6 Defining Knowledge

1. Top-grade Belief

Consider three cases in which a person does not have knowledge.

A weather forecaster, McA, is preparing to make a radio report on the weather to be expected the next day. He finds, five minutes before he must go on air, that he has lost all the data on which the report was to be based. So he tosses a coin. Heads for rain, tails for dry. Heads for winds under five miles per hour, tails for winds over five miles per hour. And so on. After five minutes he has his report, and he goes on air and predicts a hurricane, though the sky is blue, the barometer is steady, and the breezes are gentle.

McA does not *know* that there will be a hurricane the next day. Suppose that by some freak of the weather there is a hurricane the next day. That does not turn his report into knowledge; it was just guesswork, and only turned out accurate by chance.

A detective, FitzB, is investigating a murder. Near the victim's body she finds a wallet containing a driver's license in the name of Z. She goes to interview Z and finds that he has left town in a hurry. Z is tracked down and found to have traces of the victim's blood under his fingernails. Moreover, his fingerprints are all over the scene of the crime. A little investigation reveals that Z had a long-standing grudge against the victim and stood to gain a lot of money from the victim's death. There are no other

suspects. FitzB concludes that she has solved the crime: Z is the killer.

But Z is not the killer. Z was framed by a very clever enemy of the victim who managed to plant all the evidence while leaving no traces of her activities. Z is convicted and executed, and the real killer is never brought to justice. So FitzB did not know that Z was the killer. She thought she knew, and she had very good reasons for so thinking. But in fact her belief was false, so it cannot have been knowledge.

A scientist, Dr O'C, is testing a drug, magicillin, which she thinks may be a good treatment for cancer of the liver. Suppose that her suspicion is right, and magicillin does cure liver cancer in a large proportion of human cases. Suppose, moreover, that she performs experiments on both human and animal subjects, designed to measure the rate of recovery when the drug is administered and to rule out alternative explanations. Suppose that her reasoning from these experiments is faultless, so that when she publishes her results in a research paper, they provide very good reasons to believe that magicillin will save the lives of a large proportion of people with cancer of the liver. Suppose now that, although Dr O'C does not know it, one of her lab assistants has falsified part of an experiment by replacing one experimental animal which looked as if it might be diseased with a healthy animal that had not previously been part of the experiment.

When Dr O'C finds out, she will be furious and dismayed: her research has been ruined. After she learns that the experiment has been tampered with, she can no longer maintain to her colleagues that she knows that magicillin cures liver cancer. She does not know it; she did not know it even before she discovered what her assistant had done, and thereby discovered that what she had thought was knowledge was not. Yet her belief is true – magicillin does cure liver cancer – and her reasoning is perfect – her belief is justified.

In all three cases we have a belief that fails to be knowledge. But the reasons for this are different. McA does not know that there will be a hurricane the next day, because what he passes off

as reasoning is so feeble. You cannot predict the weather by tossing coins. So the fault here lies in him. FitzB does not know that Z is the killer, because her belief is false. As it happens, Z is not the killer, although FitzB's reasoning is fine. The fault here lies in the relation between her beliefs and the facts.

Looking at the McA and FitzB cases, we might think that it is obvious what knowledge is. To have knowledge, you have to have a belief which is both true and justified. Both the facts and your reasoning have to be right. But the O'C case shows that things are not that simple. Dr O'C's belief is true, and her reasoning is fine, but she does not have knowledge. The O'C case is one of many examples – known as Gettier examples after Edmund Gettier, who first pointed out their significance – which show that we cannot take 'know' to mean simply 'true and justified'. When we say that someone knows something, we are saying more than that they believe it and their belief is true and they are justified in believing it. What more?

This is a controversial question. There is very little agreement about what the right answer to it is, or even whether it is an important question. In this chapter I shall describe several approaches to the question which have been defended by philosophers since the 1970s. The chapter is thus an introduction to a part of the theory of knowledge with a very distinctive flavour of its own. Some people find the questions and approaches here fascinating; others have trouble seeing the point of them. My aim is to make a case for thinking that there are important and central questions here. But before entering this controversy, it is worth stating some simpler, uncontroversial facts about the concept of knowledge, to give us something definite to hold on to.

A belief qualifies as knowledge if, in acquiring it, one has achieved the basic aim in the enquiry that led to it. In other words, we use the term 'knowledge' as a kind of stamp of high quality for beliefs: if a belief fails to be knowledge, it is because it falls short of some expectation we might reasonably have of it. Thus McA's belief that there would be a hurricane the next day,

and the beliefs of the people who believed what he said on the radio, were not knowledge, because we expect weather forecasters to pay more attention to the evidence than he did. FitzB's belief that Z is the killer is not knowledge, because we want beliefs to be true – that is one of the most basic aims of having beliefs – and it was not. Dr O'C's belief that magicillin cures liver cancer is not knowledge, partly because we cannot take the evidence and reasoning she presents as grounds for believing it ourselves: once we know that some of the data have been tampered with, we know that the experiments are flawed and will have to be rerun and reanalysed.

This general picture holds in many ordinary cases where we say that someone knows something, too. When we classify a belief as knowledge, we are saying that it has the good qualities that a belief should have. When we say that a student knows the answers on a test, we mean that the answers she gives are true and that she is not just guessing. (We mean more than this, but at least this.) When you are asked what the capital of Colombia is and you say 'I don't know, but I think it's Bogota', you are saying that we should not treat your belief that Bogota is the capital as a completely safe basis for our beliefs. When we say that the police know that a certain person has committed a certain crime, we mean that the police have done their job successfully; they have the right person, and they have discovered him by sensible detective work. Moving on to rather different ways of talking, when we say that someone knows about marathon running, we are saying that he has beliefs which make him a good source of information for us to base our beliefs about marathons on. And when we say that someone knows which the best restaurants in town are, we are saying that she can supply descriptions of the restaurants which can both help us to get a good meal and help us to find them. Moving on even further, when we say that someone knows how to recognize a good restaurant when she sees it, we mean that in the presence of good restaurants she will form the belief that they are good. Her skills are ones that she and others can base reliable beliefs on.

These are only some of the good qualities that beliefs can have. One way of understanding the definitions of knowledge that are discussed in the rest of the chapter is as ways of classifying and organizing all these desirable features of beliefs. Which are the most fundamental and most important of them?

2. Lehrer's Principle

Dr O'C's reasoning, though correct, began with a false premiss, that the animals with which the experiment ended were those with which it began. Perhaps knowledge is true belief that lies at the end of a chain of reasoning which begins with true premisses.

This suggestion is not right either. For suppose that Dr O'C does other experiments. This time her experiments don't get fouled up. But unfortunately they don't directly indicate that magicillin cures liver cancer; instead, they indicate that it inhibits cell division in a certain complicated way. From this result she infers that it will be effective against cancer. Now as a matter of fact (in this story) magicillin is effective against liver cancer, but it does not inhibit cell division in the way that the experiments indicated. Again things have gone wrong; again Dr O'C has not got knowledge, though this time her belief is not only true, and obtained by good reasoning, but also proceeds from true premisses.

It seems, then, that for a chain of reasoning to lead to knowledge, it must not only start with true premisses but must also not have any false steps at any point along the way. It should have no detour through falsity. We have here a fairly simple principle that does seem to correspond to something important about knowledge. Let me call it **Lehrer's principle** and state it more carefully: if a belief is based on reasoning, then it qualifies as knowledge if, and only if, it does not depend on any chain of reasoning that has a false step at middle, beginning or end.

A number of philosophers have accepted Lehrer's principle. But, like many valuable philosophical ideas, it contains an enormous fudge. For it says that knowledge cannot *depend on* any reasoning that proceeds through falsity. What does 'depend on' mean?

A belief does not necessarily depend on the reasoning by which it was originally acquired. One can first acquire a belief for some reasons (perhaps bad), and later obtain other reasons (perhaps good) for holding it. For example, a person may hold a belief just because someone he dislikes believes the opposite; later he makes some investigations and finds some strong evidence for the belief he already has. Now he may have knowledge, although his first reasons were bad. (Another example is the case of Harvey's belief in the circulation of the blood, mentioned in chapter 1, section 3.)

Nor does a belief depend on just any reasoning which *could* support it, whether or not the person holding the belief uses or is aware of the reasoning. For often something that one clearly knows can be obtained as the conclusion of reasoning which has the fatal characteristics that it is logically correct, one believes each of its steps, but some step is false. This false step should not cause one's belief not to be knowledge if one does not actually use the fatal reasoning to get or support it. For example, Dr O'C may have the false belief that magicillin inhibits cell division; but suppose that, although this might be used in a persuasive argument to conclude that it cures liver cancer, she does not base her belief upon this reasoning. She bases her belief instead on a line of reasoning which does not involve any false step. She then knows that magicillin cures liver cancer; her knowledge is not undermined by the faults of the reasoning she does not use.

We must take 'depends on' to mean roughly the following: A person's belief depends on a chain of reasoning when the reasoning is used by the person, publicly or privately, to show why he should hold the belief; if the reasoning is refuted, if the person stops thinking that it is correct or that all its stages are true, then the person will either find some other reasoning to support the belief or cease to consider it as one of his well-supported opinions.

3. Knowledge not Based on Reasoning

Lehrer's principle applies to beliefs which depend on reasoning. When I tried to explain what 'depends on' must mean, it turned out that the reasoning had to be fairly explicit; the person must be able to formulate it to himself. Many beliefs are not based on any such reasoning, and most beliefs result in part from processes that are not explicit reasoning. Perception, for example, though it is far from a simple unstructured process, does not consist in reasoning that the perceiver can express. So all our beliefs based partly on perception rest partly on processes that are not explicit reasoning.

Moreover, it is rational to hold many beliefs not because they are supported by any reasoning, but because they explain, unify or predict one's experience and opinions. This is the case with many systematic or theoretical beliefs and with much of the knowledge that we get by learning from others. Consider, for example, the knowledge that the earth is round or that germs cause certain diseases. Such knowledge is useful, and one would be foolish to give it up just because one does not have at hand an argument that would convince someone who thinks that the earth is flat or that all diseases are caused by lack of faith. (See the discussion of ignorance-avoidance in chapter 5.)

No doubt in many cases one could invent an argument which would do part of the job, if one thought hard enough. It would probably be some version of the 'inference to the best explanation', discussed in chapter 10. It would run roughly: this hypothesis explains the following facts, and fits in with these other facts in the following nice way; the alternative hypotheses for explaining these things won't do, for the following reasons, so it seems to be true. (I say that such an argument might do 'part of the job' because someone who believes, say, that the earth is flat, is likely to have enough other strange beliefs about the way the world is and what a scientific hypothesis should do that he will not be much affected by any such argument.) It has often been taken as

one of the functions of philosophers to provide such sophisticated arguments for things most people believe without an argument. But in any case we saw that Lehrer's principle is not concerned with reasoning that could be used to support a belief, but rather with reasoning that is actually used; and it is clear that although most people know that the earth is round and that typhoid is caused by a bacillus, very few people can make convincing arguments for these statements.

4. Reliability

Beliefs to which Lehrer's principle cannot be easily applied are also obtained from theories which give reliable predictions although they are in fact partly false. Here is an example. Before the discovery and acceptance of the idea that the earth revolves around the sun, rather than the other way around, sailors used to compute their latitude from the position of the sun and other heavenly bodies by using a theory which falsely assumed that the earth was stationary. As a matter of fact, sailors got the right navigational results this way. Consider now an old sea captain who computes his latitude (how far north or south of the equator he is) from observations of the heavens, and imagine him to have enough astronomical lore to think things through in terms of the old false astronomy. Surely, if he gets the right answer, if he makes no mistakes and concludes that his latitude is what it really is, then he knows what his latitude is. But his reasoning involves many false assumptions. Lehrer's principle seems wrong here.

The important thing about this case is that the captain has obtained his result by a method that is reliable – it always works – and that could have been based on true assumptions, though it was not. For the captain's calculations could have been derived from the true astronomical facts, or even on the non-committal assumption 'If the apparent positions of the stars in the sky are such and such, then the latitude is the following'. It is important that it is a whole method that can be given a true foundation, and

not just a single belief. For if the captain had obtained his estimate of his latitude from the horoscopes of the crew, and by pure luck had come up with the right answer, then the fact that he could have derived his latitude from sound navigational principles does nothing to confer the status of knowledge on his guess. What is special about the captain's case is that his method is reliable, and it is no accident that it is reliable; given the way things really are, it is inevitable that the method will give true results.

The case of the old sea captain is important because it leads to another side of the concept of knowledge, as basic as the side expressed by Lehrer's principle. The captain knows where he is because his method of discovery is *reliable*: there is a good reason why it gives the right answers. However rational or irrational he is in using his method, others can rely on his results; seen as a source for other people's reasoning or as part of a larger scientific project, his method is just as it should be.

The reliability of methods is a subtle business. Another story: in this story, unlike the earlier one, Dr O'C has discovered that magicillin works against all cancers in rats, mice, gerbils and squirrels; so she concludes that it works against cancer in all rodents, and therefore will be effective in chipmunks. This reasoning is sound, and, as it happens, magicillin is effective against cancer in chipmunks. But the way that it cures cancer in most rodents is completely different from the way it works in chipmunks; by some freak of chipmunk metabolism, it is the substance in which the normally active ingredient of magicillin is dissolved, rather than that ingredient itself, that is effective against cancer in chipmunks. Dr O'C does not know that the drug works in chipmunks, although her reasoning does not violate Lehrer's principle and is trustworthy, in that it leads, and will usually lead, to a true conclusion. What is wrong with her belief is that its truth is a sort of accident: the reasons why her roughly inductive reasoning usually works do not explain why it worked this time.

'Reliable' is a word that can be interpreted in several different ways. One way is statistical: a reliable method is one that gives

true beliefs a high proportion of the time. (So we could speak of 75 per cent or 98 per cent reliability.) Statistical reliability is probably not the right way to think about knowledge. For suppose you have a ticket in a lottery and 999,999 other people have also bought tickets. (See chapter 4, question 14.) If your method of getting beliefs is to believe everything that has a more than 99 per cent chance of being true, then you should believe that you will not win the lottery. Suppose you do not win. Although this was the almost certain outcome, it does not seem right to say that you knew in advance that you would not win.

Another way is causal: a belief is reliably acquired if the laws of nature and mind explain why the process that resulted in the belief resulted in a true belief. So reliability requires a firm connection between the fact and the belief, as a result of the ways in which information is transmitted and processed and the effects of any disturbing factors. In fact, a process can be reliable, in this way, even if it operates only once and cannot be repeated. For example, I may have been blind since infancy (after a brief sighted period in which my visual system developed somewhat); then, late in life, the injection of a powerful drug allows me to see for a few moments, before it utterly rots away my retinas. During those few moments I know that there are people about me, what their movements are, and the colours of their clothes.

The idea is still vague. A belief is reliably acquired if there is an explanation of why the process that resulted in it resulted in truth. What process? Does Dr O'C's belief result from thinking, from scientific method, from the particular experimental design she used, or what? Any way of spelling out what the relevant process is will appeal to conditionals: that is, if-statements. Suppose we say that Dr O'C's belief is reliably acquired if the actual experimental design she used was reliable. This means that (given the way human minds and the rest of the world work) if that design had been used in various circumstances, the result would have been a true belief. A reliable method is one that, when used in the right circumstances, will not give a false belief. (Perhaps in many circumstances, it gives no belief at all.) So a more precise

way of talking about reliability would specify a range of circumstances and of beliefs, and then say that if the person was in those circumstances, then those beliefs would be true.

One specific form of this idea is the concept of *tracking*. Consider the most basic kind of knowledge: perception. Suppose that you are watching a bird fly across your garden and land in a tree. As the position of the bird changes, so does your perception, and thus your perceptual belief. If the bird had followed a different path, your sequence of beliefs would have been different. Your beliefs therefore track the position of the bird, much as the eyes of a predator track the position of potential prey, or a blip on a radar screen tracks the position of a plane.

Tracking can be defined in terms of conditionals, if-sentences. If the bird had followed a different course, your beliefs would have been different. And before the bird flew the way it did, it would have been true to say that if it were to fly in the way it did, then you would believe that it was flying that way. Generalizing, we can turn this into a definition:

> A person's belief *tracks* a fact when two conditions are met: (1) if that fact had occurred (perhaps under slightly different circumstances), then the person would have believed that it occurred; (2) if the fact had not occurred, the person would not have had the belief.

This definition is essentially due to Fred Dretske and Robert Nozick. They suggested that we can define knowledge in terms of tracking: a person's belief counts as knowledge when it tracks the fact that makes it true. So, for example, you know that the bird is half-way to the tree when your belief that it is half-way there tracks the fact that it is half-way there. And this will be so when, if the bird had only been a quarter of the way there, you would have believed that it was a quarter of the way there, and if it had been three-quarters of the way there, you would have believed that it was three-quarters of the way there.

This approach to knowledge is clearly a version of the reliability approach. Reliability, remember, depends on whether there are good reasons, in the workings of the mind and the routes by which information comes to it, why true beliefs should result. So if a belief is based on a reliable process, then the process should give the belief only when it is true, and should give different beliefs in different situations. In other words, the belief should track the fact that makes it true.

There is something intuitively very right and appealing about this approach. It captures the feeling that when someone knows something, then it is no accident that their belief is true: it is true because it tracks the facts. It captures the feeling that a person who has knowledge can be used by others as a reliable source of information: for if you know that the person's beliefs track the facts, then you know that if you trust what they say, your own beliefs will be true. And it captures the feeling that beliefs based on irrational reasoning or mistakes are not knowledge, even if they happen to be true. For such deviations from correct thinking usually introduce factors which make the belief independent of the facts: even if the facts had been different, the person would still have held the belief.

5. Missing Information

But there are still problems, more tricky examples. One family of problem cases involves beliefs which track the facts but do so because of an accident. Suppose, for example, that you are in a psychological experiment like one of those described in chapter 2. You are given eyeglasses which make everything look upside-down. But you do not know that this is what they are; you do not even know you are in a psychology experiment. (Perhaps the experimenters tell you that they are comparing different glasses' frames for comfort.) The glasses are put on you in the dark, and when the light comes on, you see a scene through a window. By

mistake a mirror has been left by the window which inverts the scene, exactly cancelling the effect of the glasses. In the scene a balloon is floating, and you are asked if it is rising or falling. You say that it is rising, as it in fact is. Your belief is true, and tracks the facts: if the balloon had been falling, you would have said that it was falling; if it was stationary, you would have said it was stationary, and so on. But the set-up that allowed your belief to track the fact was a fluke; all the subjects in the experiment before and after you believed that the balloon's motions were the opposite of what they in fact were. Many people when they think about this case will conclude that you do not know that the balloon is rising. Although your belief tracks the fact that it is rising, you could too easily have had a false belief. Knowledge should be more secure than this.

One feature of this case is very significant. Suppose that you had been given the true information that you were being fitted with inverting glasses. Then, when the balloon appeared to rise, you would not have believed that it was rising. You would have thought: 'I am in an experiment, and they are doing weird things to my vision; who knows what may really be happening?' So there is true information you do not have which would have prevented your having the belief, even though the belief was true. This fact is closely linked to the accidental way that the belief tracks the truth. For although there is a variety of circumstances in which it is true that if they held you would have true beliefs about them, there is also a slightly larger variety of circumstances in which you would not have true beliefs. In many cases in which tracking is the result of an accident, information about the accidental features of the situation can create circumstances in which the person's beliefs would be false.

The failure of a belief to be knowledge is often connected with the fact that it would not have survived exposure to some true information. For example, suppose that you are walking around central London thinking about the troubles of the British royal family. Around the corner walks a woman whose appearance you have often seen on television and in newspapers; you go back to

your hotel to tell your friends, 'I just saw Princess Diana'. But there are facts that you do not know. On that day anti-royalist protesters have flooded the streets with Diana look-alikes. There are hundreds of them in the area. Moreover, this fact has been well publicized; if you had been reading serious newspapers instead of tabloids, you would have realized it. Now, as a matter of fact, the woman you saw was (surprise surprise) Princess Diana. But, according to most people who reflect on cases like this, you do not know that it was her. For you could as easily have encountered one of the look-alikes, and then you would have been fooled.

Or consider a more tangled version of the same example. Again you are walking around London, and again you see the princess. But this time there are not hundreds of look-alikes out and about. Instead, there are newspaper headlines visible all over the place saying, falsely, that there are hundreds of Diana look-alikes wandering the streets. If you had seen any of these headlines, you would not have believed that the person you saw was the princess. And it is a pure accident that you do not see any of these many headlines. In this case too people's reactions are usually that you do not know that the person you are seeing is the princess. And one explanation of this reaction is that the fact that if you had seen the headlines you would not have believed that the woman you saw was the princess shows that your belief did not rest on a reliable process. Although it gives a true belief in this case, there are very similar circumstances in which it does not.

It is interesting to make two contrasts here. On the one hand contrast your situation – seeing the princess but in the unseen presence of look-alikes or misleading headlines – with that of a close friend of the princess, who could tell her from all the look-alikes. The friend would know the princess on sight, even if he knew about the look-alikes or had seen the misleading headlines. On the other hand, contrast your situation with that of someone who knows the princess no better than you do, but just happens to see her coming out of an expensive London shop, when there

are no look-alikes about or misleading headlines to be glimpsed. This person knows that it is the princess before her, just by using her ordinary powers of perception and recognition. So the concept of knowledge has a really subtle and interesting combination of strictness and tolerance. It is tolerant in that we can often know things by use of our ordinary capacities to perceive, reason and remember, even though these capacities sometimes lead to false beliefs. It is strict in that it does not allow that every use of these capacities which leads to a true belief results in knowledge. Sometimes we have to consider a much more specific process operating in a much more specific context, and ask: 'Is this reliable?'

There are many examples along these lines, in which people's beliefs are not knowledge not because of any false or irrational belief they have but because there are true beliefs they lack. Such examples are particularly troublesome for accounts of knowledge based on reliability or tracking. What they show is that if we want to define knowledge in terms of the conditions under which the process behind the belief results in truth, we are going to have a delicate job specifying the right process. If we define it too narrowly, as tracking analyses do, then a belief can fail to be knowledge because there are wider circumstances in which the person would have had a false belief. (For example, we ignore the effect of the look-alikes or misleading headlines in the Diana cases.) If we define it too widely – for example, by counting reason or perception as belief-giving processes – then we will run into the problem that these processes often give false beliefs. (So then we wrongly fail to count as knowledge the belief of a normal person under normal circumstances who sees and recognizes the princess.) We will need a definition narrow enough that the facts of the case can explain why this belief obtained by this process in these circumstances is true, and wide enough that the belief can be seen as a special case of a general truth-producing process. The right balance has not yet been found.

6. Internalism/Externalism

The many examples that philosophers have produced to support and refute definitions of knowledge have not resulted in a completely correct definition. Not yet, anyway. But they have brought out a large number of important qualities of beliefs. We want beliefs to have these qualities, and will withhold the label 'knowledge' if they are absent. We want beliefs to have all of the following qualities:

truth	justification
reliability	coherence
fact-tracking	reasonability
usability by others	not undermined by others' beliefs

There is a fundamental contrast between the qualities on the left and those on the right. The qualities on the right are *internal* qualities. They describe ways in which beliefs relate to the functioning of the individual, aspects of the individual's reasoning and perceptual processes. The qualities on the left are *external* qualities. They describe ways in which beliefs relate to the world around the individual. Some theories of knowledge, **internalist** theories, emphasize internal aspects and qualities, while others, **externalist** theories, emphasize external aspects and qualities. (See the discussion of radical externalism in section 5 of chapter 1.) The more a theory emphasizes justification, evidence and reasoning, the more internalist it will be, and the more it emphasizes reliability and the objective conditions under which a person's belief will be true, the more externalist it will be. One fundamental question is: Is knowledge basically a matter of the internal or the external qualities of a belief?

The obvious answer to this question is: Both. The old definition of knowledge as justified true belief was in agreement with this answer, since it emphasized one external quality (truth) and

one internal quality (justification). Yet this may not be a promising way to answer the question; it may duck important issues. Some considerations about scepticism will make clear why this might be. Consider two sceptical possibilities.

First, suppose that our powers to reason are not nearly as reliable as we think they are. This is easiest to imagine with inductive reasoning: suppose that very many of the generalizations that we think we have conclusive inductive evidence for are in fact false, and that the true generalizations are too subtle and complicated for humans to understand. But a similar story is on the edge of imaginability for deductive reasoning. The reason, let us suppose, why our powers of reasoning have not failed us very often in the past is that we have been exposed to only a small part of the universe. If we were to explore just a little further out in time or space, or if some simple fact about the environment were to change, we would suddenly find that our attempts at reasoning consistently produced false results.

Would this show that very few of our true beliefs – our beliefs about the little part of reality where reason is effective – are knowledge? No, according to internalists. For our beliefs are reasonable, and they result from our best possible attempts to understand the world. To claim that they are not knowledge would be like claiming that kindness and fairness are wrong just because the world is cruel and unfair. But Yes, according to externalists. These facts, if they really were so, would show that the truths among our beliefs have come about by lucky chance, not because we are well equipped to understand the world.

Now consider another improbable story. Suppose that in some part of our lives our rational explanations about how our beliefs are obtained are completely mistaken, but many of our beliefs are true all the same. For example, suppose that many of each person's beliefs about other people's minds – whether other people are angry or disappointed or in love and so on – are true, but that our beliefs that we understand these things by reasoning in a rational way from other people's behaviour and facial expressions are wrong. In fact, suppose we gain these beliefs in a purely

instinctive way based in part on how people smell and the rhythms of their body movements. Does this mean that we do not have knowledge about these things (about other people's minds)? Internalist theories allow a sceptical Yes. Although many of our beliefs on this topic may be true, they are irrational and amount to lucky hunches. Externalist theories, on the other hand, allow a more comforting No. Irrational though they may be, our methods for obtaining such beliefs work. They give us reliable truths.

So internalist and externalist approaches to knowledge can lead to sceptical conclusions, given different assumptions. Internalist approaches emphasize the proper working of our capacities to reason and process information. When these capacities are working properly, internalism will count the result as knowledge, and when they are not, it will tend to scepticism. Externalist approaches, on the other hand, emphasize the connection between belief and the world. If these are operating reliably, externalism will be ready to ascribe knowledge, and if they are not, its attitude is likely to be sceptical.

Internalism and externalism are thus fundamentally opposed attitudes to knowledge. They differ in the kinds of analysis of the concept that they present, the qualities of beliefs which they consider centrally important, and the conditions under which they will allow drastic sceptical conclusions. They are too different to be easily reconciled; a bland acknowledgement that each of them has an angle on the truth does not do justice to the conflict.

7. Knowledge and Trust

To end this chapter, I shall briefly describe some very basic qualities of beliefs that underlie the qualities in both columns of the list in the previous section. These do not lead to a definition of knowledge. But they do allow us to pull together some of the ideas of the chapter, and see why the concept of knowledge is so subtle. These ideas are controversial, though; they are more one

person's suggestions than a report of the consensus among philosophers.

People rely on others for information. An isolated individual using the evidence and ways of thinking available to us could know very very little. There are two ways of using a person as a source of information. The simple way is to trust what they say. If someone says that she saw a moose walking down the street at two in the morning, then you believe that at two in the morning a moose walked down the street. The other way is to take what someone says as evidence about their state of mind. You conclude that she believes a moose walked down the street at two in the morning, or perhaps that she believed she could get you to think that a moose had walked down the street. On the first of these ways, you assume that she knows what she is saying. On the second, you simply talk about her beliefs. So one thing we can mean when we say that what someone has is knowledge is: you can trust what they say.

Now consider a rather different use of knowledge. We build our beliefs on other beliefs, using what we already believe as evidence and as background for new beliefs. We do not use all our beliefs in the same way here. We treat some of our beliefs as definite facts about the world, safe to use as guides for future beliefs, and we treat others as conjectures, which we are prepared to revise when more information comes in. (Notice that safety is not the same as certainty or as being apriori. We can treat a scientific theory – for example, the theory of evolution – as a basis for evaluating later beliefs, although there is more evidence to consider for or against it, and although we can see that possible evidence could refute it.) The difference between the two kinds of beliefs becomes particularly clear when evidence against a belief appears, and we have to think out how to react. You believe that all mammals have live young and then you discover the platypus, which has the appearance and metabolism of a mammal but lays eggs. In reconsidering your beliefs, you think: perhaps the platypus is not a mammal, or perhaps some mammals lay eggs, but at any rate we know that all mammals in Europe and the Americas have live young, and that the placenta appeared very early in the evolution of mammals. When you think 'at any rate we know', the concept of knowledge is being used to set aside some beliefs which you do not take to be threatened by the evidence at hand. So another thing you can mean by saying that something is knowledge is: we can take this for granted in evaluating the evidence for and against other beliefs.

The most interesting aspect of these two implications of the concept of knowledge is that they are essentially the same. If we treat someone as an authority on a subject-matter, taking them as someone who has knowledge about that subject, then we will take what they say on the subject as definite fact rather than conjecture. (You are wondering what to believe about the platypus, so you consult a biologist. If you treat her as a source of knowledge, then you revise your beliefs in ways that do not challenge what she says.) And, the other way round, if we treat some of our beliefs as definite facts, then we are taking ourselves to be authorities with respect to those facts. (In the morning you seem to remember a moose walking down the street at 2 a.m. If you wonder whether you were dreaming or hallucinating, but subsequently decide to accept your own memory-beliefs as facts, you are in effect saying: I know what happened to me last night, so I'll trust what I believe.)

I suspect that this is the core of the concept of knowledge. Knowledge is based on those properties which beliefs have when we can treat them as authorities in evaluating other beliefs, holding them constant as new evidence accumulates and new arguments are considered. This feature of knowledge explains why in many cases we allow or withhold the label. For example, Lehrer's principle generally applies because when we learn that a belief is based on false evidence, then even if we do not come to doubt the belief, we cannot take it for granted in further reasoning. The evidence is flawed, so the belief cannot be trusted. Sometimes, though, we can see how false assumptions on which a person's belief is based will not make it dangerous for us to use this belief ourselves. For example, in the case of the old sea

captain the false astronomy that he uses does not make his beliefs about his latitude unreliable. We can see that as long as we stick to beliefs about navigation, the falsity of his beliefs about astronomy are not going to matter. So we say that he knows where he is. (We say this because we understand enough about both navigation and astronomy to see that he will get the same answer if he uses his false beliefs or the truth. But if it is explained to him that all his beliefs about the movements of the earth, sun and stars are wrong, he may conclude that he does not have knowledge.)

Reliability theories of knowledge fit into this picture in an obvious way. When we think that someone's belief is based on a reliable process, we think that they can be used as an authority on which to base our beliefs. We can trust what they say. We can do this as long as their beliefs are on the topic about which they are reliable. (What exactly the extent of this topic is, is a hard question, as we saw in sections 5 and 6.) The misleading true information problems for reliability theories also now make sense. Although if you are an ordinary person under ordinary circumstances, your belief that a certain very famous person is walking down the street can be used by yourself and others as authoritative, if there is misleading information around, then someone else using your beliefs as a basis for her own would have to be sure that you had much better than average recognition skills, the kind that only a really good friend of the famous person would have, before trusting your report. In fact, you yourself, when you later learn about the misleading information, will want to reconsider your report, and will wonder whether you could have distinguished the celebrity from the look-alikes.

So we have a formula: to take a belief as knowledge is to take it as something that you can trust when forming other beliefs. But this is not a definition of knowledge. It is much too vague, for one thing. In addition, it is a formula not about what actually *is* knowledge, but about which of a person's beliefs that person or another will *treat as* knowledge. The nearest it brings us to a definition is this: it is reasonable for one person, A, to take another person, B, to know something if from what A knows of

B's situation it is reasonable for A to share B's belief. Or, another way of saying the same thing, A should take B to know something when A can take B to be a reliable source of beliefs like the one in question, where what counts as 'like' is determined by B's purposes.

Notice that these definitions refer to the purposes of the person who classifies another person's beliefs as knowledge. (Or, very similar, to what it would be reasonable for that person to think.) This element of purpose is often quite clear when we attribute knowledge. Suppose, for example, one bank robber says to another: 'The police know it was us who pulled the Chase Manhattan job.' The force of 'know' here is that the police have information about their identity of a nature and reliability that is relevant to tracking them down. (Compare the examples of special knowledge about police work, restaurants and other topics towards the end of section 1.) When we try to make a general abstract definition of 'person P knows fact F', one of the problems we are tackling is that of finding a single, core purpose for having beliefs. Is there a unified aim of enquiry in terms of which we can say what counts as a top-quality belief?

Reading Questions

1 The McA case contrasts with the FitzB case. One case shows that truth is required for knowledge, and one shows that justification is required. Which is which?
2 The O'C case shows that truth and justification are not enough for knowledge. Why is Dr O'C's belief justified?
3 Suppose that a person is listening to the radio and hears McA's weather forecast. As a result she prepares for a hurricane. And a hurricane arrives. How does Lehrer's principle apply to her belief?
4 A father believes that his son is wonderful. He defends his belief with good school reports and his son's large number of friends. He also believes that any son of his has to be wonderful. How could we discover whether his belief that his son is wonderful depends (in the sense described in section 2) on the school reports and

friendships, on his belief that his son must be wonderful, on both, or on neither?

5 If the old sea captain had derived his estimate of his latitude from the horoscopes of the crew, would this method have been more or less reliable than his use of old-fashioned astronomy?

6 Why is my belief that I will not win the lottery (see section 4) statistically reliable? What reliable method is being used in this case?

7 An officer is using a radar device to measure the speeds of cars on a road. The device registers the speed of car A as 100 mph. Which of the following are relevant to the question of whether the use of the device tracks the speed of car A?
   (a) If the previous car had gone 60 mph slower, the radar would have registered 100 mph.
   (b) If the car after car A goes 100 mph, the radar will register 100 mph.
   (c) If car A goes 100 mph, the radar will register 100 mph.
   (d) The car is going 100 mph.
   (e) If car A goes 60 mph, the radar will register 60 mph.
   (f) If the car goes 60 mph, the radar will register 100 mph.
   Which of these are relevant to the question of whether the officer knows that car A is going 100 mph?

8 Does McA's belief that there will be a hurricane track the fact that there will be a hurricane?

9 Which of these are externalist properties of beliefs, and which internalist: falsity, causes, simplicity, being more or less perceptual, being apriori, being necessary?

Thinking Questions

10 Are all desirable features of belief relevant to whether a belief is knowledge? What about unoffensiveness, simplicity, comfortingness?

11 Newspaper quotation: 'The trekkies who follow the voyages of the *USS Enterprise* have known for years that there are planets of all kinds orbiting practically every star in the galaxy.' The author of this clearly does not think that the trekkies' belief is true. But he or she still says that they 'know' it. Does this show that knowledge does not have to be true?

12 How does Lehrer's principle relate to cases like the Princess Diana example in which there are relevant true beliefs which the person does not have?

13 We might try defining knowledge as belief which is true and justified and would remain justified whatever true beliefs were added to the believer's stock of beliefs. What problems would this definition run into?

14 Internalist properties might reduce to externalist ones. For example, a reasonable belief might be one that is produced by a method that most often gives true conclusions. How might an internalist object to this?

15 Externalist properties might reduce to internalist ones. For example, a true belief might be one which a reasonable person who had access to all the evidence would hold. How might an externalist object to this?

Further Reading

Good surveys of work on defining knowledge are Robert K. Shope, *The Analysis of Knowing* (Princeton University Press, 1983), and chapters 2 and 3 of Jonathan Dancy, *An Introduction to Contemporary Epistemology* (Blackwell, 1985). Dancy is probably the simpler of these.

Part 2 of Robert Nozick's *Philosophical Explanations* (Oxford University Press, 1981) is a standard source for reliability and conditional theories of knowledge. Personal favourites are Fred Dretske, 'Conclusive reasons', *Australasian Journal of Philosophy*, 66 (1971), 1–22, and Edward Craig, *Knowledge and the State of Nature* (Oxford University Press, 1990). For an extremely stimulating, unorthodox approach to knowledge, read Peter Unger, *Ignorance: A Case for Scepticism* (Oxford University Press, 1975).

The philosophical text which started the debates reported in this chapter can be found in Part I section 2 of John Cottingham, *Western Philosophy: An Anthology* (Blackwell, 1996), Plato, 'Knowledge versus Opinion'.