated by such cases. In a famous essay Edmund Gettier devised a set of cases that lead us to question whether rational justification of a true belief is sufficient for knowledge.[9] The cases are quite complex and rely upon the rules governing logical relations between certain sorts of propositions. The following **Gettier case** relies upon the rules of logical entailment for propositions that are joined by the connective word 'and'. It is fairly complex, but try to follow and see why the satisfaction of the three conditions for knowledge given by the tripartite account doesn't seem to be sufficient to say that there is knowledge in this case.

Suppose that Sherlock the detective, investigating a murder case, has the following belief:

A   Billy the Kid committed the murder and Billy the Kid has size 10 feet.

Sherlock has uncovered various evidence that makes his belief justified: forensic evidence, eyewitness accounts, taking Billy into custody and measuring his feet then comparing it with footprints found at the scene of the crime, and so on. Now, given the evidence he has for A, particularly that pertaining to the size of the footprints left at the scene, Sherlock inferentially forms a second belief:

B   The murderer has size 10 feet.

He can do this because A logically entails B: the inference from A to B is deductively valid. However, belief A is false because despite all the evidence to the contrary, Billy did not commit the crime in question. Perhaps he was actually at the scene around the time and his fingerprints were found there, because he was in the process of a burglary. The actual murderer, however, is Wild Willie Williams, who does indeed have size 10 feet. So while belief A may be false, belief B is true, the murderer does have size 10 feet. Furthermore, Sherlock's belief B fulfils the conditions for knowledge given by the tripartite account. Sherlock believes that the murderer has size 10 feet. His belief is true because Wild Willie Williams does indeed have size 10 feet. His belief is justified on the basis

---

9   'Is Knowledge Justified True Belief?', in *Analysis*, 1963.

of the evidence he has for his belief A from which he has validly inferred belief B. Thus Sherlock's belief that the murderer has size 10 feet is a justified true belief. If the tripartite account is a correct account of knowledge, we ought not to have any qualms about saying that Sherlock *knows* that the murderer has size 10 feet. Yet one feels uncomfortable concluding that Sherlock knows the size of the murderer's feet. One suspects that this has something to do with the fact that all the evidence provided for belief A, from which B was inferred, actually comes from Billy and not from the actual murderer.

A further, similar case confirms this hunch: Glitzie, Glossie, Barbie and Sindy are all actors auditioning for the lead role in a new television series, *Bimbowatch*. Understandably, they are all keen to get the part and are checking out their rivals. Glossie has evidence for the following belief:

> **C** Glitzie has paid a lot of money to surgically enhance her body to Hollywood standards and she will get the part.

Glossie's evidence for belief C is that first, the casting director has told her that Glitzie is favoured for the part; and second, Glossie uses the same plastic surgeon and has seen Glitzie at the clinic. This evidence justifies her belief C. Now, as you can no doubt see, C entails the following proposition:

> **D** The woman who will get the part has a surgically enhanced body.

Suppose that Glossie sees that if she's justified in believing C, she is also justified in believing D and thereby forms the belief D. However, unknown to Glossie, Barbie will be given the part. Also unknown to Glossie is the fact that Barbie has also had her body surgically enhanced to meet Hollywood's exacting standards. Thus D is true, but C from which, following rules of entailment, she inferred D, is false. Once again, all the conditions for knowledge according to the tripartite account are met:

> **(i)** Glossie believes that the woman who will get the part has a surgically enhanced body (D)
> **(ii)** D is true
> **(iii)** Glossie is justified in believing D (because she forms belief D on the basis of C for which she has good evidence).

However, the natural thought is that despite having met these three conditions, Glossie does not *know* that the woman who will get the part has a surgically enhanced body. This case reinforces the idea that this is

because, although she has acquired justification for belief D via her justification for belief C, her belief in D is based upon tracking the wrong facts of the matter. For D is true in virtue of the fact that Barbie has a surgically enhanced body, not in virtue of the fact that Glitzie has a surgically enhanced body which is the fact for which Glossie has evidence. Her belief D is formed on the basis of what she knows about Glitzie's body, not what she believes or knows about Barbie.

To re-cap then: C is a justified but false belief. It is justified because of the evidence of the casting director's admission and Glossie's having seen Glitzie at the plastic surgeon's clinic. It is false because Glitzie does not actually get the part in *Bimbowatch*. On the basis of this justified but false belief C, Glossie forms belief D, which is true, and justified on the basis of Glossie's evidence for C. But although D is both true and justified, it doesn't seem to count as knowledge because it doesn't track the correct fact in virtue of which it is true.

Cases such as the two considered here suggest that it remains possible for someone not to count as having knowledge that P despite the fact that their belief that P meets the conditions laid down by the tripartite account of knowledge. The implication of such cases is that knowledge is not simply justified, true belief.

That is not to say that we shouldn't seek rational justification for our beliefs. In most cases the evidence available to us will provide good reasons for holding a belief, and when that belief is true then we can reasonably claim to have knowledge. Our evidence, our justification, does tell us how likely it is that our belief is true. However, cases such as these demonstrate that rational justification can fail to generate knowledge even where it does secure its intended object, namely truth.

CHAPTER SUMMARY

It is crucial for critical thinkers to recognise that **truth is objective** and **not relative**. Otherwise the critical thinker's objective of analysing and assessing arguments with the aim of getting at the truth of a matter is deeply undermined. Some sentences that appear to be straightforward assertions are, in fact, implicitly relative expressions of subjective preferences and tastes: they are **implicitly speaker-relative**. In such cases there cannot be genuine disagreements about the facts of the matter, whereas when a factual proposition is asserted, a genuine disagreement can occur because a truth is at stake. To deny this, to accept the myth that all truth is relative, is to accept something which it seems impossible to make sense of. In the case of **moral beliefs**, or beliefs about **value**, relativism may not be readily refutable,

but (1) the consequences of denying that there is truth in this realm appear to be extremely pernicious and (2) relativism does not completely close the door to rational persuasion, because we may still demand consistency of the relativist.

If our stance is not that of not considering it, or refusing to consider it, which of the three remaining **stances** we take towards a given proposition depends upon the evidence available to us. We require sufficient evidence for the truth or falsity of a proposition in order to justify believing or rejecting it. If such evidence is unavailable, we should suspend our judgement. It is perfectly possible to be justified in holding a false belief. The evidence available to us might make it rational for us to accept a proposition despite the fact that it turns out to be false.

Merely having a true belief in a proposition is not sufficient to count as **knowing** the proposition, though it is necessary for knowledge thereof. Traditionally, philosophers have concluded that knowledge requires a true belief arrived at via the **right route**; that is, via good reasoning based upon sufficient evidential support. This leads to the formulation of an account of knowledge known as the **tripartite account** according to which **to know that P, is to have a justified true belief that P**.

However, a series of cases proposed by Edmund Gettier prompt us to question the completeness of this account. There are cases in which a justified true belief is not knowledge. Many of these cases involve someone having a justified true belief which they have arrived at via an entailment from another belief that is justified yet false. The evidence that apparently supports their true belief is evidence for the initial, false, belief but not really evidence for their true belief. This is because the evidence does not ensue from the facts that make the true belief true.

EXERCISES **1** For each of the propositions expressed by the following sentences, say whether its truth-value is implicitly speaker-relative. If it is, explain also whether it could be implicitly relative to the preferences of the person asserting it. Explain your answers. Note that in some cases, there is room for reasonable disagreement as to the correct answer. So don't worry if you are not sure of how to answer; the important thing is to try to explain the answer you do give.

a   The maximum speed limit on UK roads is 70 miles/hour.

b   Yorkshire is nearby.

c   Martina Hingis did not win the 2000 US Open Tennis Championship.

d   It's raining.

e   Euthanasia is morally acceptable.

f   The climate in Southern Italy is too warm.

g   Haggis is delicious.

h   Bill Clinton was a great president.

i   I am in pain.

j   Drink-driving should be punished by a prison sentence.

k   The government should cut the tax on petrol.

l   Bitter tastes better than lager.

m   I like bitter better than lager.

n   Robbie Williams has sold more records than Ricky Martin.

o   Robbie Williams' music is better than Ricky Martin's.

p   If you eat more fruit and vegetables you will be healthier.

q   It's going to snow tomorrow.

r   Tchaikovsky wrote the 1812 overture.

s   I am in pain.

t   That hurts!

u   France and Great Britain were Second World War allies.

v   10 is the square root of 100.

w   All pop/rock music is pathetic cacophony and a waste of any intelligent person's time.

x   All triangles have three angles totalling 180 degrees.

y   Blood is thicker than water.

---

2   Without looking back at the relevant section of the book, write out the tripartite account of knowledge. Then, using your own example,

explain why simply having a true belief that P, is insufficient for knowledge that P.

---

3   Suppose you are justified in believing a proposition that is in fact false. Could there be an argument for that proposition which is rationally persuasive for you? Would the situation change if the proposition were actually true rather than false? Explain your answer, using examples – your own examples – if you find it useful.

---

4   Consider the following case and then answer the questions that follow.

> Julia, Nicole and Jennifer have all been nominated for the Best Actress Award at this year's Academy Awards. Harry Starstruck is Hollywood correspondent for the *Daily Dirt*. Harry believes that Julia will win the Oscar. His belief is based upon evidence provided by a tip-off from a usually reliable inside source. Harry's source also provides an extra titbit of gossip: Julia has recently dyed her hair flame red. Thus, Harry also comes to believe that the woman who will win the best actress Oscar has red hair. In the event, however, Harry's usually reliable source turns out to have let him down. The Oscar was won by another famous Hollywood red-head – Nicole.

a   According to the tripartite definition of knowledge, does Harry know that the woman who will win the Oscar for best actress has red hair? Explain your answer.

b   Examine your own intuitions about knowledge, does Harry know that the woman who will win the Oscar has red hair? Why/Why not?

■ Glossary

Words written in **bold** indicate references to relevant further glossary entries.

Ambiguity   A sentence is ambiguous in a given context if there is more than one possible way to interpret it in that context. A word is ambiguous in a given context if there is more than one possible way to interpret it in that context. See also **lexical ambiguity**, **syntactic ambiguity**. Compare **vagueness**.

Antecedent   A **conditional statement** asserts a relation between two propositions, the antecedent and the **consequent**. When the antecedent of a true conditional statement is true, the consequent must also be true. Thus the antecedent of 'If John wins then Mary will cry' is 'John wins'; if one accepts both this conditional and its antecedent as true, then one must also accept that Mary will cry.

Argument   An argument is a system of propositions comprising one or more **premises** advanced by the arguer in support of a **conclusion**. Arguments may be evaluated as **valid**, **inductively forceful**, or neither, but not as **true** or false.

Argument commentary   An argument commentary is a short essay that discusses an **argument reconstruction**, covering the following points: (i) how and why the argument was reconstructed as it was; (ii) the **validity** or degree of **inductive force** of the argument; (iii) the **truth-values** of the premises; (vi) the degree of **rational persuasiveness** of the argument for relevant audiences, where this is not already addressed by (iii).

Glossary

Argument reconstruction   A presentation of an argument in **standard form** in which premises may be added, made explicit or clarified in order to make the argument either **valid** or **inductively forceful**.

Argument trees   A graphical representation of the structure of an argument. See pp. 227–229.

Causal generalisation   A causal generalisation is a **generalisation** to the effect that things of one kind tend to cause things of another kind. Such a generalisation is true if the presence of a thing of the first kind raises the probability of things of the second kind, even when other possible causes are absent.

Conclusion   An argument's conclusion is the proposition that its **premises** are intended to support. The distinctive aim of giving an argument is rationally to persuade an audience that the conclusion is true.

Conclusion indicators   These are words such as 'therefore', 'so', and 'thus' that are often used to indicate the conclusion of an argument. Such words sometimes serve other purposes, however, such as indicating a causal relationship.

Conditional   A conditional proposition is a single proposition that joins two propositions, the **antecedent** and **consequent**. Its usual function is to assert that if the antecedent is true, then so is the consequent. It is most characteristically expressed by means of 'if–then'. Conditionals, unlike **arguments**, may be evaluated as **true** or false.

Conditional probability   The conditional probability of a proposition P, given evidence (set of premises) A, is the **probability**, if A holds (that is, if all the premises in A are true), that P is true. Conditional probabilities are to be assessed ignoring any evidence not included in A that might be relevant to the truth-value of P.

Connecting premises   These are **conditionals** or **generalisations**, usually **implicitly** assumed by an arguer, that are needed in order to infer the argument's conclusion. For example if the argument is 'Mary is a doctor, therefore Mary has a university degree', a suitable connecting premise is 'All doctors have university degrees'.

Connotation, primary/secondary   The primary connotation of a term is the condition that is necessary and sufficient for something's being

a member of the **extension** of the term: it is the rule that determines whether or not a given thing is or is not correctly designated by the term. The secondary connotation of a term is the range of further attributes that a thing is commonly assumed to possess if it is thought to be correctly designated by the term. The primary connotation of 'mink coat', for example, is 'coat made from the furs of minks'; its secondary connotation might include 'expensive', 'posh', 'old-fashioned', 'warm', and 'beautiful' or 'immoral', depending on opinion. Unlike a term's primary connotation, its secondary connotation can vary from person to person.

Consequent   See **antecedent** and **conditional**.

Context   The context of an argument is the set of circumstances in which an argument is actually advanced by an arguer. Context is significant because in order to **reconstruct** an argument we often have to fill in **premises** that are only **implicitly** assumed by the arguer. To determine what an arguer is likely to have assumed, we usually need to know the circumstances in which the argument is advanced.

Cost/benefit   Potential actions are evaluated as justified or not in terms of their costs and benefits, where these indicate values, positive or negative, that need not be measurable monetarily (that is, they may be moral or aesthetic values, or some other for which the question of price may be inapplicable). Such arguments are the domain of **practical reasoning**.

Counter-example   (i) A counter-example to a **generalisation** is a particular statement – a statement that is not a generalisation – that is the negation of an **instance** of the generalisation. For example 'Darcy Bussell is a great ballerina who is tall' is a counter-example to 'No great ballerinas are tall'. (ii) A counter-example to an argument is an argument of the same form or pattern as the first argument that is clearly invalid or inductively non-forceful. It is used as an illustration to make it clear that the first argument is invalid or inductively non-forceful.

Covering generalisation   Often, where a premise of an argument is a conditional proposition of the form 'If Mary is a doctor, then Mary has a university degree', it is **implicitly** inferred from a **generalisation** of which it is an **instance** – in this case, 'All doctors have university degrees'. This relation is made more conspicuous if we express the generalisation in the form 'For any given person, if that person is a doctor, then that person has a university degree'.

Glossary

Credibility   The degree to which someone's having said something constitutes a reason to think it true. While critical reasoning requires us to focus on an argument and not on the person putting it forward, a person's character and actions are certainly relevant to their credibility.

Deductive validity   Validity can be defined according to either of the following, equivalent formulations: (1) An argument is valid if and only if it would be impossible for its premises to be true but its conclusion false. (2) An argument is valid if and only if necessarily, if its premises are true, then its conclusion is true.

Defeated argument   An **inductively forceful** argument, whose premises a person reasonably believes, is defeated for that person if he or she has good reasons to think the conclusion false.

Expected value   The expected value of a given action depends on the values and probabilities of the possible outcomes. In particular, if o_1, o_2, ... and so on are the possible outcomes of an action, $V(o)$ is the value of a given outcome, and $P(o)$ is the probability of a given outcome, then the expected value of the action is

$$[\mathbf{V}(o_1) \times \mathbf{P}(o_1)] + [\mathbf{V}(o_2) \times \mathbf{P}(o_2)] + \ldots$$

The value of each outcome must be assigned a number, but the purpose of the numbers is only to indicate the comparative values of the possible outcomes. For example if one outcome is judged to be twice as good as a second, we could assign them any two numbers so long as the first is assigned a number that is twice that assigned to the second. Expected value is the central concept of **cost/benefit** analysis. The idea is that, given a range of possible actions, one should perform the action with the highest expected value.

Explanations   We give an **argument** for something when we seek to persuade an audience that that proposition is true. By contrast, when we give an explanation of something, we know, or assume, that the audience already accepts that the proposition to be explained is true. Our aim is not to give reasons for believing that proposition, but to specify, for example, the causes of the event that it mentions. Both arguments and explanations can be described as answering 'why?' questions, but there is a crucial difference: whereas the question in the case of explanation is 'why is it so?', or 'why did it happen?', the question in the case of argument is 'why should I believe it?'.

Potentially confusing is that in order to establish that an explanation is the correct one – e.g. that it specifies the actual cause of an event – we often have to give reasons why it should be believed, i.e. an argument. That is, we sometimes have to argue *for* an explanation.

Extended arguments   An extended argument for a proposition is one containing more than one **inference**: a **conclusion** is used as a **premise** for a further argument, the conclusion of which may be used as a premise for a further argument, and so on. Conclusions used as premises for further inferences in an extended argument are called **intermediate conclusions**.

Extension   The extension of a general term such as 'cat' or 'red car' is the set or group of things designated by the term.

Factual assessment   The stage in the assessment of an argument in which we determine whether or not the argument's premises are true. If the argument is either **valid** or **inductively forceful**, then the argument is **sound** if and only if all its premises are true.

Fallacies   The term 'fallacy' encompasses certain commonly encountered failures of argumentation; it is partly because they are often effective as **rhetorical ploys** that they are commonly encountered. *Formal* fallacies are simply logical mistakes; that is, arguments that fail to be valid or inductively forceful in certain characteristic ways. *Substantive* fallacies are arguments that implicitly assume some quite general premise of a kind which, when more closely and explicitly considered, can readily be seen to be false. Some other common defects in argumentation fit neither classification; but since they involve fooling the audience in the context of argument they can be appropriately classified as fallacies.

Generalisations   A generalisation is a proposition concerning a class of things, either explicitly or implicitly involving a **quantifier** such as 'all', 'every', 'no', 'some', 'most', 'twelve', 'at least twelve', and so on. For example, whereas 'That dog is black' is not a generalisation, replacing 'that dog' by 'every dog', 'no dog', 'at least one dog', and so on, yields a generalisation. Sometimes the verb must be changed to the plural form, and likewise the predicate if it involves a noun rather than an adjective.

Gettier cases   Gettier cases are cases in which someone satisfies the conditions for knowing a proposition that are set down by the **tripartite account of knowledge**, yet fails to know it. This is usually

because the person is only accidentally **justified** in believing a **true** proposition.

Good reasons   For someone to have good reasons for believing a proposition is for that person to possess an argument for that proposition that is **rationally persuasive** for them.

Hard generalisation   A hard generalisation is one that is correctly conveyed by using a quantifier such as 'all', 'every', 'each', and 'no'. Unlike **soft generalisations**, such generalisations are true only if there are no **counter-examples**.

Implicit   A **premise** is implicit in an argument if it has been assumed but not actually stated by the arguer. **Conclusions** may also be implicit, though this is less common. Whether or not, as a matter of psychological fact, a given premise has been assumed by an arguer is often beside the point. In general, implicit propositions are those not stated by the arguer that would be included in a **reconstruction** produced in accordance with the **Principle of Charity**.

Implicit relativity   A statement is **implicitly relative** when the type of fact it expresses involves a relation to something that is not explicitly mentioned in the statement. For example 'John is tall' is implicitly relative because what it really means is 'John is taller than the average man' (if John is a man). The relation to the average man is not explicit in the original statement.

Implicit speaker-relativity   An implicitly speaker-relative statement is one that is **implicitly relative**, where the implicit term of the relation is the person making the statement. Thus the statement is **speaker-relative**, but only implicitly so. For example 'Chocolate ice cream tastes better than strawberry ice cream' is implicitly speaker-relative because what it really means is 'Chocolate ice cream tastes better to me than strawberry ice cream does'. (Some might say that the implicit term here should not be 'me' but something like 'most people').

Inductive force   The inductive force of an argument is the **conditional probability** of its conclusion relative to its premises.

Inductive inference   To draw an inductive inference is to conclude, on the basis that a certain proportion of a **sample** of a population possesses a certain feature, that the same proportion of the whole population possesses that feature. The inference is **inductively forceful** to the degree that the sample is representative of the population.

Inductive soundess   See **soundness**.

Inference   An inference is a step in reasoning from one or more premises to a conclusion. All **arguments** contain at least one inference. Inferences are evaluated not as **true**/false, but as **valid**/**invalid**, **inductively forceful**/inductively non-forceful.

Inference bar   In an **argument reconstruction**, an inference bar is a line written between a **premise** and **conclusion**, indicating that the proposition expressed below the line has been inferred from one or more of the propositions above it. It can be read as 'therefore'. Every reconstructed **argument** contains at least one inference bar.

Instance   An instance of a **generalisation** is a proposition about an individual that is directly implied by the generalisation. Normally this will be a **conditional**. For example, 'If Socrates is a philosopher, then he is wise' is an instance of 'All philosophers are wise'. The **inference** from a hard generalisation to a corresponding instance is always **deductively valid**. That from a soft generalisation to an instance is **inductively forceful**.

Intermediate conclusion   In an **extended argument**, an intermediate conclusion is a **conclusion** inferred from some set of **premises** that is used, in the same argument, as a premise for a further **inference**.

Invalid   Synonymous with 'not **deductively valid**'. An argument may be invalid yet **inductively forceful**, or neither valid nor inductively forceful.

Justification   One's degree of rational justification for believing a proposition is the degree to which one is entitled to think it true. In many cases this will depend on whether or not one has good reasons – arguments that are **rationally persuasive** for one – for thinking so. Rational justification can be distinguished from pragmatic justification, according to which one is justified in believing something if believing it has desirable consequences. Rational justification and pragmatic justification do not always coincide.

Knowledge   See **tripartite account of knowledge**.

Lexical ambiguity   An **ambiguous** sentence is lexically ambiguous if it contains an ambiguous word.

Logic   Logic is the systematic study of arguments, especially deductive validity and inductive force.

Glossary

Logical assessment   Logical assessment is the stage of argument assessment at which it is determined whether the argument is **valid** or **invalid**, and at which the degree of **inductive force** of invalid arguments is determined. Except where **inductive inferences** are concerned, it is sharply to be distinguished from **factual assessment**.

Practical reasoning   Practical reasoning is the use of arguments whose conclusions recommend an action. See **expected value**.

Premise   A premise in an **argument** is advanced as a reason for **inferring** the argument's **conclusion**.

Premise indicators   These are expressions such as 'since, 'because', and 'for the reason that' that are often used to indicate a **premise** in an **argument**. As in the case of **conclusion indicators**, however, such words sometimes serve other purposes, such as indicating a causal relationship.

Principle of Charity   According to this principle, if our aim is to discover the truth about a given issue, then we should reconstruct arguments so as to yield the maximum degree of rational persuasiveness for the relevant audience (which normally will include ourselves).

Probability   The probability of a proposition is the degree to which it is likely to be **true**, where this degree is expressed as a fraction or decimal between 0 and 1. There are different ways of explaining this, such as proportion and frequency; but in this book the degree to which a proposition is likely to be true is taken to be the degree to which it would be perfectly rational to expect it to be true. Since this obviously depends on the evidence one has, the key concept is that of **conditional probability**: the degree to which it is rational to expect a proposition to be true given such-and-such evidence.

Proposition   A proposition is the factual content expressed by a declarative sentence on a particular occasion of using (writing or uttering) the sentence. In particular, it is what is expressed that admits of being **true** or false. Different sentences can express the same proposition. For example 'Antony kissed Cleopatra' expresses the same proposition as 'Cleopatra was kissed by Antony'. Different propositions may be expressed by means of the same sentence. For example if Antony and Cleopatra each utter the sentence 'I'm hungry', they express different propositions, since they talk about different people. A sentence's prepositional content is independent of its rhetorical or emotive content.

Quantifiers Quantifiers are expressions such as 'all', 'some', 'every', 'many', 'twelve', 'not very many', and 'no', used in the explicit statement of **generalisations**. Often, where both hard and soft generalisations are concerned, they are left **implicit**.

Rational persuasiveness This is the concept we use to characterise a person's having 'good reason' for accepting a **conclusion**. An **argument** is rationally persuasive for a person if, but only if: (i) the person accepts its **premises**, and is **justified** in doing so; (ii) the argument is either **deductively valid** or **inductively forceful** and (iii) if it is inductively forceful, the argument is not **defeated** for that person.

Refutation by counter-example This is a method of criticising an argument. One fabricates an argument that embodies the same form or strategy as the argument one means to criticise, but which is clearly **invalid**, not **inductively forceful**, or **fallacious**. This shows that the strategy or form of the argument one means to criticise is defective in the same way.

Relevance The falsity of a premise can be more or less relevant, where the aim is to discover reasons for or against the conclusion in question. Sometimes a false premise can simply be removed from an argument without destroying its **validity** or degree of **inductive force**; the resulting argument may thus be sound. Other times the false premise can be replaced by another, similar premise that is true, again without compromising validity or inductive force. For example, the **scope** of a generalisation may be reduced.

Representative See **samples**.

Rhetoric A verbal or written attempt to persuade someone to believe, desire or do something that does not attempt to give **good reasons** for the belief, desire or action, but attempts to motivate that belief, desire or action solely through the power of the words used, is an instance of rhetoric.

Rhetorical force Not part of the **proposition** that a sentence expresses, but the emotive or otherwise suggestive window-dressing that surrounds that **proposition** and is used to persuade us to believe or do something by appeal to our non-critical faculties.

Rhetorical ploys Commonly encountered instances of rhetorical use of language these include: appeal to novelty, appeal to popularity, appeal

to compassion, pity or guilt, appeal to cuteness, appeal to sexiness, appeals to wealth, status, power, hipness, coolness, etc., appeal to fear (also known as scare tactics), the direct attack and hard sell, buzzwords, scare quotes, trading on an equivocation, smokescreen (changing the subject).

Rhetorical question An interrogative sentence that is not really intended as a question, but as a statement, usually of a **proposition** with which the speaker or writer assumes the audience will agree.

Samples An **inductive inference** is inductively forceful if, but only if, the sample cited in the premise of the inference is representative of the population cited in the conclusion of the inference. Suppose that n is the proportion of the sample known to bear a certain trait. We can reasonably conclude that the proportion of the population bearing that trait is n only if we have good reason to think that nothing has caused that proportion in the sample that does not equally affect the population as a whole.

Scope To speak of the scope of a generalisation – for example one of the form 'All X are Y' – is a way of discussing the size of the class X. For example the scope of 'All dogs are friendly' is wider than that of 'All beagles are friendly'. See **relevance**.

Sham-reasoning Attempts to persuade which appear to give reasons for accepting/rejecting a claim but in fact do not give us any reason for doing so are instances of sham-reasoning. **Fallacies** and **rhetorical ploys** are types of sham-reasoning, the former argumentative sham-reasoning, the latter non-argumentative.

Soft generalisation A soft generalisation such as 'Most dogs are friendly' is contrasted with a **hard generalisation** such as 'All dogs are friendly'. Unlike hard generalisations, soft generalisations cannot be refuted by a single counter-example. Soft generalisations are frequently expressed without an explicit **quantifier**, as in 'Siamese cats meow a lot'.

Soundness An argument is sound if, but only if, its premises are all **true**, and it is either **deductively valid** or **inductively forceful**.

Speaker-relativity A kind of statement is speaker-relative if it expresses a different **proposition** depending on who makes the statement. For example if John says 'I am left-handed' he says that *he* is left-handed,

whereas if Mary says 'I am left-handed' she says that *she* is left-handed. The two propositions, then, could have different **truth-values**.

Stances (towards a proposition)   There are four possible stances that can be taken towards a **proposition** with which one is presented: believing it, not believing it, suspending judgement, not engaging with it. Believing and not believing a **proposition** admit of degrees dependent upon the evidence available to provide **justification** for belief in that **proposition**.

Standard form   This is a style of displaying an argument's **premises**, **conclusion**, and inferences, in which each proposition is enumerated, and each **inference** indicated by an **inference bar**. For example:

> P1) If Mrs McFee were murdered, there would be signs of a struggle.
> P2) There are no signs of a struggle.
> _____
> C1) Mrs McFee was not murdered.
> P3) If Mrs McFee was not murdered, then Inspector Radcliffe ought to leave the premises.
> _____
> C2) Inspector Radcliffe ought to leave the premises.

Syntactic ambiguity   An **ambiguous sentence** is syntactically ambiguous when the arrangement of its words is such that the sentence could be understood in more than one way (as expressing more than one proposition.) See also **lexical ambiguity**.

Tripartite account of knowledge   According to this account of knowledge, someone counts as knowing a proposition if and only if (i) they believe it, (ii) it is true and (iii) they are justified in believing it. See also **Gettier cases**.

Truth   To say, of a **proposition**, that it is true, is to say that things are as the proposition says it is. For example, to say that it is true that snow is white is to say that snow is as that proposition says it is, namely white. Another way to put the point: To say that it is **true** that snow is white is equivalent to saying that snow is white. Unless a statement is **speaker-relative**, the proposition that it expresses cannot be true for one person but false for another.

Glossary

Truth-value   Sometimes it is convenient to speak of the truth-value of a proposition (see for example **speaker-relativity**): The truth-value of a **true** proposition such as 'snow is white' is truth, and that of a false proposition such as 'snow is green' is falsity.

Vagueness   An expression is vague if (i) its extension is indefinitely bounded, e.g. 'bald', 'tall' or (ii) in a given context it is unclear what is meant by it. In the first case the meaning of the term may be clear, but the extension is 'fuzzy' because there is no clear point at which things arranged on a single dimension (e.g. degree of baldness) cease to have the quality denoted by the term. In the second case particular uses of a term such as 'political' are vague: it may not be clear exactly what is meant by calling something political in a given context. Because of this, a word such as 'political' may be vague in both senses: because its meaning is vague, so is its extension.

Validity   See **deductive validity**.

■ Index

Index