

One of the remarkable conclusions of contemporary epistemology is that the rational thought responsible for our great intellectual achievements is not different in kind from the rational thought involved in routine epistemic procedures, like determining the color of an object seen in broad daylight, remembering your mother's name, discovering that most objects fall to the ground when unsupported, or summing 12 and 25. If we can understand how rational thought enables us to solve these routine epistemic problems, we can understand the discovery of DNA as the result of stringing together a large number of routine problems. What is extraordinary about human thought is already present in our ability to solve routine epistemic problems. Thus we begin by focusing on the mundane, in hopes that it will lead us to the sublime.

2. Skeptical Problems

Rather than ask, "How is it possible to discover DNA, or find a cure for cancer?", the epistemologist has traditionally begun by asking, "How is knowledge possible at all?" The philosophically inexperienced reader might find this puzzling, thinking to herself, "I know how to tell what color something is and how to remember my mother's name. I don't care about that. I want to know how to find a cure for cancer." Indeed, we all know how to perform simple epistemic tasks, but there is an important difference between knowing how to do it and knowing how it is done. We all know how to pick up a cup of coffee without spilling it, but imagine trying to give a precise description of how to do this sufficient to enable an engineer to build an industrial robot to accomplish the same task. In fact, engineers spent years trying to solve this very problem. Similarly, although we know how to perform simple epistemic tasks, it is extremely difficult to explain how we do and why what we do yields knowledge. Historically, philosophers have often motivated the study of simple epistemic tasks with the help of *skeptical arguments*. These are initially compelling arguments that seem to show that even simple epistemic tasks are impossible. Consider the following tale:

It all began that cold Wednesday night. I was sitting alone in my office watching the rain come down on the deserted streets outside, when the phone rang. It was Harry's wife, and she sounded terrified. They had been having a late supper alone in their apartment when suddenly the front door came crashing in and six hooded men burst into the room. The men were armed and they made Harry and Anne lay face down on the floor while they went through Harry's pockets. When they found his driver's license one of them carefully scrutinized Harry's face, comparing it with the official photograph and then muttered, "It's him all right." The leader of the intruders produced a hypodermic needle and injected Harry with something that made him lose consciousness almost immediately. For some reason they only tied and gagged Anne. Two of the

men left the room and returned with a stretcher and white coats. They put Harry on the stretcher, donned the white coats, and trundled him out of the apartment, leaving Anne lying on the floor. She managed to squirm to the window in time to see them put Harry in an ambulance and drive away.

By the time she called me, Anne was coming apart at the seams. It had taken her several hours to get out of her bonds, and then she called the police. To her consternation, instead of uniformed officers, two plain clothed officials arrived and, without even looking over the scene, they proceeded to tell her that there was nothing they could do and if she knew what was good for her she would keep her mouth shut. If she raised a fuss they would put out the word that she was a psycho and she would never see her husband again.

Not knowing what else to do, Anne called me. She had had the presence of mind to note down the number of the ambulance, and I had no great difficulty tracing it to a private clinic at the outskirts of town. When I arrived at the clinic I was surprised to find it locked up like a fortress. There were guards at the gate and it was surrounded by a massive wall. My commando training stood me in good stead as I negotiated the 20 foot wall, avoided the barbed wire, and silenced the guard dogs on the other side. The ground floor windows were all barred, but I managed to wriggle up a drainpipe and get in through a second-story window that someone had left ajar. I found myself in a laboratory. Hearing muffled sounds next door I peeked through the keyhole and saw what appeared to be a complete operating room and a surgical team laboring over Harry. He was covered with a sheet from the neck down and they seemed to be connecting tubes and wires to him. I stifled a gasp when I realized that they had removed the top of Harry's skull. To my horror one of the surgeons reached into the open top of Harry's head and eased his brain out, placing it in a stainless steel bowl. The tubes and wires I had noted earlier were connected to the now disembodied brain. The surgeons carried the bloody mass carefully to some kind of tank and lowered it in. My first thought was that I had stumbled on a covey of futuristic Satanists who got their kicks from vivisection. My second thought was that Harry was an insurance agent. Maybe this was their way of getting even for the increases in their malpractice insurance rates. If they did this every Wednesday night, their rates were no higher than they should be!

My speculations were interrupted when the lights suddenly came on in my darkened hidey hole and I found myself looking up at the scariest group of medical men I had ever seen. They manhandled me into the next room and strapped me down on an operating table. I thought, "Uh, oh, I'm in for it now!" The doctors huddled at the other end of the room, but I couldn't turn my head far enough to see what they were doing. They were mumbling among themselves, probably deciding my fate. A door opened and I heard a woman's voice. The deferential manner assumed by the medical malpractitioners made it obvious who was boss. I strained to see this mysterious woman but she hovered just out of my view. Then, to my astonishment, she walked up and stood over me and I realized it was my secretary, Margot. I began to wish I had given her

that Christmas bonus after all.

It was Margot, but it was a different Margot than I had ever seen. She was wallowing in the heady wine of authority as she bent over me. "Well Mike, you thought you were so smart, tracking Harry here to the clinic," she said. Even now she had the sexiest voice I have ever heard, but I wasn't really thinking about that. She went on, "It was all a trick just to get you here. You saw what happened to Harry. He's not really dead, you know. These gentlemen are the premier neuroscientists in the world today. They have developed a surgical procedure whereby they remove the brain from the body but keep it alive in a vat of nutrient. The Food and Drug Administration wouldn't approve the procedure, but we'll show them. You see all the wires going to Harry's brain? They connect him up with a powerful computer. The computer monitors the output of his motor cortex and provides input to the sensory cortex in such a way that everything appears perfectly normal to Harry. It produces a fictitious mental life that merges perfectly into his past life so that he is unaware that anything has happened to him. He thinks he is shaving right now and getting ready to go to the office and stick it to another neurosurgeon. But actually, he's just a brain in a vat."

"Once we have our procedure perfected we're going after the head of the Food and Drug Administration, but we needed some experimental subjects first. Harry was easy. In order to really test our computer program we need someone who leads a more interesting and varied life—someone like you!" I was starting to squirm. The surgeons had drawn around me and were looking on with malevolent gleams in their eyes. The biggest brute, a man with a pockmarked face and one beady eye staring out from under his stringy black hair, was fondling a razor sharp scalpel in his still-bloody hands and looking like he could barely restrain his excitement. But Margot gazed down at me and murmured in that incredible voice, "I'll bet you think we're going to operate on you and remove your brain just like we removed Harry's, don't you? But you have nothing to worry about. We're not going to remove your brain. We already did—three months ago!"

With that they let me go. I found my way back to my office in a daze. For some reason, I haven't told anybody about this. I can't make up my mind. I am racked by the suspicion that I am really a brain in a vat and all this I see around me is just a figment of the computer. After all, how could I tell? If the computer program really works, no matter what I do, everything will seem normal. Maybe nothing I see is real. It's driving me crazy. I've considered checking into that clinic voluntarily and asking them to remove my brain just so that I can be sure. Frankly, I don't know if even that would put my worries to rest.

Mike is luckier than most brain-in-a-vat victims. He at least has a clue to his precarious situation—Margot told him he is a brain in a vat. Of course, it could all be contrived. Perhaps he is not a brain in a vat after all. There is no way he can be sure. Meditating about this case, it may occur to you that you might be a brain in a vat, too. If you are, there is no way you could ever find out. Nor, it seems, is there any way

you can be sure you are not a brain in a vat, because everything would seem just the same to you in either case. But if you cannot be sure you are not a brain in a vat, how can you trust the evidence of your senses? You have no way of knowing that they are not figments of a computer. It seems that you cannot *really* know anything about the world around you. It could all be an illusion. You cannot rule out the possibility that you are a brain in a vat, and without being able to rule out that possibility, knowledge of the material world is impossible.

This is a typical example of a skeptical problem. Mike's plight involves fanciful technology and thinking about it is entertaining, but conundrums of this sort have a serious philosophical point. Skeptical problems seem to show that we cannot have the kinds of knowledge we are convinced we have, including the most mundane sorts of knowledge that we take for granted on a daily basis. If skeptical problems challenge our most basic kind of knowledge, then they appear also to easily and completely undermine the sophisticated knowledge that is distinctive of human beings. Such problems have played a central role in epistemology. It is tempting to become caught up in the task of refuting the skeptic, and at one time epistemologists took that to be their principal goal. Descartes was concerned with finding beliefs that he could not reasonably doubt and to which he could appeal in justifying all the rest of his beliefs, and Hume was nonplussed by his inability to answer his own skeptical dilemma about induction. In the *Critique of Pure Reason*, Kant wrote:

It still remains a scandal to philosophy ... that the existence of things outside of us ... must be accepted merely on *faith*, and that, if anyone thinks good to doubt their existence, we are unable to counter his doubts by any satisfactory proof.¹

But contemporary epistemology tends to take a different attitude toward skepticism. If we consider a variety of skepticism that confines itself to some limited class of beliefs, it *might* be possible to answer the skeptic by showing that those beliefs can be securely defended by appeal to other beliefs not among those deemed problematic. But for any very general kind of skepticism, that is impossible in principle. Every argument must proceed from some premises, and if the skeptic calls all relevant premises into doubt at the same time then there is no way to reason with him. The whole enterprise of refuting the skeptic is ill-founded, because he will not allow us anything with which to work.

The proper treatment of skeptical arguments requires looking at them in a different light. We come to philosophy with a large stock of beliefs. Initially, we regard them all as knowledge, but then we discover that they conflict. They cannot all be true because some are inconsistent with others. One instance of this general phenomenon is represented by skeptical arguments. Starting from premises in which we are initially confident,

1. Kant (1958), p. 34. This passage is quoted by G. E. Moore (1959), p. 126.

the skeptical argument leads us to the conclusion that we cannot possibly have certain kinds of knowledge. But we are also initially confident that we do have such knowledge. Thus our original confidently held beliefs form an inconsistent set. We cannot reasonably continue to hold them all.

Upon discovering that our system of beliefs is inconsistent, the initial reaction might be that we should throw them all away and start over again. Descartes pursued this strategy at the beginning of his *Meditations on First Philosophy*. But that will not solve the problem. The skeptic is not just questioning our beliefs. He is also questioning the cognitive processes by which we arrive at our beliefs, and if we start all over again we will still be employing the same cognitive processes. We cannot dispense with *both* the beliefs and the cognitive processes, because then we would have nothing with which to begin again. As Otto Neurath (1932) put it in an often-quoted passage, "We are like sailors who must rebuild their ship upon the open sea."² We must start with the beliefs and cognitive processes we have and repair them "from within" as best we can. The legitimacy of beginning with what we already have was urged by G. E. Moore in a famous passage:

I can prove now, for instance, that two human hands exist. How? By holding up my two hands, and saying, as I make a certain gesture with the right hand, "Here is one hand," and adding, as I make a certain gesture with the left, "and here is another." ... But now I am perfectly well aware that, in spite of all that I have said, many philosophers will still feel that I have not given any satisfactory proof of the point in question. ... If I had proved the propositions which I used as *premisses* in my two proofs, then they would perhaps admit that I had proved the existence of external things. ... They want a proof of what I assert *now* when I hold up my hands and say "Here's one hand and here's another." ... They think that, if I cannot give such extra proofs, then the proofs that I have given are not conclusive proofs at all. ... Such a view, though it has been very common among philosophers, can, I think, be shown to be wrong. ... I can know things which I cannot prove; and among things which I certainly did know, ... were the premisses of my two proofs. I should say, therefore, that those, if any, who are dissatisfied with these proofs merely on the ground that I did not know their premisses, have no good reason for their dissatisfaction. (1959, 144ff)

Even though Moore's complete hostility to skepticism has difficulties of its own, there is something to his point. If we reflect upon our beliefs, we will find that we are more confident of some than of others. It is reasonable to place more reliance on those beliefs in which we have greater confidence, and when beliefs come in conflict we decide which to reject by considering which we are least certain of. If we have to reject something, it is reasonable to reject those beliefs we regard as most

2. "Wie Schiffer sind wir, die ihr Schiff auf offener See umbauen müssen." This passage has been immortalized by Quine (1960) who refers to it repeatedly.

doubtful.³ Now consider how these observations apply to skeptical arguments. An argument begins from premises and draws a conclusion:

$$\begin{array}{c} P_1 \\ P_2 \\ \vdots \\ P_n \\ \hline \text{Therefore, } Q. \end{array}$$

Presented with an argument whose premises we believe, the natural reaction is to accept the conclusion, even if the conclusion is the denial of something else we initially believe. But that is not always the reasonable response to an argument. In the above argument, Q is a deductive consequence of P_1, \dots, P_n , but all that really shows is that we cannot reasonably continue to believe all of P_1, \dots, P_n and $\sim Q$. The validity of the argument does not establish which of these beliefs should be rejected, because we can convert the argument into an equally valid argument for the denial of any one of the premises. For instance, the following is also a valid argument (where $\sim Q$ is short for 'It is false that Q ')

$$\begin{array}{c} P_2 \\ \vdots \\ P_n \\ \sim Q \\ \hline \text{Therefore, } \sim P_1. \end{array}$$

Faced with a skeptical argument, we believe all of the premises P_1, \dots, P_n , but we also believe $\sim Q$ (the denial of the conclusion, the conclusion being that we do not have the knowledge described). The argument establishes that we must reject one of these beliefs, but it does not tell us which we should reject. To determine that, we must reflect upon how certain we are of each of these beliefs and reject the one of which we are least certain. In typical skeptical arguments, we invariably find that we are more certain of the knowledge seemingly denied us than we are of some of the premises. Thus it is not reasonable to adopt the skeptical conclusion that we do not have that knowledge. The rational stance is instead to deny one or more of the premises. In other words, a typical skeptical argument is best viewed as a *reductio ad absurdum* of its premises, rather than as a proof of its conclusion.⁴

3. This is what John Rawls (1971) calls "the method of reflective equilibrium".

4. There is no logical necessity that this should be the case. It is conceivable that

To illustrate, consider inductive reasoning. When we reason inductively, we draw general conclusions from the observation of a finite number of instances. For example, having observed many swans, and noting that they were all white, we may infer that all swans are white. This is typical of the kind of reasoning that is performed in science when we test scientific theories by testing whether particular instances of them are true. Virtually all of our general beliefs about the world are held on the basis of induction. But an obvious feature of inductive reasoning is that the truth of the premises does not logically guarantee that the general conclusion is true. For example, in concluding that all swans are white we would be mistaken, because it turns out that there are black swans in Australia.

In the *Treatise on Human Nature*, Hume used the above observations to propound a skeptical argument against induction. He reasoned as follows:

1. The premises of an inductive argument do not logically entail the conclusion.
 2. If the premises of an argument do not logically entail the conclusion, then it is not reasonable to believe the conclusion on the basis of the premises.
-
3. Therefore, inductive reasoning is illegitimate—one cannot acquire knowledge of general truths by reasoning inductively.

The conclusion of this argument is something that we initially disbelieve. That is, a little reflection convinces us that we know many general truths about the world on the basis of induction. However, the premises of the argument can also seem convincing. The first premise is illustrated by the swan example. And there was a time in the not too distant past when virtually all philosophers took the second premise to be equally obvious. If the premises of an argument can be true without the conclusion being true, then how (they asked) could the premises give us any reason for believing the conclusion?

This is another typical skeptical argument. The premises seem initially compelling, but the conclusion that we cannot have inductive knowledge seems absurd. The strategy suggested above is to treat this as a *reductio ad absurdum* argument. Although the premises seem initially compelling, they seem less certain than the conclusion seems certainly false, so we should take the argument as an argument to the effect that one of the

there should be a skeptical argument whose premises we believe more firmly than we believe that we have the putative knowledge the argument denies us. The claim we are making here is a contingent one about those skeptical arguments that have actually been advanced in philosophy.

premises is false. The epistemological problem then becomes that of deciding which premise is false, and why.

A few philosophers have tried to respond to Hume's argument by denying the first premise. Bertrand Russell (1912) and C. I. Lewis (1956) both suggested that inductive reasoning is based upon an additional premise that turns the inductive argument into a deductive argument. The premise they sought was called "the uniformity of nature". This premise was sometimes formulated rather vaguely as the principle that the future will be like the past. The intention was for this principle of the uniformity of nature to be sufficiently strong that, when conjoined with the inductive evidence, it would provide us with premises logically entailing the inductive generalization. So, for example, because we have in the past observed only white swans, we should expect that swans observed in the future will also be white.

The most obvious difficulty for trying to resolve Hume's skeptical dilemma in this way is that of trying to give a precise formulation of an appropriate principle of the uniformity of nature. It is not sufficient to say that the future will be like the past, because obviously the future isn't *always* like the past. Things change.

But an ultimately more compelling objection is that no principle of the uniformity of nature could do the job required of it. It is supposed to turn inductive reasoning into deductively valid reasoning. Thus in the case of the swans, our actual argument will be as follows:

Swan #1 is white.
Swan #2 is white.
⋮
Nature is uniform.

Therefore, all swans are white.

The difficulty is that the conclusion of this argument is false. If it is deductively valid, it must have a false premise. But all the premises other than that of the uniformity of nature are true, so the premise of the uniformity of nature must be false. As such, it becomes useless to us in reasoning about the world. It is not reasonable to accept conclusions drawn by reasoning from a premise known to be false. Thus if this were the correct form of inductive reasoning, all inductive reasoning would be based upon a false premise, and we would be led to Hume's skeptical conclusion all over again, by a different route.

If we are to avoid Hume's conclusion, we cannot do it by denying the first premise of his argument. Thus it must be the second premise that is false. Contrary to what Hume supposed, it must be possible for the premises of an argument to support a conclusion without logically entailing it. In other words, the premises of the argument can provide a reason for believing the conclusion without the reason being logically

conclusive. Such reasons are *defeasible* in the sense that, while they can justify us in believing their conclusions, that justification can be “defeated” by acquiring further relevant information. In the case of induction, if we observe that all the A’s in our sample are B’s, this may provisionally justify us in believing that all A’s are B’s. But if we subsequently encounter another A and note that it is not a B, that is sufficient to defeat the original justification and it makes it totally unreasonable to continue believing that all A’s are B’s. Thus the skeptical argument about induction points to the existence of defeasible reasons in epistemology. In our opinion, one of the most important advances in epistemology in the last half of the twentieth century was the recognition of defeasible reasons. It is now generally acknowledged that most of our reasoning proceeds defeasibly rather than deductively. Induction is one example of this, but other examples are equally obvious. For example, our beliefs about our surroundings are based upon perception. When we perceive the world around us, it looks various ways to us, and in the absence of conflicting information we conclude that it is the way it looks. But obviously, our senses can mislead us. The world does not *have* to be the way it looks. So our reasoning is defeasible rather than deductive. Defeasible reasoning will be discussed at some length in the next chapter.

This discussion of Hume’s argument illustrates the importance of skeptical arguments for epistemology. What makes skeptical arguments important is not that their conclusions might be true. They are important for what they show about knowledge rather than because they make us doubt that we have knowledge. The task of the contemporary epistemologist is to *understand* knowledge. For this she need not refute the skeptic—we already know that the skeptic is wrong. Nevertheless, important conclusions about the nature of knowledge and epistemic justification can be gleaned from the investigation of skeptical arguments. This is because such an argument constitutes a *reductio ad absurdum* of its premises, and its premises consist of things we initially believe about knowledge and justification. Thus in deciding which of those premises is wrong we are learning something new about knowledge and correcting mistaken beliefs with which we begin. In short, the task of the epistemologist is not to show *that* the skeptic is wrong but to explain *why* he is wrong. The difference between these endeavors is that in the latter we can take it as a premise that we have various kinds of knowledge (i.e., we can assume $\sim Q$) and see what that requires. For example, we might ask, “Given that we have perceptual knowledge, what must the relationship be between our perceptual beliefs and our sensory experience?” The fact that we do have perceptual knowledge will impose important constraints on that relationship and can lead us to significant conclusions about epistemic justification. This reasoning has the form, “We do have such-and-such knowledge; we could not have that knowledge if so-and-so were the case; therefore, so-and-so is not the case.” This kind of reasoning is very common in contemporary epistemology. Note that such reasoning results from contraposing the premises and conclusion of a skeptical argument.