

Can I Know
That I Am Not Dreaming?

DAVID BLUMENFELD
JEAN BEER BLUMENFELD



IN MEDITATION I Descartes offers an argument to show that he cannot know that he is not dreaming. This argument has occupied a central place in the history of modern philosophy: it forcefully raises the problem of the external world, and, at the same time, leads to considerations that Descartes uses in his proof of dualism. Yet it has seemed to many philosophers that some more or less simple maneuver is all that is needed to refute the argument. It will be our contention that this is not so. Some of the criticisms miss the mark entirely; others show that the argument needs to be reformulated in certain ways. But none, in our view, disposes of the fundamental skeptical worry that lies at its heart. Our aim in this paper, however, is not to establish any form of skepticism about the senses. What we do hope to show is that Descartes's argument is a great deal more difficult to refute than has been commonly thought.

I

Descartes begins by considering the possibility that his senses deceive him. A line of argument which he finds compelling is contained in this well-known passage:

How often has it happened to me that in the night I dreamt that I found myself in this particular place, that I was dressed and seated near the fire, whilst in reality I was lying undressed in bed! At this moment it does indeed seem to me that it is with eyes awake that I am looking at this paper; that

this head which I move is not asleep, that it is deliberately and of set purpose that I extend my hand and perceive it; what happens in sleep does not appear so clear nor so distinct as does all this. But in thinking this over I remind myself that on many occasions I have in sleep been deceived by similar illusions, and on dwelling carefully on this reflection I see so manifestly that there are no certain indications by which we may clearly distinguish wakefulness from sleep that I am lost in astonishment. And my astonishment is so great that it is almost capable of persuading me that I now dream. [HR 1:145-46]

As we shall see, the argument stated above is subject to some interesting variations. But as it stands its basic thrust is clear. Descartes thinks that the fact that his current experience has such and such perceptual features is never sufficient to show that he is awake. The fact that he has had dreams whose qualitative character was indistinguishable from that of waking experience convinces him that there can be no experiential criterion whereby he can determine that he is not dreaming. He does not pause to consider a long list of possible criteria, for he thinks that any qualitative feature of a waking experience could be a feature of a dream as well. Hence the presence of such a feature will never provide an answer to the question: Am I awake? One plausible way of setting out Descartes's argument is this:

Argument (D)

- (1) I have had dreams which were qualitatively indistinguishable from waking experiences.
- (2) Therefore, the qualitative character of my experience does not guarantee that I am not now dreaming.

It should be pointed out that the notion of the "qualitative character of my experience" which we use here to capture the intent of Descartes's argument is to be construed quite broadly. It includes, for example, my current recollections (memory impressions) of how things have been in the past. Given this broad sense of the phrase, one item that would count as an aspect of the qualitative character of my experience is its apparent coherence with my past. But it would be Descartes's (Meditation I) view that the satisfaction of this criterion is not sufficient to distinguish dreaming from waking. Since it could be that I am only dreaming that my present experience is coherent with my past experience, this criterion fails to provide a way of establishing that I am awake.

Descartes does not ultimately conclude from (2) that he cannot know that he is not dreaming. For although he thinks that he cannot know this simply on the basis of the qualitative character of his experience, he believes that he can go on to prove the existence of a veracious God,

and that one who is in possession of this knowledge can know that he is not dreaming. Now, his strategy is notoriously dubious, and it has been criticized by so many others that we propose to ignore it. Clearly, Descartes's contribution to the topic of dreaming lies more in the skeptical considerations he raised than in the manner in which he proposed to combat them. (D) can be converted into a fully skeptical argument by the addition of the following (non-Cartesian) steps:

- (3) If the qualitative character of my experience does not guarantee that I am not now dreaming, then I cannot know that I am not now dreaming.
- (4) Therefore, I cannot know that I am not now dreaming.
- (5) If I cannot know that I am not now dreaming, then I cannot know that I am not always dreaming.
- (6) Therefore, I cannot know that I am not always dreaming.
- (7) If I cannot know that I am not always dreaming, then I cannot know to be true any belief which is based on my experience.
- (8) Therefore, I cannot know to be true any belief which is based on my experience.

Two qualifications with respect to (7) and (8) are necessary. Many skeptics would allow that there is a special sort of belief which might be said to be "based" on my experience, but whose certainty is unimpugned by the possibility that I am dreaming. These are beliefs which concern the character of my current experience itself, and which can be expressed by sentences of the form "It seems to me that . . .". Since we do not wish to discuss the issue of whether such beliefs do indeed have a special status, we want to stipulate that (7) and (8) are to be understood as excluding these beliefs. Secondly, a skeptic influenced by Descartes might hold that my belief that I exist is "based" on my experience of thinking, but that the certainty of this belief is unimpugned by the possibility that I am dreaming. Again, we do not wish to discuss the issues connected with the cogito, and so we want to stipulate that (7) and (8) are to be read as making an exception of my belief that I exist. Finally, since our primary concern is with the problems that Descartes left us, we shall direct our attention to the fully skeptical argument which includes steps (3) through (8); it is this that we shall refer to as argument (D).

II

Premise (1) of argument (D) is apparently denied by J. L. Austin in *Sense and Sensibilia*. In criticizing A. J. Ayer's claim that "'there is no intrinsic difference in kind between those of our perceptions that

are veridical in their presentation of material things and those that are delusive,'" Austin says,

Is it the case that 'delusive and veridical experiences' are not 'qualitatively different'? Well, at least it seems perfectly extraordinary to say so in this sweeping way. Consider a few examples. I may have the experience (dubbed 'delusive' presumably) of dreaming that I am being presented to the Pope. Could it be seriously suggested that having this dream is 'qualitatively indistinguishable' from *actually being* presented to the Pope? Quite obviously not.¹

But this is just the sort of thing Descartes *is* seriously suggesting. On his view, it is possible to have a dream which is qualitatively indistinguishable from actually being presented to the Pope. To refute such a suggestion, one would need an argument, and Austin does offer one in the passage that immediately follows. He says,

After all, we have the phrase 'a dream-like quality'; some waking experiences are said to have this dream-like quality. . . . But of course, if the fact here alleged *were* a fact, the phrase would be perfectly meaningless, because applicable to everything. If dreams were not 'qualitatively different' from waking experiences, then *every* waking experience would be like a dream.²

Austin clearly has Descartes (among others) in mind here, as he adds in a footnote, "This is part, no doubt *only* part, of the absurdity in Descartes' toying with the notion that the whole of our experience might be a dream."³ Now, while it seems that the second passage cited from Austin contains an argument directed against something like our premise (1), it is not altogether clear what the conclusion of this argument is meant to be. Nevertheless, we believe that there are only three possibilities:

- (A1) Every dream is qualitatively different from waking experience.
- (A2) Most dreams are qualitatively different from waking experience.
- (A3) Some dreams are qualitatively different from waking experience.

Each of these three interpretations of Austin's meaning renders his argument problematic, however. (A2) and (A3) may be plausible in themselves, but each of these is consistent with premise (1). Premise (1) says that some dreams are qualitatively indistinguishable from waking experience, and neither (A2) nor (A3) contradicts this. (A1), on the other hand, does contradict premise (1). But (A1), unlike (A2) and (A3), is dubious. If Austin meant to endorse (A1), then his argument for it is apparently this:

- (A4) If it is not the case that every dream is qualitatively different from waking experience, then every waking experience is dream-

like, and hence the phrase "a dreamlike quality" is meaningless because applicable to everything.

But it seems perfectly possible that the following situation should obtain. Some dreams are vague and unclear, and in virtue of this fact are said to possess a dreamlike quality. Other dreams are clear and vivid, and thus do not possess a dreamlike quality. Most waking experiences are clear and vivid, although some few are vague and fuzzy. In such a situation, it might be that the clear and vivid dreams are qualitatively indistinguishable from the clear and vivid waking experiences, i.e., the antecedent of (A4) would be true. But it certainly would not follow that the clear and vivid waking experiences possess a dreamlike quality, nor that the phrase "a dreamlike quality" is meaningless. Hence (A4)—the support for (A1)—is false, and Austin has not succeeded in undermining argument (D).

III

Anthony Kenny attempts to show that premise (3) is false. According to (3) I can know that I am awake only if the qualitative character of my experience guarantees that this is so. In other words, I can know that I am not now dreaming only if there is some feature of my experience from which I can infer that I am awake. It is Kenny's view, on the other hand, that I can know that I am awake without appealing to the character of my experience. In fact, on Kenny's view I know that I am awake without appealing to any criterion at all. With respect to the question "Am I awake?" Kenny says, "there is a true answer to the question—namely, 'I am awake.' Moreover, I know this answer. . . . When I say 'I am awake,' I do so without grounds, but not without justification."⁴ This justification is the alleged fact that in order to make a judgment or entertain a belief, one must be awake. Kenny agrees with Norman Malcolm that one cannot make judgments during dreams.⁵ Hence, any time that I judge or believe that I am awake, it will be true that I am awake. For, any time that I make any judgment whatsoever, it will be true that I am awake. Kenny's claim is, "The judgment 'I am awake' cannot be mistaken. . . . The question 'Am I awake?' . . . is pointless . . . to the extent that if a man is in a position to ask the question, he is also in a position to answer it."⁶ His argument to show that one can know that he is not dreaming can be summarized as follows:

- (K1) If I judge that I am awake, I know that I am awake.
- (K2) I now judge that I am awake.
- (K3) Therefore, I now know that I am awake.

It might be objected that people do indeed make judgments during sleep, including some true ones. For example, during a dream one sometimes makes a correct mathematical judgment. And, it is quite common for a sleeping person to make the judgment that he is dreaming. So, it is not obvious that the support which Kenny gives for (K1) is satisfactory. But suppose that we grant (K1) for the sake of argument. It is nevertheless difficult to see what establishes that (K2) is true. One might be inclined to think that it is evident to me when I am judging. But in thinking this over I must remind myself that on many occasions in sleep it has seemed to me that I was judging. That is, it has seemed evident that I was judging when I was really only dreaming that I judged. So, the fact that I seem to judge is not sufficient to show that I do judge, nor to establish that (K2) is true. If we grant that judgments cannot occur in sleep, the question for Kenny is: How do you know that you are now *judging* that you are awake? Perhaps you are only dreaming and so it merely *seems* to you that you are judging.

Of course, many philosophers—Descartes included—have held that if it seems to me that I judge, then I do judge. To adhere to this thesis consistently, however, one must also hold that in dreams, when I seem to judge, I do judge. Kenny clearly cannot enjoy the benefits of this view, for according to (K1) judgments are not made during dreams. Thus, for Kenny, a necessary condition of making a judgment is that one be awake, and it is impossible to know that one is judging simply on the basis of the fact that it seems to one that he is judging. Kenny's claim that (K1) is true forces him to justify his claim that (K2) is true.

Unless he were to maintain that in dreams we do not seem to judge,⁷ only one reply is available to Kenny. This is that one *can* know that he judges: he need only judge that he judges. To this, the Cartesian question will be: How does one know that he judges that he judges? Perhaps he is only dreaming that he judges that he judges. Kenny, of course, can make his move again, as can Descartes, ad infinitum. But the regress is clearly fatal to Kenny's position, not Descartes's. For Kenny has claimed to be able to dispose of the Cartesian question. In fact what he has done is only to push it back to another level. At each level the question reemerges, and, for this reason, Kenny cannot justifiably claim to have answered it.

IV

Another interesting critique of argument (D) is offered by G. E. Moore in his paper "Certainty." In trying to set out the reasoning

which leads many philosophers to the view that one does not know that he is not dreaming. Moore says, "one premiss which they would certainly use is this: 'Some at least of the sensory experiences which you are having now are similar in important respects to dream-images which actually have occurred in dreams'."⁸ Moore thinks that this premise is true (and indeed harmless), but he also thinks,

there is a very serious objection to the procedure of using it as a premiss in favour of the derived conclusion. For a philosopher who does use it as a premiss, is, I think, in fact *implying*, though he does not expressly say, that he himself knows it to be true. He is *implying* therefore that he himself knows that dreams have occurred. And, of course, I think he would be right. . . . But can he consistently combine this proposition that he knows that dreams have occurred, with his conclusion that he does not know that he is not dreaming? Can anybody possibly know that dreams have occurred, if, at the time, he does not himself know that he is not dreaming? If he *is* dreaming, it may be that he is only dreaming that dreams have occurred; and if he does not know that he is not dreaming, can he possibly know that he is *not* only dreaming that dreams have occurred? Can he possibly know therefore that dreams *have* occurred? I do not think that he can; and therefore I think that anyone who uses this premiss and also asserts the conclusion that nobody ever knows that he is not dreaming, is guilty of an inconsistency.⁹

Now it seems to us that Moore's point is correct, and further, that it is successful in refuting argument (D). For, if I cannot know that I am not now dreaming, then I cannot know that my belief that I have had dreams in the past is true. Therefore, if I can know that (1) is true, subconclusion (4) is false. But this shows that argument (D) is invalid. For it shows that the truth of (1), (2), and (3) does not guarantee the truth of (4): (4) is false if I can know (1) to be true.

It might be thought that the Cartesian could salvage his argument simply by denying that (1) can be known to be true. Moore, however, points out the difficulty with such a tactic: "a philosopher who uses it [(1)] as a premiss is, I think, in fact *implying*, though he does not expressly say, that he himself knows it to be true" (p. 249). If (1) cannot be known to be true, there is an oddity in using it as a premise in an argument. For if the truth value of one of the premises of an argument is unknowable, then the argument does not provide knowledge of the conclusion.

But while Moore's criticism of (D) is correct, we believe that the skeptical point about dreaming can be made without the claim that deceptive dreams have occurred. One who does not know whether he has actually had any dreams may still have the *concept* of a dream. That is, he may have the notion of a nonveridical experience which occurs during sleep¹⁰—even though he is in doubt whether he has

undergone such an experience or, in fact, ever been asleep. Given this concept he is in a position to argue—quite apart from whether he has had any dreams—that it is logically possible that one should have a dream which is qualitatively indistinguishable from waking experience. The skeptic could then repair argument (D) by substituting the following in place of premise (1):

(1') It is logically possible that my experience should have just the qualitative character that it does have, and yet that I be dreaming.

From (1') he could legitimately move to (2) and from there the argument would proceed exactly as stated originally. Let us label the amended argument, consisting of (1') and (2) through (8) of argument (D), argument (D').

Moore very perceptively anticipates this move on the skeptic's part. Having (in effect) countered argument (D), he goes on to ask:

But what if our skeptical philosopher says . . . It is logically possible *both* that you should be having all the sensory experiences you are having, and also that you should be remembering what you do remember, and *yet* should be dreaming. If this *is* logically possible, then I don't see how to deny that I cannot possibly know for certain that I am not dreaming: I do not see that I possibly could.

Moore's reply is as follows:

But can any reason be given for saying that it *is* logically possible? So far as I know nobody ever has, and I don't know how anybody ever could. And so long as this is not done my argument, 'I know that I am standing up, and therefore I know that I am not dreaming', remains at least as good as his, 'You don't know that you are not dreaming, and therefore don't know that you are standing up'.¹¹

Now Casimir Lewy, the editor of Moore's *Philosophical Papers*, tells us that Moore was particularly dissatisfied with the last four paragraphs of the paper "Certainty," in which this rebuttal to the skeptic occurs.¹² Lewy does not mention the source of Moore's dissatisfaction, but perhaps it was something like the following. Moore was interested in refuting skepticism. That is, he was interested in showing that one can assert firmly and without reservation that he knows that he is not dreaming. But Moore's argument does not show this. For even if we grant him everything that he claims, he still has not shown that premise (1') is false, nor that his own argument is *better* than (D'). If Moore is right in claiming that his own argument is "at least as good" as (D'), this will at best achieve a stand-off. Barring further argument, the upshot of the stand-off is a slightly weaker form of skepticism, namely, that I do not know whether I can know that I am

not dreaming.¹³ Moore's strategy has placed him in a position that would have delighted the old Pyrrhonian skeptics. They argued, it will be remembered, that for every plausible argument, an equally plausible counterargument can be given. It was for this reason that they thought one must suspend judgment. It appears, then, that a version of skepticism remains untouched by Moore's attack. In order to avoid Pyrrhonism, one would have to refute argument (D'), and not simply offer an argument that is equally plausible.

V

An initially attractive attempt to refute (D'), which employs the notion of probability, runs as follows. Premise (3) says that if the qualitative character of my experience does not guarantee that I am not now dreaming, I cannot know that I am not now dreaming. But why suppose that a guarantee is needed? The fact that I can have no such guarantee only shows that I cannot be absolutely certain that I am not now dreaming; it does not show that my experience fails to provide me with grounds for asserting that it is extremely probable that I am awake. Unless the skeptic wishes to defend the implausible view that (for empirical propositions) "X knows that *p*" entails "X currently has experience whose qualitative character guarantees that *p*," there is no support for (3). For if my experience is sufficient to render it highly probable that I am not dreaming, then I can know that I am awake, even though my experience does not guarantee that this is true. It is quite possible, then, to concede the truth of (1') and (2) while denying (3).

The criticism, however, is superficial. Once having conceded that he cannot rule out the possibility that he is dreaming, how can the critic of (3) go about supporting the claim that it is probable that he is awake? To justify it he must appeal to his past experience, citing, for example, the fact that his past beliefs of this kind (when grounded on experiences such as he has now) have been true more often than not. But if he allows that it is possible that he is dreaming, it appears that he is barred from relying on the data of the past. For he has conceded that all such data may be merely dream-data, and hence totally unreliable. If his memory impressions are merely states of a dream, then, for all he knows, they are no index at all of what has actually occurred. Thus, the possibility that one may be dreaming throws into question the reliability of memory, and with it, all of that material from the past on which a judgment of probability must ultimately depend. To find a basis for a probability judgment, it appears that one must first be able to rule out the possibility that he is dreaming.

It might be objected that it is not true that in order to be justified in relying on memory one must rule out the possibility that he is dreaming. For, all one needs to know is that his memory is *probably* reliable. If he knows this, then he can use memory in making empirical judgments. The difficulty with this suggestion, however, is that it is not clear how one can establish that his memory is probably reliable. It would seem that in order to do this he must again make use of memory, assuming that it has been generally reliable in the past. But this justification is obviously circular.¹⁴

The skeptic's reply to the probability argument does *not* commit him to holding that (for empirical propositions) "X knows that *p*" entails "X currently has experience whose qualitative character guarantees that *p*." His reply commits him only to the following: if X's claim to know that *p* is based on his view that *p* is probable, then he can support the former claim by means of the latter only if he can rule out the possibility that he is dreaming. The skeptic can allow that if one is justified in claiming that *p* is probable, it might very well be true that he knows that *p*—even though his experience does not guarantee that *p*. The skeptical difficulty is that unless one can rule out the possibility that he is dreaming, he cannot ever be justified in claiming that *p* is probably true (where *p* is an empirical proposition whose truth is not guaranteed by the character of his experience).¹⁵

VI

We have so far considered a number of attacks on the first four steps of argument (D'). Attempts to forestall its conclusion, however, have also consisted in attacks on (5). One such attempt is nicely stated by Bernard Williams in a recent discussion of the philosophy of Descartes. In criticizing Descartes's skepticism, Williams comments:

even if . . . considerations of past error were allowed to show that *any* occasion of supposed perception might be illusory (because, for instance, I might be dreaming), there would be no valid inference from this to the supposition that *every* supposed occasion of perception might be illusory.¹⁶

Williams's claim is that (6) does not follow from (4), that is, that premise (5) is false. According to Williams, even if I cannot know that I am not now dreaming, I can know that I am sometimes awake.

We want to argue, however, that the inference from (4) to (6) is actually legitimate. (4) says that at any time, for all I know, I may be dreaming at that time. But take any time, say now—if I am dreaming now, then my belief that not all my experiences have been dreams is itself a belief held in a dream, and hence it may be mistaken. If I am

dreaming now, then my recollection of having been awake in the past is merely a dreamed recollection and may have no connection whatever with reality. The point we are making could be described as the mirror image of the one made so effectively by Moore against argument (D). Moore pointed out that if I cannot know that I am not now dreaming, then I cannot know that I have had dreams in the past. I might only be dreaming that I have had dreams. Thus Moore showed that the skeptic is not entitled to the claim that he has had dreams. Our point is this: if I cannot know that I am not now dreaming, then I cannot know that I have been awake in the past. I might only be dreaming that I have been awake. Thus the antiskeptic is not entitled to assert that he has been awake. The truth of (4) generates a special epistemic context: given (4), I am, for all I know, "inside a dream," and my belief that I have been awake may be mistaken.

It should be noted that our defense of (5) depends on the idea that if I believe that I have been awake in the past, this belief is to be justified by my apparent memories of having been awake. The argument is that if I am dreaming now, then my current memories are not to be trusted—they are not a reliable index of what has actually occurred, and hence cannot provide the required justification. It might be objected, however, that this argument still does not show that (6) follows from (4). What it shows is that, if (4) is true, there is no way that I can show on a posteriori grounds that I have ever been awake. To defend the claim that (6) follows from (4) it would have to be shown as well that (6) is not false a priori. For while the truth of (4) would undermine my beliefs whose justification depends on my past experience, it is not clear that it would undermine beliefs which are known a priori. Even Descartes held that a priori propositions (such as " $2 + 2 = 4$ ") can be known to be true in a dream (HR 1:147). It appears, then, that we still need to rule out the possibility that (6) can be known to be false a priori.

But how plausible is it to suppose that I can know on a priori grounds that I have not always been dreaming? The negation of (6), of course, could be established if there were some a priori incoherence in the very idea of a being who slept (and dreamt) all of his life. But clearly there is none. We can easily imagine a malevolent experiment in which an infant is drugged in utero so that it will sleep—and dream—all of its life. The child is born, artificially sustained for a period of years, and finally dies, having undergone nothing but dream experiences. But perhaps it will be said that while there is no incoherence in the idea of some being or other (e.g., a baby with only rudimentary sorts of experiences) dreaming throughout its entire life, this does not show that it is possible for a being with experiences such

as I have to be dreaming all of its life. To show that (6) follows from (4), the case of the baby will not suffice. For perhaps I can know a priori that no one with experiences of a certain highly complex type could be continually dreaming. But this also seems very doubtful. To see this, let us imagine our evil experimenter equipped with a machine conceived by Keith Lehrer:

The machine operates by influencing the brain of a subject who wears a special cap, called a 'braino cap'. When the braino cap is placed on the head of a subject, then the operator of the braino can affect the brain of the subject in such a way as to produce any hallucination in the subject that the operator of the braino wishes. The braino is a superhallucination-producing machine. The hallucinations produced by it may be as complete, systematic, and coherent as the operator of the braino desires to make them.¹⁷

If (6) can be known to be false independently of experience, then I can know a priori that I have not been the victim of a uterine braino experiment. Since it is clear that I cannot know this a priori, we believe that the inference from (4) to (6) cannot be blocked on a priori grounds. And, since we have already argued that it cannot be blocked on a posteriori grounds, we take it that the inference is valid. Some philosophers, however, have offered independent argument to show that (6) is a priori false. We shall consider these arguments, along with some closely related variants, in the next section.

VII

Perhaps the most widespread criticism of (D') is what might be called the contrast argument, (CA). There are various versions of (CA), but what they have in common is the idea that the skeptic's conclusions, (6) and/or (8), could not be true because it is not possible that one should always be mistaken. The idea of being mistaken, it is claimed, requires the existence of the contrasting state of being correct. The first version we want to discuss, (CA1), is mentioned by Anthony Kenny:

Critics have argued that sense deception is only possible against a background of veridical perception. There cannot be errors, it is reasoned, where there is no possibility of correction, for if it makes no sense to talk of something's being corrected, then it makes no sense to talk of its being wrong.¹⁸

J. L. Austin gives much the same argument in the following passage from *Sense and Sensibilia*:

Next, it is important to remember that talk of deception only *makes sense* against a background of general non-deception. (You can't fool all of the

people all of the time.) It must be possible to *recognize* a case of deception by checking the odd cases against the more normal ones.¹⁹

Apparently, Austin is making two claims here. His first claim is that talk of deception only makes sense if it is possible to recognize cases of deception. If this is true, both (6) and (8) are senseless, since each implies that an undetectable deception is possible. His second claim is that it is possible to recognize cases of deception only if there is a background of general nondeception. We do not wish to discuss the second claim. But the first claim seems to us to depend on a principle which is highly suspect. The principle is: no fallibility without detectability. A claim that a belief may be in error *makes sense* only if it is possible to determine whether it really is in error. In other words, unless it is possible to determine whether a statement is true, that statement is meaningless. This is, of course, (one form of) the verification principle. The objections to this principle, however, are well known and it would be pointless to try to review them here. But in view of the criticisms of verificationism which are available, (CA1) cannot be regarded as very powerful.

Gilbert Ryle, on the other hand, gives a version of the contrast argument which does not depend upon the verification principle. In *Dilemmas* he comments,

I must say a little about the quite general argument from the notorious limitations and fallibility of our senses to the impossibility of our getting to know anything at all by looking, listening and touching.

A country which had no coinage would offer no scope to counterfeiters. There would be nothing for them to manufacture or pass counterfeits of. They could, if they wished, manufacture and give away decorated discs of brass or lead, which the public might be pleased to get. But these would not be false coins. There can be false coins only where there are coins made of the proper materials by the proper authorities.

In a country where there is a coinage, false coins can be manufactured and passed; and the counterfeiting might be so efficient that an ordinary citizen, unable to tell which were false and which were genuine coins, might become suspicious of the genuineness of any particular coin that he received. But however general his suspicions might be, there remains one proposition which he cannot entertain, the proposition, namely, that it is possible that all coins are counterfeits. For there must be an answer to the question 'Counterfeits of what?'²⁰

Ryle's claim is that in order for counterfeit coins to exist, real coins must exist as well. He does not explicitly apply his point to the case of dreaming, and there are several ways in which the analogy could be used. One interpretation of Ryle, (CA2), is that he wants to maintain that not all of my experiences could be dream experiences. In

order for me to have dreams, I must have veridical experiences as well. If this is correct, then I know a priori that (6) is false.

(CA2) argues from an analogy with the case of currency to the conclusion that not all *my* experiences could be dream experiences. In order to derive this conclusion from the currency example, however, it would have to be claimed that not all my money could be counterfeit. But even if we grant that it is impossible that all coins should be counterfeit, this does not entail that it is impossible that all of my coins should be counterfeit. Surely we can imagine circumstances in which I possess nothing but counterfeit coin. So, the analogy provides no reason to deny that it may be the case that all my experiences are dreams. For all Ryle has shown, it is possible that I should have been asleep and dreaming since birth.

Perhaps (CA2) does not capture Ryle's intention, however. He says that not all coins of the realm (i.e., all the coins there are) could be counterfeit. So perhaps he means to claim only that it could not be the case that *everyone's* money should be counterfeit. In this case, all that is to be inferred from the analogy is that it is not possible that all the experiences there are should be dreams. If there are any dream experiences, then someone, at least, must have experience of real things. Let us dub this way of pushing the analogy (CA3). (CA3) makes a very weak claim against skepticism. Its success would not imply that any of my experiences are veridical, and so it does not place the truth of (6) in jeopardy. Nevertheless, (CA3) would guarantee that someone or other has veridical experience. So, if it succeeds then at least one of my beliefs based on my experience is true, namely, my belief that in addition to myself and my experiences, there exists an external world which is independent of any dream I may have. For, in order for anyone to have veridical experience, there must be an independent world which he experiences. Hence, if (CA3) is correct, it refutes premise (7) of argument (D').

(CA3) rests on the claim that it is logically impossible that everyone's money should be counterfeit. But this claim is debatable. It is not evident that there are no conceivable circumstances under which we would say that all money is counterfeit. Consider the following possibility. Just as the first money is about to be printed, a band of criminals seizes the presses and issues its own currency, which is a facsimile of the original design. Later, when the shady origins of the currency are exposed, the community forever drops the institution of money. In such a situation, it would appear that all the money that has ever existed has been counterfeit.

Still, there is probably something to be learned from Ryle's analogy that has not been brought out by our discussion so far. We suspect that

the intuitive idea behind (CA3) comes from imagining the following sort of situation. Suppose that the criminals described above are never caught, their plot never uncovered. The money they issue is accepted by the people as legitimate and is used by all. If we imagine this situation continuing indefinitely, then perhaps many philosophers would find it odd to describe such money as "counterfeit." They would be inclined to say that what functions systematically and continuously as money *is* money (real money), whatever its origins. Intuitions about this rather fanciful case may differ, and there is, if course, room for dispute.²¹ But let us grant, for the sake of argument, that the money should be regarded as real. If so, then surely this would be for one fundamental reason: the currency is used by everyone as money in a systematic way. Similarly, the opponent of the skeptic may want to argue that if everyone were always "dreaming," and if everyone's experiences were systematically correlated, then it would be incorrect to describe such experiences as "dreams." Such a world would be quite like the one Berkeley thinks is the actual world (leaving theological considerations aside). In Berkeley's world, an object is real if it can be experienced intersubjectively: it is not required that it have material existence. In the currency example, the money is real if it is commonly used: it is not required that it be produced by the government. In both cases it is the *common* coin which is the real one.

Now, we agree that if the Berkeleyan situation did obtain, it would not be one in which all experiences are delusive. Given the Berkeleyan assumption that there are other minds which have experiences similar to mine, sense can be given to the idea that there are real objects. (CA3) can be resuscitated, then, if we can find some way of using the analogy with counterfeit money to show that we are justified in believing that this assumption is true. But it is precisely here that the analogy is weakest. For in the coin case there is no difficulty in principle in determining whether the coins are commonly used. In the dreaming case, however, there is a difficulty in determining whether the experiences I have are commonly shared. In fact, we cannot assume that there are other people without begging the question against the skeptic. Remember that (CA3) is a denial of premise (7); it does not purport to deny (6). But if I cannot know that I am not always dreaming, then I cannot know that other people's experiences are correlated with my own; nor can I even know that other people exist. Without this knowledge, however, I cannot know that the Berkeleyan situation obtains, and thus cannot use the idea of commonly shared experiences to resuscitate (CA3).

There is yet a fourth version of the contrast argument, (CA4). According to (CA4), a dream is, by definition, a merely subjective experi-

ence occurring during sleep. To call something a "dream" is to contrast it with a "real" or "nondream" state of affairs to which one might, conceivably at any rate, wake up. The skeptical supposition that there might exist nothing but the dreamer and his dream experiences is contradictory. For it is impossible that there should be dream experiences unless there is also an external world independent of the dreamer. So, if I am always dreaming, that in itself would imply that my belief that there is an external world is true.

Perhaps it is true that the word "dream" is normally used in such a way that it implies that the dreamer is asleep.²² But this fact cannot seriously damage the skeptic's position. He uses the word because he thinks of dreaming as a purely subjective state in which one may be deceived about the existence of objects. When he denies that one can know that he is awake, the point he is making is that one cannot know that there is an external world. If he is told that the applicability of the words "dream" and "sleep" implies the existence of an independent world, then surely he is free to stipulate that he employs these words without any such connotation. This decision will involve no intellectual confusion as long as he makes clear how he uses the terms. Having done this, the skeptic need say no more to dispose of (CA4), which makes a purely verbal point.

VIII

A frequent charge against skepticism is that it shows that we cannot have knowledge only by adopting an implausibly strong definition of knowledge, viz., one that makes certainty a necessary condition. According to this view, although the certainty of empirical beliefs is unattainable, it is not required for knowledge. All that is required for knowledge is the justification of the belief and its truth.²³ In section V we argued that empirical propositions cannot be justified on the basis of probability, and, further, that this skeptical claim does not imply that in order to know that p , one must be certain that p . But, of course, there could be some other sort of epistemic justification for empirical beliefs, and in particular, for the belief that an external world exists. In this section we would like to pursue this suggestion. We will assume, along with the opponent of skepticism, that certainty is not a requirement for knowledge.

It should be noted, first, that the idea that my experiences are caused by an external world provides me with an explanation of the existence and character of my experiences, whereas the belief that these experiences are mere illusions or dreams does not. This fact might be used as a basis for the claim that the external-world hypothesis is

epistemically preferable to the dreaming hypothesis. But, following Descartes's strategy in Meditation I, the skeptic could say that, for all one knows, there is an evil demon who, out of the desire to deceive him into believing there is an external world, causes his experiences (his dreams) to be as they are. In this case, the skeptic too will have a logically possible account of all of his sensory data. One would then have to show that the hypothesis of an external world of physical things is epistemically superior to that of a powerful and deceitful demon. One might think that this could be argued on grounds of the greater simplicity of the external-world hypothesis. But it is hard to see in what respect the external-world hypothesis is simpler than that of the demon. The latter is committed to the existence of the demon (a spirit) with the means of and a motive for producing sense experiences, to a mind in which these experiences are produced, and to the sense experiences themselves. The external-world hypothesis, on the other hand, is committed to all of the above, except the existence of the demon. But it is committed, in addition, to a physical world with the capability of producing sense experience. So, it is hard to see *how* the external-world hypothesis is simpler.

If a case for the greater epistemic reasonableness of the external-world hypothesis over that of the demon can be made out, it is surprising how little effort there has been to show this in detail. A notable exception, however, is an intriguing argument given by Michael Slote. Slote's aim is to establish that there is a valid principle of rational scientific inquiry which the hypothesis of a demon violates, but which the hypothesis of an external world does not.²⁴ He hopes to prove thereby that there is epistemic justification for accepting the latter hypothesis over the former. To prepare the way he introduces the notion of an "inquiry-limiting hypothesis." An hypothesis is inquiry-limiting as an explanation of certain phenomena just in case one who accepts it "ensures the impossibility of his coming to have rationally justified or warranted belief (consistent with his other beliefs) in more and more true explanations of various aspects of or facts about the phenomena in question." Slote argues that, other things equal, it will be unreasonable from the scientific point of view to accept an inquiry-limiting hypothesis over one which is not inquiry-limiting. Science, he says, is an enterprise which

seeks to give explanations of events, processes, etc., but also to give, as far as possible, explanations of all the various aspects of the very things it posits in its explanations. For it is a goal of scientific enterprise to gain deeper and deeper and more and more explanations of whatever things there are in the world, wherever possible.

With this conception of the goals and purposes of science in mind, Slote proposes the following two-part principle, which he refers to as the *Principle of Unlimited Inquiry*:

(a) that it is scientifically *unreasonable* for someone to *accept* what (he sees or has reason to believe) is for him at that time an inquiry-limiting explanation of a certain phenomenon, other things being equal; and (b) that there is *reason* for such a person to *reject* such an explanation in favour of an acceptable non-inquiry-limiting explanation of the phenomenon in question, if he can find one.

According to Slote, this principle provides the scientific and epistemic rationale for preferring the hypothesis of the external world over that of the demon. Clearly there is no reason to suppose that by believing in the external world one prevents himself from gaining more and more warranted beliefs in true explanations of various aspects of nature. So, there is no reason to suppose that this hypothesis is inquiry-limiting. But the hypothesis of a demon who causes one's experiences is inquiry-limiting. For the demon theory implies that all the data one wants to explain are supplied by a powerful deceiver. Hence, for any further explanation one conceives, he has no more reason to suppose that it is correct than he has to suppose that it is merely another product of the demon's deception. Even if there is a demon, a belief in his existence cuts one off from the possibility of gaining more and more warranted beliefs. As Slote puts it, "whatever in fact is the explanation of our sense experiences, if we accept the above sort of demon-hypothesis . . . we thereby frustrate our purposes as scientists." Consequently, we have rational scientific justification for rejecting the demon-hypothesis and believing in the external world.

Let us suppose, with Slote, that the demon-hypothesis is inquiry-limiting and that the external-world hypothesis is not. Allowing this, however, it is still a matter for debate how much of a threat this would pose to skepticism. For what Slote's argument purports to show is that there is *some* reason to reject the demon-hypothesis in favor of a belief in the external world. Even if certainty is not required for knowledge, it would have to be shown that the degree of justification conferred by his argument is adequate for knowledge. But we shall not pursue this issue, for we believe it can be shown that the argument does not provide *any* epistemic justification for a belief in the external world.

The situation is this. We have granted that the demon-hypothesis is inquiry-limiting. We are also willing to grant that this provides some sort of reason for rejecting the hypothesis. But it is absolutely essential for Slote to be able to show that the reason or justification in question really is epistemic justification, i.e., justification of the sort which could provide support for a truth claim. As he is well aware,

to combat skepticism, it will not do to show that there are non-epistemic reasons for rejecting the demon-hypothesis. It is no part of argument (D') (nor of most other forms of skepticism) to deny that there could be various sorts of nonepistemic reasons for believing a given proposition. The question, then, is: in what *sense* does the fact that the demon-hypothesis is inquiry-limiting make it reasonable to reject it? It seems to us that the sense in which it does this is a practical (or pragmatic) rather than a theoretical (or epistemic) one. Slote's argument shows that one who has as his goal the obtaining of more and more warranted beliefs, but who accepts the demon-hypothesis, does something which makes it impossible for him to achieve this purpose. As he puts it, "if we accept the . . . demon-hypothesis . . . we thereby frustrate our purposes as scientists." But an argument of the form "You will frustrate your purposes unless you do such and such" is one that provides a practical reason for adopting a certain course of action. It is not normally one which provides an epistemic reason for belief (i.e., rational grounds that count in favor of the truth of the belief). It is an interesting feature of Slote's argument, however, that the "course of action" for which it gives support is the adoption of a belief. What he has produced is a practical justification for adopting a certain epistemic or theoretical attitude if one can (or, presumably, for trying to maintain the attitude if one already has it). Since the conclusion of his argument is that holding a certain belief would be reasonable, one may be misled into thinking that he has given epistemic support for this belief. But in fact he has only shown that holding the belief would be *practically* reasonable.

It may be objected, however, that although the argument is practical (indicates what is reasonable considering certain ends), the ends in this case are one's goals as a scientist or truth-seeker, and that this makes a crucial difference. If the argument shows that by accepting the demon-hypothesis one defeats his purposes as a truth-seeker, the support it provides is surely epistemic as well as practical. Now, we admit that if Slote's argument actually established that a belief in the external world is likelier to be *true*, it would provide the epistemic support which he desires. But it plainly does not do this. Recall that the demon-hypothesis was *not* ruled out because it is less likely to be true. It was ruled out because, even if it is true, its acceptance would cut one off from the possibility of having warranted beliefs about the nature of more and more phenomena. But an argument which gives no reason for supposing that its conclusion is true cannot provide epistemic support for that conclusion.

Slote does offer a direct argument to show that the considerations he has adduced are specifically epistemic. The argument is based on

a definition of 'epistemic reasonableness' which he attributes to Chisholm. Slote says that Chisholm "defines '*p* is epistemically reasonable in believing *q*' as 'if *p* were a rational being, and if his concerns were purely intellectual, it would be reasonable (i.e., a good thing) for him to believe *q*.'"²⁵ But Slote thinks this will suffice only if "one understands 'his concerns were purely intellectual' to mean 'his concerns were those of intellectual (or theoretical) understanding and knowledge in *general*, and included no other concerns or interests.'" So understood, however, he finds the definition plausible and conducive to his purposes. For the demon-hypothesis is an inquiry-limiting explanation; and, on the part of one whose concerns are purely intellectual, it is reasonable (a good thing) to reject such an explanation in favor of one that is not inquiry-limiting. Slote concludes that "it is *epistemically* reasonable to have tentative belief in an external world."

Note that this account of epistemic reasonableness is much broader than the one we have relied on. Slote defines the concept in terms of purely intellectual interests, but his definition does not imply that if one belief is epistemically more reasonable than another, then the former is likelier to be true. It is for this reason, we believe, that his definition is open to counterexamples. Consider the following case. I am a citizen in a totalitarian state in which it is a condition of my entering school, or of my doing any sort of serious research or study, that I accept a certain proposition, *c*. Let us suppose that I have no reason for believing *c* to be true or for believing it to be false except this: I will arrive at a much larger body of warranted true explanations of phenomena if I accept *c* than if I do not. From the point of view of my purely intellectual concerns it would be reasonable (a good thing) for me to accept *c* if I can. If I reject *c*, I will be allowed to learn nothing; if I accept *c*, I will almost certainly learn a great deal. Thus on Slote's definition I would have good epistemic or theoretical grounds for accepting *c*. But we submit that this is clearly not the case. Although I may have a very good practical reason to accept *c* if I can, I have no theoretical justification for doing so. For I have no indication that *c* is likelier to be true than it is to be false. Slote's argument, therefore, gives no epistemic reason for believing in the external world.

It appears, then, that (D') is capable of withstanding the criticisms brought against it by a number of contemporary philosophers. Although it has often been thought to be relatively weak, the dreaming argument is actually quite strong, and it continues to present a serious challenge to our supposed knowledge of the external world. Of course, it may be that this challenge can be met. But if so, we believe that this will have to involve means considerably different from those discussed here.

We want to thank Laurence Bon Jour for helpful comments on an earlier version of this paper.

1. J. L. Austin, *Sense and Sensibilia* (Oxford: Oxford University Press, 1962), p. 48.
2. *Ibid.*, pp. 48–49.
3. *Ibid.*, p. 49n.
4. Anthony Kenny, *Descartes: A Study of His Philosophy* (New York: Random House, 1968), p. 30.
5. Norman Malcolm, *Dreaming* (London: Routledge & Kegan Paul, 1959), and "Dreaming and Skepticism," *Philosophical Review* 65 (1956):14–37.
6. Kenny, *Descartes*, p. 31.
7. Kenny appears not to hold this view, but Norman Malcolm clearly endorses it. He says, "to a person who is sound asleep, 'dead to the world', things cannot even seem" ("Dreaming and Skepticism," p. 26). On Malcolm's view, it is logically impossible that one have any sensations, thoughts, or feelings during sound sleep. We acknowledge that Malcolm's position, if it were acceptable, would refute premises (1) and (3) of argument (D). Malcolm sees how far one would have to push the present line of attack to make it actually get at the dreaming argument. But his views have been widely discussed and criticized over the past twenty years, and space does not permit a review of the controversy here. A few of the many relevant articles to be consulted are: A. J. Ayer, "Professor Malcolm on Dreams," *Journal of Philosophy* 57 (1960):517–35; Donald Kalish, "Review: *Dreaming* by Norman Malcolm," *ibid.* 57 (1960):308–11; John V. Canfield, "Judgments in Sleep," *Philosophical Review* 70 (1961):224–30; D. F. Pears, "Professor Norman Malcolm: Dreaming," *Mind* 70 (1961):145–63; Hilary Putnam, "Dreaming and 'Depth Grammar'," in *Analytical Philosophy*, ed. by R. Butler (Oxford: Blackwell, 1962), pp. 211–35; Charles Chihara, "What dreams are made on," *Theoria* 31 (1965):145–58; E. M. Curley, "Dreaming and Conceptual Revision," *Australasian Journal of Philosophy* 53 (1975): 119–41.
8. G. E. Moore, *Philosophical Papers* (London: George Allen & Unwin, 1959), p. 248.
9. *Ibid.*, p. 248–49.
10. This is not intended as a definition of "dream." See note 22.
11. Moore, *Philosophical Papers*, p. 250.
12. *Ibid.*, editor's note, p. 251.
13. Of course, many philosophers would hold that knowing implies knowing that one knows. If this thesis is correct (and if Moore's argument is the best one can offer against skepticism), then it would follow that one cannot know that he is not dreaming. In this case the skepticism of (D') would be completely untouched by Moore's argument. But we cannot take up this complex and controversial thesis here.
14. Cf. James W. Cornman and Keith Lehrer, *Philosophical Problems and Arguments: An Introduction* (New York: Macmillan, 1974), pp. 114–18.
15. It should be pointed out that the problem raised here is distinct from one often noted in connection with the frequency interpretation of probability. According to this interpretation, a probability statement makes a factual claim concerning the frequency of a given property relative to a given reference class. But if a probability statement makes a factual claim, then it itself must be evaluated in terms of its probability, yielding another factual claim, which must in turn be evaluated for its probability, and so on. The upshot is that in order to know that a given statement is probably true, one would have to know to be probably true an infinite number of other statements. The problem concerning the reliability of memory is not peculiar to the frequency interpretation of probability statements: it is common to all theories which consider empirical data relevant to probability determinations. C. I. Lewis, in *An Analysis of Knowledge and Valuation*, distin-

guishes the two problems and discusses each of them (La Salle: Open Court, 1946), chs. 10 and 11. While we think Lewis's solution to the memory problem is unsuccessful, it is too complex to consider here. But certain essential elements of his position which are also endorsed by other authors are criticized in various sections of this paper.

16. Bernard Williams, "Descartes, René," in *The Encyclopedia of Philosophy* vol. 2 (New York: Macmillan and The Free Press, 1967), p. 346.
17. Lehrer, *Philosophical Problems and Arguments*, p. 81.
18. *Descartes: A Study of His Philosophy*, p. 25.
19. *Sense and Sensibilia*, p. 11.
20. Gilbert Ryle, *Dilemmas* (Cambridge: Cambridge University Press, 1960), pp. 94–95.
21. Perhaps this difference of intuition can be accounted for by the fact that there seem to be two parts of the concept of genuine money. One part is that it is money produced by a legitimate government, and the other is that it is universally and systematically used as money. Normally, of course, that money which fulfills one of these conditions fulfills the other as well. When we are asked to imagine money that meets one of these conditions and not the other, confusion arises. For while it is clear that these conditions are jointly sufficient, it is not clear that either one is by itself either necessary or sufficient.
22. But it should be pointed out that we also speak of daydreams, and the dreams of an opium eater (some of whose hallucinatory experiences occur while he is awake). So it is not clear that an experience must occur during sleep in order to count as a dream.
23. Further conditions for knowledge may be required in order to avoid Gettier-type counterexamples, but this issue is clearly not relevant to the present discussion.
24. Michael A. Slote, *Reason and Scepticism* (London: George Allen & Unwin, 1970), the quotations that follow are from pp. 66, 65, 67, 68–69; emphasis added. O. K. Bouwsma has also attacked the demon-hypothesis in "Descartes' Evil Genius," *Philosophical Review* 58 (1949):141–51. For criticism of Bouwsma's view, see *Reason and Scepticism*, p. 63; Cornman and Lehrer, *Philosophical Problems and Arguments*, pp. 87–92; and Kenny, *Descartes: A Study of His Philosophy*, pp. 36–37.
- According to Slote, the idea that our experiences arise through chance also violates this principle. For the sake of a more manageable exposition, we do not describe the application of his argument to the hypothesis of chance, which will probably be obvious in any case. Also, we consider only what we regard as his most forceful argument against the demon. Our brief remarks on simplicity, which help set up the problem of this section, are derived from Slote.
25. Slote, *Reason and Scepticism*, pp. 85–86; Slote cites pp. 21f. of Roderick M. Chisholm, *Theory of Knowledge* (Englewood Cliffs, N.J.: Prentice-Hall, 1966). But it is not at all clear that Chisholm intends to *define* epistemic reasonableness. In the passage cited, he seems to be using this notion as a primitive in terms of which other notions are defined.