

BELIEFS AND THEIR QUALITIES

1. Defending and Attacking Beliefs

Any person has many beliefs. You believe that the world is round, that you have a nose and a heart, that $2 + 2 = 4$, that there are many people in the world and some like you and some do not. These are beliefs about which almost everyone agrees. But people disagree too. Some people believe that there is a god and some do not. Some people believe that conventional medicine gives the best way of dealing with all diseases and some do not. Some people believe that there is intelligent life elsewhere in the universe and some do not. When people disagree they throw arguments, evidence, and persuasion at one another. Very often they apply abusive or flattering labels to the beliefs in question. “That’s false,” “That’s irrational,” “You haven’t got any evidence,” or “That’s true,” “I have good reason to believe it,” “I know it.”

We use these labels because there are properties we want our beliefs to have: we want them to be true rather than false, we want to have good rather than bad reasons for believing them. The theory of knowledge is concerned with these properties, with the difference between good and bad beliefs. Its importance in philosophy comes from two sources, one constructive and one destructive. The constructive reason is that philosophers have often tried to find better ways in which we can get our beliefs. For example, they have studied scientific method and tried to see whether we can describe scientific rules which we could follow to give us the greatest chance of avoiding false beliefs. Rationalism and empiricism and Bayesianism, described later in this book, are constructive philosophies of this sort. The destructive reason is that philosophy has often been caught up in the conflict between one set or system of beliefs and another. For example, people with religious faith sometimes try to find philosophical reasons for believing in God, and anti-religious people sometimes try to find philosophical reasons why it is irrational to believe in God. So the theory

of knowledge – or epistemology as it is also called, from the Greek word *epistēmē* meaning “knowledge” – can get involved both in trying to find better ways of acquiring beliefs and in criticizing the beliefs people already have.

This chapter introduces the basic ideas and terminology of the theory of knowledge. It connects the search to improve our beliefs and to referee the conflicts between different systems of beliefs with the fundamental ideas of the subject. The central idea here is the importance of questions about the kinds of beliefs we want to have. The chapter ends with two extreme views, deep skepticism and radical externalism, to show the importance of these questions.

2. Epistemic Ideals

Until recently most philosophers working in the theory of knowledge have not paid much attention to the different ways in which beliefs and the ways we acquire them can be satisfactory or deficient. They have not asked: What qualities do we want our beliefs to have and what qualities do we want them not to have? One reason for this is that philosophers have often thought that the answer was obvious: We want our beliefs to be true and not false. As we will see later (especially in chapter 5) this answer is *not* obvious. But, focusing on the desire for truth, most philosophers until recently have described various ideals for beliefs: ways in which our beliefs and the ways we get them could be perfectly organized. Rationalists described an ideal in which arguments as forceful as those in a mathematical proof could demonstrate the truth of all the beliefs we need to know. Empiricists described an ideal in which evidence from what is seen, heard, or otherwise perceived could give adequate evidence for all our beliefs. A contemporary epistemic ideal, defended by the Bayesian movement in probability theory and the theory of knowledge, aims to describe ways in which we can discover exactly how probable each of our beliefs is, given the evidence we possess.

In the course of this book we will discuss each of these ideals. One important question to ask about each of them is “Are human beings capable of satisfying this ideal: can we have beliefs that are like this?” But another equally important question is “What would the price be for satisfying this ideal: in order to have beliefs like this would we have to lose something else of value?”

A very simple epistemic ideal is that of **coherence**. Coherence is having beliefs that not only make sense individually but which hang together in a coherent pattern. If I believe that all cats are intelligent, believe that my neighbour's cat is a cat, and also believe that my neighbour's cat is stupid, then my beliefs are incoherent. They cannot all be true, and I can start with some of them and give good reasons for disagreeing with others. My beliefs could be incoherent for other reasons too. I might believe many things

which amount to strong evidence for something and yet believe the opposite. This often happens when people deceive themselves. Suppose for example that someone knows that their child gets into fights at school, knows that the teachers are afraid of the child, knows that many other children are not permitted to play with the child, and still deceives herself into thinking that her child is a well-behaved little angel. Such a person's beliefs will not be coherent.

Why should we want our beliefs to be coherent? One reason is that incoherent beliefs tend to include many false beliefs. Another is that incoherent beliefs are hard to defend against people who challenge or attack them. So coherence is an ideal that we could set ourselves. We could try to make our beliefs as coherent as possible. This does not mean that any person's beliefs could ever be totally coherent. Every human being will probably always be subject to bad reasoning and self-deception. That's the way we are. But it is an ideal we can try to achieve. It is also an ideal that someone might decide not to aim for, probably because it was thought to conflict with some other ideal, such as the ideal of having interesting new ideas. So even for this very simple epistemic ideal there are questions to answer about what the price of aiming at it is, and how near to achieving it human beings can actually come.

Differences of epistemic ideal lie beneath many disagreements in everyday life. For example, some people have a lot of faith in various kinds of alternative medicine, while others are very skeptical of such claims. A popular book on herbalism may say things like "A man riddled with cancer was told he had four weeks to live. He heard about herb X and decided to drink a quart of X tea daily. Three years later he was still alive." The story may give some people faith in the powers of herb X, while others may refuse to be impressed. The unimpressed people will not only fail to believe that X has curative powers. They will also think that something is wrong with the way that people who are influenced by the story are forming their beliefs. They will think that these people are too ready to believe something on inadequate evidence. They will be loyal to an epistemic ideal according to which beliefs should be related in defensible ways to evidence. A very strong form of this ideal would say that you should only have beliefs for which there is strong evidence like the evidence that supports scientific theories. People who have more faith in alternative medicine are very rarely loyal to any such ideal. They are more likely to be loyal to an ideal according to which a person's beliefs should be part of a pattern that gives meaning to the person's life and fits their own moral and spiritual ideas. These ideals are so different that it is not surprising that people who are loyal to them often find each other's beliefs incredible. The important point for now is that many differences of opinion are the result of different ideals of what a person's beliefs should be like. One task of the theory of knowledge is to provide ways in which these differences of epistemic ideal can be discussed and even resolved.

3. The Basic Concepts

In discussing our beliefs, philosophers use a number of central concepts. It is hard to define any of them in terms that all philosophers would accept, for the definitions are usually tied up with theories about knowledge, just the theories that we have to compare in this book. But if you consider the examples below you will almost certainly conclude that they are concepts that you are already familiar with. Everyone uses them in everyday life, in a general and imprecise way, even if they do not use the words for them that philosophers do. (In what follows I will often write a word in italics to indicate that it is a word worth noticing or in boldface if it is included in the list of definitions of epistemological terms at the end of the book.)

Rational/irrational

George has a date with Shoshana, who is blonde. She decides at the last minute not to go out with George that evening but to stay home and study instead. George is furious, and decides that all blonde women are evil. He will never trust a blonde again. This is especially strange since his mother and sister, who have always treated him with the greatest kindness, are blonde. But from that day on, however friendly, considerate, or helpful a blonde woman is, George always interprets her behavior as evil.

The way that George comes to have the belief that all blonde women are evil is **irrational**. It consists not in careful thinking but in a sudden angry impulse that continues to grip him. Speaking loosely, we may say that the belief itself is irrational. We might say that it was a crazy or a stupid belief, though George may be neither crazy nor stupid. Many philosophers think that a lot of people's beliefs, including beliefs handed down through the generations, are irrational. In some important ways they are like George's belief about blondes. Superstitious beliefs, like for example the belief that it is bad luck to have a black cat cross your path, are good candidates for being irrational. That is, it is irrational to believe that something will cause you harm just because not very well-informed people sometimes say it will, although they can produce no good evidence for this. Some philosophers argue that all religious beliefs are irrational, and some philosophers think that it is irrational to believe in an objective difference between right and wrong. Many other philosophers, needless to say, disagree.

Contrast George with Sonya. Sonya has a cruel father, and her brother is a drug dealer. Neither shows any affection or consideration for her. Indeed, with the exception of two of her teachers in school, all the men who have played any role in her life have been baddies. Yet when asked her attitude to men she says "There are a lot of bad ones. But I've met a few decent ones, so I have some slight hope for them." This does not seem irrational. The way

Sonya gets her belief is **rational**, to the extent that it does not go beyond what the evidence available to her suggests, and leaves open possibilities that are not closed off by that evidence. One conclusion to draw is that very often a rational belief, one that is acquired in a rational way, will have to be more subtly expressed than an irrational one in response to the same evidence. The rational belief will less often say “all” and “never.”

True/false

Suppose that there is some mechanism in the universe which ensures that whenever a black cat crosses a person's path something bad happens to that person in the near future. No human being knows of the existence of this mechanism, which works by physical principles which humans will never understand. Then the belief that it is bad luck to have a black cat cross your path is **true**. (So the belief that the probability of bad events is unchanged by a black cat's crossing your path is **false**.) But still there may not be any good evidence for it: the combinations of cat-crossings and bad occurrences are too subtle for us ever to notice them. Then the belief that black cats are bad luck is true even though there is no good evidence for it; true even if the belief that it is true is irrational. So irrational beliefs can be true. That may be a surprising conclusion, but it is clearly right. A less surprising conclusion is that rational beliefs can be false. There can be very strong reasons for believing something although it is false. (Your best friend who has never lied to you before can tell you that your house was not struck by lightning last night, but in fact . . .)

Evidence

Scientists do experiments to find **evidence** for and against scientific theories. Detectives hunt for evidence showing who committed crimes. Evidence can take many different forms. The behavior of animals in a learning task, the pattern of light in the viewpiece of a telescope or microscope, a letter confessing to an action; all these and many others could count as evidence in suitable circumstances. Very often when evidence is produced it is in order to convince someone to change their mind, from belief to disbelief, from disbelief to belief, or from neutrality to either. The evidence then has to be believable by the people who are to be convinced, and in addition has to be such that when they think about it they will, if they are rational, find some tendency to change their beliefs. So evidence produced by the defense in a court case might be testimony that even a juror who was inclined to convict would take seriously. Similarly, evidence for a scientific theory might be the result of an experiment that even someone who believed a rival

theory would have to admit did definitely occur and did definitely give the result that it did.

Reasoning and argument

When evidence supports a belief it makes people think that the belief might be true. Because of the evidence they perform some reasoning that tends towards the belief. There are many kinds of reasoning. Sometimes to persuade someone you do not produce any evidence at all but say "Suppose for the sake of argument that . . ." and then go on to draw conclusions. A defence lawyer says "suppose that someone besides my client was hiding in the house that night" and then shows how that mystery person could have committed the crime and planted her client's stolen wallet near the scene. The jury thinks about this, and is led through steps of *reasoning* by the lawyer's *argument*. Then they may conclude "someone else just might have done it" or "if someone else could have got into the house then someone else could have done it." Sometimes reasoning can show that a belief is true without using any evidence at all. For example suppose a student says to a librarian "This book was due on February 1st and it is now March 1st, and the fine is \$1 per day, so I owe \$29" and the librarian replies "But this is a leap year, so you owe \$30." The librarian's reasoning shows that the student's belief is wrong, without producing any evidence they didn't both already have.

Justified/unjustified

Consider Toshiro who lives in Tokyo and knows nothing about North American animals. He has never seen a moose or a chipmunk and does not know what sorts of animals they are. His family is taking a trip to Seattle and on the plane they give him a book, in English, with the title *Animals of North America*. In fact the book is a joke and most of the information in it is wrong. In particular, the photo and description of a moose are of a chipmunk and the photo and description of a chipmunk are of a moose. In Seattle he goes to a zoo and sees a moose. He believes that he is seeing the animal called "chipmunk." That is clearly not a silly thing for him to believe, given what he has read and what he sees. On the basis of that information his belief that he is seeing what is called a chipmunk is **justified**. His sister Sumiko, who has much more information about North American animals, looks at the big ruminant with the wide horns and at once thinks "That's a moose," and her belief is also justified. A belief is justified on the basis of information which makes believing it a better strategy for getting to the

truth than not believing it. Often people draw conclusions that are not justified by the information they have. For example if Sumiko thinks that since moose have horns and since “mouse” sounds like “moose” then mice have horns, then her belief is unjustified on the basis of her information, unless she also has some reason to believe that animals whose names in English sound alike are similar.

Many theories of knowledge are theories of when a person’s beliefs are justified. So they give information about when people are rational in acquiring beliefs. A person acquires a belief rationally when the acquisition is based on or guided by information that justifies the belief. That is the simplest kind of case; there are also more complicated situations in which for example a person acquires a belief irrationally and then later finds evidence which supports it. (A person dreams that they will win a lottery and so goes out next day and buys an expensive car; then that evening their ticket is announced as the winner.) In complicated cases it is simplest to think of rationality as a property of people and their thinking, and justification as a property of beliefs.

Knowledge/ignorance

Toshiro was *ignorant* of the names of North American animals. We are all ignorant of many things: there are many questions we do not know the answers to. Probably no human knows whether there is life on other planets. Probably no human knows how to achieve universal peace. Probably no human knows whether there are infinitely many twin prime numbers (like 3 and 5, 11 and 13, 1,001 and 1,003). There are many people who have beliefs about all these things, and some of their opinions are rational, and some are justified. But it does not follow that any of these opinions count as *knowledge*. To know that there is life on other planets a person would have to have a powerful theory of how life develops or have direct evidence produced by such life. To know how to achieve universal peace someone would have to have a recipe for producing peace and a very convincing reason that showed how and why it would work. To know that there are infinitely many twin primes a person would have to have a correct mathematical proof of this fact. In short and very roughly, to know something your mind has to be linked to the fact, and that link has to be a top quality reliable one. This makes knowledge seem very special and very rare. Yet we talk as if we know many things. Just about every person knows the names of his or her friends, and knows that they have a nose on their face. Most people know that the earth revolves around the sun and that 12 multiplied by 13 is 156. One sign of this is that most people can be relied on to give reliable information about these things. Seen this way, it is not

surprising that there are many controversies about knowledge: about what knowledge is and how much knowledge we have. For it is at the same time something that seems very hard to achieve and something that we think we have quite a lot of. (This suggests that disputes between those who want to use the theory of knowledge to criticize the beliefs of others and those who want to use it to understand or improve the beliefs they already have, may go very deep.)

None of these explanations was a real definition. There were too many vague and unexplained terms in them, for example the idea of a “top quality reliable” link between a mind and a fact. Later in this book more precise definitions of some of these terms will be considered. (See chapter 6 for more about knowledge, and see question 16 at the end of this chapter for more about the difference between justified and rational beliefs.) But the explanations probably will have reminded you of enough about these concepts that you can understand them. The important point to grasp now is that all of these words can be used to describe desirable and undesirable, good and bad, features of our beliefs. It is not at all obvious that there is only one kind of desirable feature of beliefs, so these good and bad aspects can cut across one another in complicated patterns. In particular note the following three complications.

A good result can be obtained by a bad method. For example, a true belief can be got by irrational reasoning. You see a spider and because you are scared by its hairy legs you think it must be poisonous. That is bad reasoning, but it may be that the spider is poisonous all the same. There are many examples of this in the history of science. For example, William Harvey in the seventeenth century came up with the theory that the blood in the body circulates, leaving the heart by the arteries and returning by the veins. He got to this conclusion by thinking: The heart is like the sun and the blood is like the earth, so since the earth revolves around the sun the blood must revolve around the heart. This is not very convincing reasoning, to put it mildly, but it led him to a true conclusion. (He later did experiments and found better evidence for his idea.)

A bad result can be obtained by a good method. For example, Toshiro in the example above was not reasoning badly when he thought that the large animal in front of him was called a “chipmunk.” Or consider a scientist who tests a million samples of a drug on a dozen species of animals and finds no harmful side-effects. In the absence of contrary evidence she is justified in concluding that the drug is harmless. But it may turn out that in some species under some conditions the drug is fatal. Her justified belief was false. (This example is related to questions about induction, discussed in chapter 4.)

Opposing beliefs can both be justified. Toshiro and Sumiko both had justified beliefs, in terms of the different evidence available to each of them. People thousands of years ago were not stupid when they thought that the earth was flat and the sun revolved around it, just as we are not stupid in believing that it is spherical and revolves around the sun. Relative to the evidence available to ancient people their belief was a sensible one. Note, though, that opposing beliefs cannot both count as knowledge. If the earth is really flat then we are wrong in thinking that we know that it is spherical.

These three complications are similar to complications that arise whenever we are applying several different kinds of good and bad qualities. In particular, they are similar to complications that arise in ethics, when we are trying to understand the good and bad qualities of human actions. There too we find that a good result can be obtained by a bad method, as when someone attacks a rival out of petty jealousy and thereby accidentally prevents the rival from committing a murder. And we find that a bad result can be obtained by a good method, as when a person saves the life of a drowning swimmer who then commits several murders. Opposite actions can also be justified, too, as when two people are in a burning building and one rushes out so that she can survive to take care of her children and the other rushes further in to save some children who are trapped inside.

The analogy with ethics is far-reaching. We apply many of the same labels when evaluating actions and reasoning: careful, sloppy, reliable, clumsy, (in)accurate, (ir)responsible, effective, pointless, and so on. We have in everyday life standards and criteria for both the ways we act and the ways we form our beliefs. And in both cases there is a tension between those who want philosophical reflection to show how often we fail to meet appropriate standards, and those who want it to help us understand our success in meeting the standards we normally set.

4. The Basic Questions of the Theory of Knowledge

Philosophers contributing to the theory of knowledge have been trying to judge how good our beliefs are and how good they could be. They want to evaluate the beliefs we actually have and to suggest ways in which we could get better beliefs. These two aims are obviously connected: if a philosopher thinks that our beliefs are generally rational and true then he or she is less likely to suggest radical changes in the ways we get new beliefs, and if a philosopher thinks that our beliefs are a mass of confusion and falsity then he or she is likely to suggest either very different ways of getting new beliefs, or despair. So there are three central questions that the theory of knowledge tries to answer.

What qualities should our beliefs have?

What qualities do our actual present beliefs have?

What qualities could our beliefs have?

The answers to the first question might seem obvious. We want our beliefs to be true, rational, and based on evidence. But suppose that a philosopher persuaded you that we are not capable of getting very many true beliefs; then you might settle for aiming at rationality rather than truth. Or suppose that a philosopher persuaded you that using reasoning to base beliefs on evidence will result in far fewer true beliefs than some other method, for example relying on the authority of some tradition. Then you might settle for aiming at true beliefs rather than beliefs based on evidence.

Nearly all philosophers want us to aim at both truth and rationality. They do differ, though, on the relative importance of these and other good qualities of beliefs. (For more on this see chapter 5.) Philosophers divide, though, into what might be called conservative and radical camps on the question of how much better our beliefs could be than they are. In the early days of the scientific revolution philosophers were very optimistic about the possibilities for human knowledge. They proposed ways of basing beliefs on reason and evidence which they hoped would give beliefs which were both more rational and more often true. (Some of these ways are discussed in chapters 2, 3, and 4.) Many of these philosophers were looking for ways of using reason and evidence which would make it unnecessary to rely on faith or on tradition and authority. Earlier philosophers, and many philosophers writing in more recent times, are less prone to suggest radically new ways of obtaining beliefs. The most radical recent suggestions have tended to come from Bayesian epistemology, discussed in chapter 10.

5. Two Extreme Views

To see how different answers to the three questions above can be combined, consider two very extreme views, *deep skepticism* and *radical externalism*. (**Skepticism** – sometimes spelled “scepticism” – and **externalism** are standard terms in the theory of knowledge. I add the labels “deep” and “radical” to show that I am describing particular forms of these positions.)

Deep skepticism

This answers the question “What qualities should our beliefs have?” with “We should be able to give conclusive reasons which show that they are true.” It answers the question “What qualities do our actual present beliefs

have?” with “We cannot give conclusive reasons why they are true.” And it answers the question “What qualities could our beliefs have?” with “Human beings are not capable of having beliefs which they can know to be true.” Deep skepticism thus gives a very strongly pessimistic description of the possibilities for human knowledge; it suggests that we have no, or very little, knowledge. (It is thus a global skepticism, since it covers all of our beliefs. Some milder forms of skepticism might be less pessimistic, in particular they might be local rather than global, applying to just some of our beliefs. For example one might be a religious skeptic, claiming just that we cannot know which of our beliefs about God and immortality are true.)

What reasons might there be for deep skepticism? Here are three arguments for deep skepticism.

Mistakes in reasoning. Most of what we believe is based on evidence. For example we believe that some medicines cure some diseases by reasoning from evidence about particular people with those diseases. But it is easy to make mistakes in reasoning. One small mistake can make a whole chain of reasoning go wrong. One safeguard is to check our reasoning. But this is not really a safeguard as the checking is itself reasoning and can just as easily go wrong. Suppose that there were some profound flaw in the way human beings think. This would spoil all our reasoning, but since it would also spoil the reasoning we use to check our reasoning we would never know that our thinking had gone wrong.

Perceptual illusions. Very often things are not the way they seem. Often when we rely on seeing or hearing we form false beliefs. One reason for this is that there are many illusions: mirages, tricks of perspective and of light, ways in which the human visual system does not work perfectly. (These are discussed at greater length in chapter 2.) And then there are dreams and hallucinations, in which people often think they are perceiving things which do not even exist. We are aware of illusions and hallucinations because they contradict the rest of our experience. But that only means that small errors can be caught. Big errors, in which a large range of our perception is illusory, are much less likely to be caught. (In a dream you usually do not know you are dreaming.) So there could be illusions that permeate right through our perception, which we will never recognize as illusions.

Our poor track record. Humans have often been wrong in the past. We once thought that the earth was flat and now we think it is spherical. The ancient Greeks thought that matter was composed of atoms, and then science thought that it was continuously divisible like a fluid until in the late nineteenth century scientists began to believe in atoms again. Newton, in the eighteenth century, thought that light was composed of particles, but later scientists decided that it was composed of waves instead, until Einstein

convinced the scientific world that there are light particles, photons. Any scientific theory will eventually be shown to be false. And the theory that succeeds it will also eventually be shown to be false. Yet scientific beliefs are the ones we have most reason to think true; if these beliefs are false then there is little hope for all our other beliefs. So we have reason to conclude that almost all our beliefs are false.

I shall not evaluate these three arguments. You should, though, consider seriously how convinced by them you are. (See questions 8 and 12 at the end of this chapter.) Instead, I shall describe a very different position, radical externalism.

Radical externalism

This answers the question “What qualities should our beliefs have?” with “Truth is the most important quality of beliefs; other features such as rationality are simply ways of making it more likely that a belief is true.” It answers the question “What qualities do our actual present beliefs have?” with “Many of them, especially beliefs about the world around us, are true.” And it answers the question “What qualities could our beliefs have?” with “There are many ways in which human beings, individually and in cooperation with one another, can be reliable sources of true information about their environment.”

Here are three arguments for radical externalism.

Evolution. Human beings evolved in an environment not too different from the one we now inhabit. Our ancestors lived on the surface of this planet, dealing with objects of roughly the sizes and shapes that we now encounter. Our human and non-human ancestors were sensitive to the same frequencies of light and sound as us. If our senses and our ability to use the information we get from them were not generally accurate we would have died out thousands of years ago. But we are still here, a testimony to our ability to form true beliefs about our environment.

Intuition. When people defend their beliefs with conscious chains of reasoning which they express with words they make many mistakes. But in the course of everyday life we learn many routines of thinking which work well, especially if we do not stop to think about them. For example, people are pretty good at simple arithmetic, as long as it stays simple. (But if you ask a person to state the reasons *why* they believe that $340 - 89 = 251$ they are likely to give a confused and erroneous answer.) Moreover we have skills, such as the ability to find our way from one place to another or the ability to understand one another's moods and facial expressions, which we do not

understand very well but which get us satisfactory results. We usually do not get lost on the way home and we usually understand when another person is angry with us. These everyday routines and skills are sources of beliefs, such as the belief that $340 - 89 = 251$, that to get home you must turn left at Elm Street, or that the person you are speaking to is on the verge of violent anger. Since our everyday lives depend on these beliefs and since our everyday lives work reasonably well, we can conclude that the routines and skills we use are fairly reliable sources of true beliefs.

Co-operation. People acting in groups can easily do many things that would be impossible for individual people. This applies not only to hunting and building houses but also to knowing. A member of a society can know information learned generations before and passed down from one person to another. Language is essential to this, and so is the disposition to trust what other people say unless there is some reason to disbelieve them. In modern time a lot of our beliefs depend on very complex networks of co-operation. For example, you believe that television sets pick up radio waves coming through the air, but you probably cannot give a good description of how they do this or even what radio waves really are. And you believe that antibiotics such as penicillin are effective against many bacterial diseases, but if you are like most people you do not have any good idea of how antibiotics are made or how they work against bacteria. But this information is available to you, via your links with other people: you can read books, ask for advice, and consult with experts. So the whole community in a way has a more complete knowledge of many things that individual people do not. Functioning as part of a community, trusting others so that they can put their different bits of information together, an individual person can have a reliable access to many true beliefs. Individual people can rarely have this knowledge using just their individual resources, and they usually cannot give a fully convincing explanation or justification of it. But that does not deny the fact that the beliefs are trustworthy and true.

Deep skepticism and radical externalism are very different views. But notice that they are not simple opposites. For they start from different answers to the first question "What qualities should our beliefs have?" Deep skepticism assumes that the most important quality of our beliefs is that we can be sure that they are true. It then tries to show that we cannot be sure that many of our beliefs are true. On the other hand, radical externalism assumes that the most important quality of our beliefs is just that they are true, whether or not we can be sure of it. (And whether or not the factors that make us capable of having true beliefs are ones we are aware of.) It then tries to show that many of our beliefs are true, even if we cannot be sure exactly which ones. Both deep skepticism and radical externalism may be wrong. But until we have decided which are the most important qualities

for our beliefs to have, we should not feel that if one were to be right the other has to be wrong. Given the assumptions each one makes about the aims of belief, they could both be right.

Issues connected with deep skepticism and radical externalism will get more discussion in later chapters. (See especially chapters 5 and 7.) For now the important point is how answers to the three questions shape theories about knowledge. Do we want our ways of acquiring beliefs to produce true beliefs, rational beliefs, beliefs that are true and rational, beliefs that we can be sure are true and justified, or what? How much more true, more based on evidence, or more rational could the beliefs of human beings be? Which of these questions you find important will depend on whether your aim is to criticize beliefs, to find good ways of acquiring them, or to understand the ways we normally acquire them. But whatever the aim, before making elaborate theories it is a wise idea to ask what questions they are supposed to answer.

Reading Questions

(For the distinction between reading questions and thinking questions see the foreword for students.)

- 1 Section 1 mentioned constructive and destructive functions of the theory of knowledge. It also gave examples of searching for better ways of acquiring beliefs and of criticizing beliefs: which of these was the constructive function and which was the destructive function?
- 2 Section 2 gives examples of epistemic ideals. It then suggests that we should consider whether they can be achieved by human beings. How might the rationalist ideal not be achieved by human beings?
- 3 Section 3 claims that rational beliefs will tend not to be wide generalizations expressed with words like “all” and “never.” Could we express any rational beliefs using these words?
- 4 Section 3 had a remark that we should not really say that a belief is rational but instead say that it is rational for a person to acquire a particular belief at a particular time in a particular way. Give an example of a belief that is rational for one person to acquire in one situation but not rational for another person to acquire in another situation.
- 5 Section 3 said that there can be very strong reasons for believing something although it is false. Give an example of this.
- 6 Consider the following four cases. (They are all Canadian cases, for no special reason.)
 - (a) Alberta is enjoying a rare warm break in a cold winter. It is January and she is in Calgary. She is about to travel to Edmonton, 250 km to the north, for a week. She thinks “I’ll leave my winter coat behind; it’s so nice today it’s sure to stay warm.”
 - (b) Victoria is waiting for the ferry from Vancouver to Nanaimo. She is wearing an expensive watch which has never gone wrong and which is reset every

noon by a time signal from the British Columbia observatory. She looks at it and sees that it says noon. She thinks “It is noon, so I had better hurry to get the ferry.”

- (c) Arthur is waiting for the plane to take him from Thunder Bay to Winnipeg. He has a watch which works one day in three. He looks at it and thinks “It is noon, so I have plenty of time before the plane leaves.”
- (d) William is waiting for the plane to take him from Winnipeg to Thunder Bay. The airport clock says 11:42 a.m. and he thinks “The plane boards at 1 p.m. so I can spend fifteen minutes finishing my coffee and still have more than an hour before I have to be at the gate.”

In which of these is the person’s belief rational? In which is it justified? In which is it knowledge? Classify them into Yes, No, and Perhaps.

- 7 Section 4 said that if you thought that we were not capable of getting very many true beliefs; then you might want to have rationality rather than truth as your aim. It also said that you might want to have truth rather than evidence as your aim if you thought that basing beliefs on evidence would result in few true beliefs. Why?
- 8 In section 5 the second argument for deep skepticism said “We are aware of illusions and hallucinations because they contradict the rest of our experience. But that only means that small errors can be caught. Big errors, in which a large range of our perception is illusory, are much less likely to be caught.” Why might someone think this?
- 9 The epistemic ideals described in section 2 were ones that scientifically minded people could accept: coherence, rationality, truth. But might there not be a price for having beliefs that are rational and aim at being true? Describe some advantages that might come from giving up on rationality and truth.
- 10 The first argument for radical externalism said “If our senses and our ability to use the information we get from them were not generally accurate we would have died out thousands of years ago.” Why might someone think this?

Thinking Questions

- 11 Consider canny Cassie. She is a great guesser. She has hunches about which teams are going to win basketball or football games; she has intuitions sometimes that a news bulletin or a newspaper headline is false; she often has a conviction about unannounced schedule changes of television programs; and just occasionally she has a terrible foreboding of an impending disaster. We all experience these things; but Cassie is different in that her hunches, intuitions, convictions, and forebodings are nearly always right. When she has a belief about the future it is nearly always true. When she is asked about this she says “I don’t know where my beliefs come from and I know no reason why they should be true. In fact, I don’t expect that my beliefs about the future are any more reliable than other people’s.” But she is wrong: her guess-based beliefs are more reliable than other people’s.

Suppose you are about to travel to New York by plane. Cassie says to you “Don’t take that flight. I don’t know why I feel this, but I really don’t want you

to take that flight.” Is Cassie’s belief rational? Suppose you take her seriously and believe that you shouldn’t take the flight. Is your belief rational? Suppose that while taking Cassie seriously you continue to believe that no person can know the future. Is that belief then irrational?

- 12 Here is part of the first argument for deep skepticism. “Suppose there were some profound flaw in the way human beings think. This would spoil all our reasoning, but since it would also spoil the reasoning we use to check our reasoning we would never know that our thinking had gone wrong.” Is it so clear that we could never tell if there was a flaw that affected all our reasoning? Suppose for example that people always estimate probabilities wrongly. Would we never notice this fact about ourselves? Might there be other kinds of problems about our reasoning which would be forever hidden from us?
- 13 Here is part of the first argument for radical externalism. “If our senses and our ability to use the information we get from them were not generally accurate we would have died out thousands of years ago.” Is this true? Can you think of kinds of inaccurate perception which would not interfere with a species’ survival? Might there be kinds of inaccurate perception which would actually increase the survival chances of a species?
- 14 Consider the ultimate virtual reality machine. It plugs into both your sensory and motor nerves and connects them to a hyper-powerful computer, which stimulates all your senses with exactly the input they would get from a real environment, and changes these stimulations exactly the way they would change if you were really moving your body while interacting with this real environment. Of course the environment the computer simulates is completely different from your actual environment. Suppose you had been in such a machine from birth. Would you have any reason to believe you were not living a real life in the environment you seem to find around you? Would your beliefs (for example that you are reading a philosophy book at this moment) be rational? How does this relate to deep skepticism?
- 15 How does the ultimate virtual reality scenario relate to radical externalism?
- 16 In section 3 rational beliefs were described in terms of the way a person acquires the belief at a particular time. Justified beliefs, on the other hand, were described in terms of the relation between the belief and the information on which it is based. Very often rational beliefs will be justified, and rationality and justification will coincide. But not always. Consider two examples. Neither of them makes an open-and-shut case, but discuss their implications for the relation between rationality and justification. (It is worth considering variations on the cases, or different ways of spelling out the details in them.)
 - (a) Genevieve wants to know whether the coin in her hand is fair. (That is, whether there are equal chances of heads and of tails if it is tossed.) She tosses it four times and it comes down heads three times. A friend who has a Ph.D. in mathematics tells her that if the coin is fair there is a 2/16 chance that it will come down heads at least three times in four throws. She thinks that since this probability is pretty small the coin is probably biased. In fact her friend is wrong, and there is a 5/15 chance that a fair coin will come down heads at least three times in four throws. Is Genevieve’s belief rational? Is it justified?

- (b) Genevieve gives a dollar to an old tramp on the street. He gives her a long complicated rambling explanation of why the probability of two people giving him a dollar on a Tuesday is 0.00347. He looks like such a sweet old man that she believes him. In fact he is a once-brilliant mathematician who has had a nervous breakdown and his explanation is completely correct, though Genevieve did not understand it. Is Genevieve's belief rational? Is it justified?
- 17 Assume that an epistemological theory we may call "traditionalism" is true. According to traditionalism it is rational to hold a belief which many others in your culture believe and have believed, even if there is no strong evidence that it is true, as long as there is no strong evidence that it is false. (Take "strong evidence" to be evidence that would make it rational to hold the belief independently of the traditions of your culture.) Then, as in question 16, we have a way in which it might be rational to hold a belief even though you had no justification for believing it. It all depends on what you mean by "justified." Suppose that a person lives in a culture in which some religious beliefs are widely held. Suppose moreover that there is no direct evidence for the existence of God – no evidence that should convince a determined atheist – and there are no logical arguments proving the existence of God. Contrast the following two definitions of "justified."
- (a) A person is justified in holding a belief B if she has evidence that would force anyone who had it to accept B or if she has a logical argument that leads from her present beliefs to B.
- (b) A person is justified in holding a belief B if given her other beliefs it is rational to add B to them.
- Assume that traditionalism is right. Describe in detail a person's beliefs and the person's situation (evidence, beliefs of other people, etc.) which would make that person's belief in God rational but not justified, according to (a) but not according to (b).
- 18 At the very end of the chapter we find the remark that different questions we ask about beliefs are related to different purposes in studying belief. Which of the questions listed in the last paragraph of the chapter are most relevant to the aim of criticizing beliefs? Which to improving belief-acquisition? Which to understanding it?
- 19 Gunther is taking part in a psychological experiment. The experimenters tell Gunther that they are giving him a powerful drug which will make all people look like his mother. Gunther believes them, although the "drug" they give him is actually pineapple juice. The point of the experiment is to test whether he will think that women look like his mother just because he expects them to. Gunther's mother is a telephone repair person and has been called to fix a telephone in the psychology department. By mistake she walks into the laboratory where Gunther is. "Hello mom" says Gunther, for a moment thinking it is his mother, before doubt sets in. Is Gunther's momentary belief that the woman in front of him is his mother justified? Does he know that it is his mother in front of him? (You might have noticed that I said nothing in this chapter about the relation between justification and knowledge. Chapter 6 takes up the question, and chapter 7 focuses on issues raised by this example.)

Further Reading

Accessible discussions of belief, reason, and the aims of the theory of knowledge are found in chapters one and two of Martin Hollis, *Invitation to Philosophy* (Blackwell, 1985), and in chapter one of W. V. Quine and Joseph Ullian *The Web of Belief* (Random House, 1978). Chapters 1 and 2 of Adam Morton, *Philosophy in Practice* (Blackwell, 1996) are also relevant. Chapter 1 of Jonathan Dancy, *An Introduction to Contemporary Epistemology* (Blackwell, 1985) begins with skepticism as a basis for covering some of the same ground as this chapter. Dancy is not an introductory book, though. Issues about how the aims of epistemology have changed during its history are discussed in Mary Tiles and Jim Tiles, *An Introduction to Historical Epistemology* (Blackwell, 1993). Fundamental distinctions between rational and justified beliefs are made and challenged in chapter one of Alvin Plantinga, *Warrant: The Current Debate* (Oxford University Press, 1993), which also discusses the resemblances between standards in morals and in epistemology. Relevant recent points of view are discussed in Michael Williams' "Skepticism," and Keith Lehrer's "Rationality" in John Greco and Ernest Sosa (eds), *The Blackwell Guide to Epistemology* (Blackwell, 1999).

Selections from classic philosophical works relevant to this chapter can be found in John Cottingham, *Western Philosophy: An anthology* (Blackwell, 1996). See part I section 7, David Hume, "Scepticism versus Human Nature," and part I section 10, G. E. Moore, "Against Scepticism." (These are Cottingham's titles, not Hume's and Moore's.)

Electronic resources: *Routledge Encyclopedia of Philosophy*: Epistemology; Epistemology, history of; Rational beliefs; Normative epistemology; Reasons for belief; Skepticism; Certainty; Doubt. *The Stanford Encyclopedia of Philosophy*: reasons; justification vs. explanation; skepticism.