

tude' is very satisfactory. He would rather put things, he says there, in terms of some conditional statements about what he would say or think right now if or when he imagines things now as happening. But, actually, this latter suggestion is much the poorer and, indeed, Malcolm's choice of the word 'attitude' is quite apt and satisfactory.

11. Peter Unger, "A Defense of Skepticism", *The Philosophical Review*, Vol. LXXX, No. 2, (April 1971), Sections II-IV. In a later issue of this journal, James Cargite replied to the skeptical suggestions in that paper of mine: "In Reply to a Defense of Skepticism", *The Philosophical Review*, Vol. LXXXI, No. 2, (April 1972). Perhaps the present paper may be taken as deepening the debate between myself and this critic in a way that would not be possible in a brief and direct rejoinder on my part.

12. A vivid characterization and illustration of this necessary condition of being certain, involving both himself and his wife, is given by Harry G. Frankfurt in "Philosophical Certainty", *The Philosophical Review*, Vol. LXXI, No. 3 (July 1962), sections IV and V. The main difference here between Frankfurt and myself is that he thinks this complete willingness to risk is not implied by a meaning of 'certain' or even 'absolutely certain' but is only a philosopher's idea; which deserves a new expression, 'philosophically certain'. My own view of course is that no new expression is needed here.

13. Moore, *loc. cit.*

14. Unger, *op. cit.*, sec. IV.

15. For a rather different but quite congenial description of unsettling experiences, see Edward Erwin's "The Confirmation Machine", in *Boston Studies in the Philosophy of Science*, Vol. VIII, Roger C. Buck and Robert S. Cohen, eds.

16. Malcolm suggests this sort of view in his lecture "Memory and The Past", *Knowledge and Certainty*, p. 201. He considers it in a somewhat different context, being most concerned there with the proposition that the earth has existed for no more than five minutes. I will treat such propositions as that in the section following this one. My thoughts on this view owe something to conversation with Michael Slote.

17. Bertrand Russell, *The Analysis of Mind*, New York: The Macmillan Company, 1921, pp. 159-160.

18. I look into this question, or part of it, in the following two companion papers: "The Wages of Scepticism", *American Philosophical Quarterly*, Vol. 10, No. 3 (July 1973) and "Two Types of Scepticism" (Vol. 25, No. 2) *Philosophical Studies*, February, 1974.

19. The main points of this paper were presented as part of a Symposium entitled "Perception, Observation and Skepticism" on March 30, 1973 in Seattle, Washington to the Pacific Division of the American Philosophical Association. My fellow symposiasts were Gilbert Harman and William P. Alston. Switching gears, I would like to thank the many people who have conversed helpfully with me about the ideas of this paper, and to give special thanks to Gilbert Harman, Saul Kripke and Michael Slote for their very great assistance in that regard.

Certainty

G. E. MOORE

I am at present, as you can all see, in a room and not in the open air; I am standing up, and not either sitting or lying down; I have clothes on, and am not absolutely naked; I am speaking in a fairly loud voice, and am not either singing or whispering or keeping

quite silent; I have in my hand some sheets of paper with writing on them; there are a good many other people in the same room in which I am; and there are windows in that wall and a door in this one.

Now I have here made a number of different assertions; and I have made these assertions quite positively, as if there were no doubt whatever that they were true. That is to say, though I did not expressly say, with regard to any of these different things which

I asserted, that it was not only true but also *certain*, yet by asserting them in the way I did, I *implied*, though I did not say, that they were in fact certain—implied, that is, that I myself knew for certain, in each case, that what I asserted to be the case was, at the time when I asserted it, in fact the case. And I do not think that I can be justly accused of dogmatism or over-confidence for having asserted these things positively in the way that I did. In the case of some kinds of assertions, and under some circumstances, a man can be justly accused of dogmatism for asserting something positively. But in the case of assertions such as I made, made under the circumstances under which I made them, the charge would be absurd. On the contrary, I should have been guilty of absurdity if, under the circumstances, I had *not* spoken positively about these things, if I spoke of them at all. Suppose that now, instead of saying “I am inside a building,” I were to say “I *think* I’m inside a building, but perhaps I’m not: It’s not *certain* that I am,” or instead of saying “I have got some clothes on,” I were to say “I think I’ve got some clothes on but it’s just possible that I haven’t.” Would it not sound rather ridiculous for me now, under these circumstances, to say “I *think* I’ve got some clothes on” or even to say “I not only think I have, I know that it is very likely indeed that I have, but I can’t be quite sure”? For some persons, under some circumstances, it might not be at all absurd to express themselves thus doubtfully. Suppose, for instance, there were a blind man, suffering in addition from general anaesthesia, who knew, because he had been told, that his doctors from time to time stripped him naked and then put his clothes on again, although he himself could neither see nor feel the difference: to such a man there might well come an occasion on which he would really be describing correctly the state of affairs by saying that he *thought* he’d got some clothes on, or that he knew that it was very likely he had, but was not quite sure. But for me, now, in full possession of my senses, it would be quite ridiculous to express myself in this way, because the cir-

cumstances are such as to make it quite obvious that I don’t merely think that I have, but know that I have. For me now, it would be absurd to say that I *thought* I wasn’t naked, because by saying this I should imply that I didn’t know that I wasn’t, whereas you can all see that I’m in a position to know that I’m not. But if *now* I am not guilty of dogmatism in asserting positively that I’m not naked, certainly I was not guilty of dogmatism when I asserted it positively in one of those sentences with which I began this lecture. I knew then that I had clothes on, just as I know now that I have.

Now those seven assertions with which I began were obviously, in some respects, not all of quite the same kind. For instance: while the first six were all of them (among other things) assertions about myself, the seventh, namely that there were windows in that wall, and a door in this one, was not about myself at all. And even among those which were about myself there were obvious differences. In the case of two of these—the assertions that I was in a room, and the assertion that there were a good many other people in the same room with me—it can quite naturally be said that each gave a partial answer to the question what sort of *environment* I was in at the time when I made them. And in the case of three others—the assertions that I had clothes on, that I was speaking in a fairly loud voice, and that I had in my hand some sheets of paper—it can also be said, though less naturally, that they each gave a partial answer to the same question. For, if I had clothes on, if I was in a region in which fairly loud sounds were audible, and if I had some sheets of paper in my hand, it follows, in each case that the surroundings of my body were, in at least one respect, different from what they would have been if that particular thing had not been true of me; and the term “environment” is sometimes so used that any true statement from which it follows that the surroundings of my body were different, in any respect, from what they might have been is a statement which gives *some* information, however little, as to the kind of environment

I was in. But though each of these five assertions can thus, in a sense, be said to have given, if true, *some* information as to the nature of the environment at the time when I made it, one of them, the assertion that I was speaking in a fairly loud voice, did not *only* do this: it also, if true, gave some information of a very different kind. For to say that I was speaking in a fairly loud voice was not only to say that there were audible in my neighbourhood fairly loud sounds, and sounds of which it was also true that they were words; it was also to say that some sounds of this sort were *being made by me*—a causal proposition. As for the sixth of the assertions which I made about myself—the assertion that I was standing up—that can hardly be said to have given any information as to the nature of my environment at the time when I made it: it would be naturally described as giving information only as to the posture of my body at the time in question. And as for the two assertions I made which were not about myself at all—the assertions that there were windows in that wall and a door in this one—though they were, in a sense, assertions about my environment, since the two walls about which I made them were, in fact, in my neighbourhood at the time; yet in making them I was not expressly asserting that they were in my neighbourhood (had I been doing so, they would have been assertions about myself) and what I expressly asserted was something which might have been true, even if they had not been in my neighbourhood. In this respect they were unlike my assertion that I was in a room, which could not have been true, unless some walls had been in my neighbourhood. From the proposition that there is a door in that wall it does not follow that that wall is in my neighbourhood; whereas from the proposition that I am in a room, it does follow that a wall is in my neighbourhood.

But in spite of these, and other, differences between those seven or eight different assertions, there are several important respects in which they were all alike.

(1) In the first place: All of those seven or eight different assertions, which I made at

the beginning of this lecture, were alike in this respect, namely, that every one of them was an assertion, which, though it wasn't in fact false, yet *might have been false*. For instance, consider the time at which I asserted that I was standing up. It is certainly true that at that very time I *might* have been sitting down, though in fact I wasn't; and if I *had* been sitting down at that time, then my assertion that I was standing up would have been false. Since, therefore, I might have been sitting down at that very time, it follows that my assertion that I was standing up was an assertion which *might have been false*, though it wasn't. And the same is obviously true of all the other assertions I made. At the time when I said I was in a room, I might have been in the open air; at the time when I said I had clothes on, I might have been naked; and so on, in all the other cases.

But from the fact that a given assertion might have been false, it always follows that the negation or contradictory of the proposition asserted is not a self-contradictory proposition. For to say that a given proposition might have been false is equivalent to saying that its negation or contradictory might have been true; and from the fact that a given proposition might have been true, it always follows that the proposition in question is not self-contradictory, since, if it were, it could not possibly have been true. Accordingly all those things which I asserted at the beginning of this lecture were things of which the *contradictories were not self-contradictory*. If, for instance, when I said "I am standing up" I had said instead "It is not the case that I am standing up," which would have been the contradictory of what I did say, it would have been correct to say "That is not a self-contradictory proposition, though it is a false one"; and the same is true in the case of all the other propositions that I asserted. As a short expression for the long expression "proposition which is not self-contradictory" philosophers have often used the technical term "contingent proposition." Using the term "contingent" in this sense, we can say, then, that one respect in which all those seven propositions which I asserted at

the beginning of this lecture resembled one another was that *they were all of them contingent*.

And before I go on to mention some other respects in which they were all alike, I think I had better now at once say some things about the consequences of this first fact that they were all of them contingent—things which are very relevant to a proper understanding of the use of the word which forms the title of this lecture, the word “Certainty.”

The first thing I want to say about the consequences of the fact that all those propositions were contingent is this: namely, that from the mere fact that they were all of them contingent, it does not follow that they were not all *known to be true*—nay more, it does not follow, in the case of any particular person whatever, that *that* person did not know them to be true. Some philosophers have in fact suggested that no contingent proposition is ever, as a matter of fact, known to be true. And I am not *now* disputing that suggestion, though I do in fact hold it to be false, and intend, in the course of this lecture to dispute it. All that I am asserting *now* is that, even if it is a fact that no contingent proposition is ever known to be true, yet in no case does this *follow* from the mere fact that it is contingent. For instance, that I am now standing up is a contingent proposition; but from the mere fact that it is so, from that fact *alone*, it certainly does not *follow* that I do not know that I am standing up. If it is to be shown—as many philosophers think they can show—that I do *not* know now that I am standing up, some other argument must be brought forward for this contention, over and above the mere fact that this proposition is contingent; for from this fact, by itself, it certainly does not *follow* that I don't know that I'm standing up. I say that this is certain, and I do not know that anyone would dispute it. But if I were asked to defend my assertion, I do not know that I could give any better defence than merely to say that the conjunctive proposition “I know that I am at present standing up, and yet the proposition that I am is contingent” is

certainly not itself self-contradictory, even if it is false. Is it not obvious that if I say “I know that I am at present standing up, although the proposition that I am is contingent,” I am certainly not contradicting myself, even if I *am* saying something which is false?

The second thing I want to say about the consequences of the fact that all those seven propositions were contingent is something which follows from the first: namely that from the fact that they were contingent it does not follow, in the case of any single one among them, that it was *possible* that the proposition in question was false. To take, for instance, again, the proposition that I was then standing up: from the fact that the proposition was contingent, it does not follow that, if I had said “It is possible that it is not the case that I am standing up,” I should have been saying something true. That this is so follows from my former contention that the contingency of the proposition in question does not entail that it was not known to be true, because one, at least, of the ways in which we use expressions of the form “It is possible that *p*” is such that the statement in question cannot be true if the person who makes it knows for certain that *p* is false. We very, very often use expressions of the form “It is possible that *p*” in such a way that by using such an expression we are making an assertion of our own ignorance on a certain point—an assertion namely that we do not *know* that *p* is false. This is certainly one of the very commonest uses of the word “possible”; it is a use in which what it expresses is often expressed instead by the use of the word “may.” For instance, if I were to say “It is possible that Hitler is dead at this moment” this would naturally be understood to mean exactly the same as if I said “Hitler *may* be dead at this moment.” And is it not quite plain that if I did assert that Hitler *may* be dead at this moment part at least of what I was asserting would be that I personally did not know for certain that he was not dead? Consequently if I were to assert now “It is possible that I am not standing up” I should naturally be understood to be asserting that

I do not know for certain that I am. And hence, if I do know for certain that I am, my assertion that it is possible that I'm not would be false. Since therefore from the fact that "I am standing up" is a contingent proposition it does not follow that I do not know that I am, it also does not follow from this fact that it is possible that I am *not*. For if from the contingency of this proposition it did follow that it is possible that I am not standing up, it would also follow that I do not know that I *am* standing up: since from "It is possible that I am not standing up" there follows "I do not know that I am standing up"; and if p entails q , and q entails r , it follows that p entails r . Since, therefore, our p ("the proposition 'I am standing up' is contingent") does not entail our r ("I do not know that I am standing up"), and since our q ("It is possible that I am not standing up") does entail our r , it follows that our p does not entail our q : that is to say, the fact that the proposition "I am standing up" is contingent does not entail the consequence that it is possible that it is false that I am standing up. In no case whatever from the mere fact that a proposition p is contingent does it follow that it is possible that p is false. But this, of course, is not to deny that it may, *as a matter of fact*, be true of every contingent proposition that it is possible that it is false. This *will* be true, if no contingent proposition is ever known to be true. But even if this is so, it still remains true that from the mere fact that a proposition is contingent it never follows that it *may* be false; this remains true because from the mere fact that a proposition is contingent it never follows that it is not known to be true, and never follows, either, in the case of any particular person, that that person does not know it to be true.

In the above paragraph I confined myself to saying that there is at least one common use of expressions of the form "It is possible that p ," such that any person who makes such an assertion is asserting that he personally does not know that p is false; and hence the only conclusion to which I am so far entitled is that the mere fact that a given proposition p is contingent does not entail

the consequence that what is expressed by "it is possible that not- p " will be true, *when "possible" is used in the way in question*. And it may be thought that there is another use of "possible" such that from " p is contingent" there does follow "it is possible that p is false." The fact is that the expression "logically possible" has often been used by philosophers in such a way that many might be tempted to think that it is a mere synonym for "not self-contradictory." That it is not a mere synonym for this can, I think, be seen from the fact that the expression "it is not self-contradictory that I am not standing up" is not English at all, whereas the expression "It is logically possible that I am not standing up" certainly is English, though it may be doubted whether what it expresses is true. If, however, we consider the expression "the proposition that I am not standing up is not self-contradictory" I think it would not be incorrect to say that the words "logically possible" are so used that *in this expression* they could be substituted for "not self-contradictory" without changing the meaning of the whole expression; and that the same is true whatever other proposition you might take instead of the proposition that I am not standing up. If this be so, then it follows that, in the case of any proposition whatever, from the proposition that that proposition is not self-contradictory it will follow that the proposition in question is also logically possible (and *vice versa*); in other words, for any p , " p is not self-contradictory" entails " p is logically possible." But this being so, it is very natural to think that it follows that you can also take a further step and say truly that, for any p , " p is not self-contradictory" entails "It is logically possible that p "; for surely from " p is logically possible" it must follow that "it is logically possible that p ." Certainly it is very natural to think this; but for all that, I think it is a mistake to think so. To think that " p is logically possible" must entail "It is logically possible that p " is certainly a mere mistake which does not do justice to the subtlety of the differences there may be in the way we use language. And I think it is actually a mistake to say that " p is not self-contradictory" entails "It is logically possible that p ."

dictory" entails "It is logically possible that p ," even though it does entail " p is logically possible." Consider the following facts. "It is logically possible that I *should have been* sitting down now" certainly does entail "The proposition that I am sitting down now is not self-contradictory." But if this latter proposition did entail "It is logically possible that I *am* sitting down now" then it would follow that "It is logically possible that I *should have been* sitting down now" entails "It is logically possible that I *am* sitting down now." But does it? Certainly it would be quite unnatural for me, who know that I am standing up, to say the latter, whereas it would be quite natural for me to say the former; and I think perhaps we can go further and say that if I said the latter I should be saying something untrue, whereas if I said the former I should be saying something true; just as if I said "I *might have been* sitting down now," I should be saying something true, whereas if I said "I *may* be sitting down now," I should be saying something false. In short I think that even the expression "It is *logically* possible that so-and-so *is* the case" retains the characteristic which we have seen to belong to one ordinary use of the expression "It is possible that so-and-so *is* the case," namely that it can only be said *with truth* by a person who does not know that the so-and-so in question is *not* the case. If I were to say now "It is logically possible that I am sitting down" I should be implying that I don't know that I'm not, and therefore implying something which, if I do know that I'm not, is false. I think that perhaps philosophers have not always paid sufficient attention to the possibility that from the mere fact that a given proposition, p , is not self-contradictory, it perhaps does not follow that any person whatever can say with truth "It is logically possible that p is true." In the case of a non-self-contradictory proposition such as the proposition that I am at present sitting down, if there be a person, for instance some friend of mine in England, who does not know that this proposition is false, then, in his case, from the *conjunction* of the fact that the proposition is not self-contradictory with

the fact that he does not know it to be false, it does follow that he could say with truth "It is logically possible that Moore is at present sitting down"; but if there be another person, myself for instance, who does know that the proposition is false, it is by no means clear that from the mere fact that the proposition is not self-contradictory—from that fact *alone*—it follows that I can truly say "It is logically possible that I am at present sitting down." From the conjunction of the fact that the proposition is logically possible with the fact that I know it to be false, it does follow that I can truly say "It is logically possible that I *should have been* sitting down at this moment"; but from the fact that I can truly say this, it certainly does not follow that I can *also* truly say "It is logically possible that I *am* sitting down"; and it is certain that in fact the two are incompatible: that, if I can truly say "It is logically possible that I *should have been* sitting down now" then it follows that I *cannot* truly say "It is logically possible that I *am* sitting down now." Perhaps, however, our use of the expression "It is logically possible that so-and-so *is* the case" is not clearly enough fixed to entitle us to say this. What is important is to insist that if "It is logically possible that p is true," is used in such a way that it does follow from " p is not self-contradictory," by *itself*, then from "It is logically possible that \bar{p} is true," it does not follow that p is not known to be false. And if a philosopher does choose to use "It is logically possible that p is true" in such an unnatural way as this, there will be a danger that he will sometimes forget that that is the way in which he has chosen to use it, and will fall into the fallacy of thinking that from "It is logically possible that p is true" there *does* follow " p is not known to be false."

The third thing which I wish to say about the consequences of the fact that those seven assertions with which I began this paper were assertions of contingent propositions, is this: that this fact is quite compatible with it being true that every one of those seven things that I asserted was not only true but *absolutely certain*. That this is so again follows from the fact that the mere contingency of

a given proposition, p , never entails, in the case of any person whatever, that that person does not know p to be true. It follows from this fact, because if any person whatever does at a given time know that a given proposition p is true, then it follows that that person could say with truth at that time "It is absolutely certain that p ." Thus if I do know now that I am standing up, it follows that I can say with truth "It is absolutely certain that I am standing up." Since, therefore, the fact that this proposition is contingent is compatible with its being true that I know that I am standing up, it follows that it must also be compatible with its being true that it is absolutely certain that I am standing up.

I think that possibly some people might be inclined to object to what I have just said on the following ground. I have just said that if a person can ever say with truth, with regard to any particular proposition p , "I know that p is true," it follows that he can also truly say "It is absolutely certain that p is true." But an objector might perhaps say: "I admit that if a person could ever truly say 'I know with absolute certainty that p is true' then it would follow that he could also truly say 'It is absolutely certain that p is true.' But what you said was not 'know with absolute certainty' but 'know'; and surely there must be some difference between 'knowing' and 'knowing with absolute certainty,' since, if there were not, we should never be tempted to use the latter expression. I doubt, therefore, whether a mere 'I know that p ' does entail 'It is absolutely certain that p .'" To this objection I should reply: I do not think that the only possible explanation of the fact that we sometimes say "I know with absolute certainty that so-and-so" and sometimes merely "I know that so-and-so" is that the latter can be properly used to express something which may be true even when what is expressed by the former is not true: I doubt therefore whether "I know that p " does not always entail "I know with absolute certainty that p ." But even if "I know that p " can be sometimes properly used to express something from which "I know with absolute cer-

tainty that p " does *not* follow, it is certainly also sometimes used in such a way that if I don't know with absolute certainty that p , then it follows that I don't know that p . And I have been and shall be concerned only with uses of "know" of the latter kind, i.e. with such that "I know that p " does entail "I know with absolute certainty that p ." And similarly, even if there are proper uses of the word "certain," such that a thing can be "certain" without being "absolutely certain," there are certainly others (or at least one other) such that if a thing is not absolutely certain it cannot be truly said to be certain; and I have been and shall be concerned only with uses of "certain" of this latter kind.

Another comment which might be made upon what I have said is that, even if there is *one* use of "absolutely certain" such that, as I said, it is never logically impossible that a contingent proposition should be absolutely certain, yet there is another use of "absolutely certain" such that this is logically impossible—a sense of "absolutely certain," that is to say, in which only propositions whose contradictories are self-contradictory can be absolutely certain. Propositions whose contradictories are self-contradictory have sometimes been called "necessary truths," sometimes "a priori propositions," sometimes "tautologies"; and it is sometimes held that the sense in which such propositions can be "certain," and therefore also the sense in which they can be "known to be true," must be different from the sense (if any) in which contingent propositions are sometimes "certain," and "known to be true." That this may be so, I do not wish to deny. So far as I can see, it may be the case that, if I say, "I know that" or "It is certain that" "it is not the case that there are any triangular figures which are not trilateral," or "I know that" or "It is certain that" "it is not the case that there are any human beings who are daughters and yet are not female," I am using "know that" and "it is certain that" in a different sense from that in which I use them if I say "I know that" or "It is certain that" "I have some clothes on"; and it may be the case that only necessary truths

can be known or be certain in the former sense. Accordingly, my statements that from the fact that a given proposition, p , is contingent it does not follow that p is not known and is not certain, should be understood to mean only that there is at least one sense in which "known" and "certain" can be properly used, such that this does not follow; just as all that I asserted positively before about the phrase "It is possible that" was that there is at least one sense in which this phrase can be properly used, such that " p is contingent" does not entail "It is possible that p is false."

Finally, there is one slightly puzzling point about our use of the phrases "it is possible that" and "it is certain that," which might lead some people to suspect that some of the things I have been saying about the consequences which follow from the fact that a given proposition is contingent are false, and which therefore I think I had better try to clear up at once.

There are four main types of expression in which the word "certain" is commonly used. We may say "I feel certain that . . .," or we may say "I am certain that . . .," or we may say "I know for certain that . . .," or finally we may say "It is certain that . . ." And if we compare the first of these expressions with the two last, it is, of course, very obvious, and has been pointed out again and again, that whereas "I feel certain that p " may quite well be true in a case in which p is not true—in other words that from the mere fact that I feel certain that so-and-so is the case it never follows that so-and-so is in fact the case—there is at least one common use of "I know for certain that p " and "It is certain that p " such that these things can't be true unless p is true. This difference may be brought out by the fact that, e.g., "I felt certain that he would come, but in fact he didn't" is quite clearly not self-contradictory; it is quite clearly logically possible that I should have felt certain that he would come and that yet he didn't; while, on the other hand, "I knew for certain that he would come, but he didn't" or "It was certain that he would come but he didn't" are, for at least one common use of those phrases, self-

contradictory: the fact that he didn't come *proves* that I didn't *know* he would come, and that it wasn't certain that he would, whereas it does not prove that I didn't *feel* certain that he would. In other words, "I feel certain that p " does not *entail* that p is true (although by saying that I feel certain that p , I do *imply* that p is true), but "I know that p " and "It is certain that p " do entail that p is true; they can't be true, unless it is. As for the fourth expression "I am certain that . . ." or "I am quite sure that . . ." (it is perhaps worth noting that in the expressions "I feel certain that . . ." and "I am certain that . . ." the word "sure" or the words "quite sure" can be substituted for the word "certain" without change of meaning, whereas in the expressions "I know for certain that . . ." or "it is certain that . . ." this is not the case) these expressions are, I think, particularly liable to give rise to fallacious reasoning in philosophical discussions about certainty, because, so far as I can see, they are sometimes used to mean the same as "I feel certain that . . ." and sometimes, on the contrary, to mean the same as "I know for certain that . . ." For instance, the expression "I was quite sure that he would come, but yet he didn't" *can*, it seems to me, be naturally used in such a way that it is not self-contradictory—which can only be the case if it is in that case merely another way of saying "I felt quite sure that he would come . . ."; but if on the other hand a philosopher were to say to me now (as many would say) "You can't be quite sure that you are standing up," he would certainly not be asserting that I can't *feel* certain that I am—a thing which he would not at all wish to dispute—and he certainly would be asserting that, even if I do feel certain that I am, I don't or can't *know for certain* that I am.

There is, therefore, a clear difference in meaning between "I feel certain that . . ." on the one hand, and "I know for certain that . . ." or "It is certain that . . ." on the other. But the point with which I am at present concerned is whether there is not also a difference of importance between each of these expressions "I feel certain that . . .," "I am

certain that . . .," and "I know for certain that . . ." on the one hand, and "It is certain that . . ." on the other. The first three expressions are obviously, in spite of the important difference I have just pointed out between the first and the last of them, alike in one important respect—a respect which may be expressed by saying that their meaning is relative to the person who uses them. They are alike in this respect, because they all contain the word "I." In the case of every sentence which contains this word, its meaning obviously depends on who it is that says that sentence; if I say "I am hot," what I assert by saying this is obviously something different from what any other person would be asserting by saying exactly the same words; and it is obvious that what I assert by saying so may quite well be true even though what another person asserts by saying exactly the same words at exactly the same time is false. "I am hot" said by me at a given time, does not contradict "I am not hot" said by you at exactly the same time: both may perfectly well be true. And in the same way, if I say "I feel certain that there are windows in that wall" or "I know for certain that there are windows," I, by saying this, am making an assertion different from, and logically independent of, what another person would be asserting by saying exactly the same words at the same time: from the fact that I feel certain of or know for certain a given thing it *never* follows, in the case of any other person whatever, that he feels certain of or knows the thing in question, nor from the fact that he does does it ever follow that I do. But if we consider, by contrast, the expression "It is certain that there are windows in that wall," it looks, at first sight, as if the meaning of this expression was *not* relative to the person who says it: as if it were a quite impersonal statement and should mean the same whoever says it, provided it is said at the same time and provided the wall referred to by the words "that wall" is the same. It is, indeed, obvious, I think, that a thing can't be certain, unless it is *known*: this is one obvious point that distinguishes the use of the word "certain" from that of the word

"true"; a thing that nobody knows may quite well be true, but cannot possibly be certain. We can, then, say that it is a necessary condition for the truth of "It is certain that *p*" that somebody should know that *p* is true. But the meaning of "Somebody knows that *p* is true" is certainly not relative to the person who says it: it is as completely impersonal as "The sun is larger than the moon," and if two people say it at the same time, then, if the one by saying it is saying something true, so must the other be. If, therefore, "It is certain that *p*" meant merely "Somebody knows that *p* is true," then the meaning of "It is certain that *p*" would *not* be relative to the person who says it, and there would then be an important difference between it, on the one hand, and "I feel certain that *p*" or "I know for certain that *p*" on the other, since the meaning of these two is relative to the person who says them. But though "Somebody knows that *p* is true" is a necessary condition for the truth of "It is certain that *p*," it can be easily seen that it is *not* a sufficient condition; for if it were, it would follow that in any case in which somebody did know that *p* was true, it would always be false for anybody to say "It is not certain that *p*." But in fact it is quite evident that if I say now "It is not certain that Hitler is still alive," I am not thereby committing myself to the statement that nobody knows that Hitler is still alive: my statement is quite consistent with its being true that Hitler is still alive, and that he himself and other persons know that he is so. The fact is, then, that all that follows from "Somebody knows that *p* is true" is that *somebody* could say with truth "It is certain that *p*": it does not follow that more than one person could; nor does it follow that there are not some who could say with truth "It is *not* certain that *p*." Two different people, who say, at the same time about the same proposition, *p*, the one "It is certain that *p* is true," the other "It is not certain that *p* is true," may both be saying what is true and not contradicting one another. It follows, therefore, that, in spite of appearances, the meaning of "It is certain that *p*" is relative to the person who says it. And this,

I think, is because, as I have implied above, if anybody asserts "It is certain that p " part of what he is asserting is that he himself knows that p is true; so that, even if many other people do know that p is true, yet his assertion will be false, if he himself does not know it. If, on the other hand, a person asserts "It is *not* certain that p " his assertion will not necessarily be true merely because he personally does not know that p is true, though it will necessarily be false if he personally does know that p is true. If I say "It is certain that p ," that I should know that p is true is both a necessary and sufficient condition for the truth of my assertion. But if I say "It is *not* certain that p ," then that I should *not* know that p is true, though it is a necessary, is not a sufficient condition for the truth of my assertion. And similarly the expression "It is possible that p is true" is, though it looks as if it were impersonal, really an expression whose meaning is relative to the person who uses it. If I say it, that I should not know that p is false, is a necessary, though not a sufficient, condition for the truth of my assertion; and hence if two people say it at the same time about the same proposition it is perfectly possible that what the one asserts should be true, and what the other asserts false: since, if one of the two knows that p is false, his assertion will necessarily be false; whereas, if the other does not know that p is false, his assertion may be, though it will not necessarily be, true. On the other hand, if it were right to use the expression "It is *logically* possible that p " as equivalent to " p is not self-contradictory," then the meaning of "It is *logically* possible that p " would *not* be relative to the person who says it.

To sum up this digression. What I have said about the consequences of the fact that all those seven propositions with which I opened this lecture were contingent, is firstly (1) that this fact does *not* entail the consequence that I did not, when I made them, know them to be true; (2) that it does *not* entail the consequence that I could then have said with truth about any of them "It is possible that this is false"; and (3) that it does

not entail the consequence that I could then have said with truth about any of them "It is not absolutely certain that this is true." It follows that by asserting that those seven propositions were contingent, I have not committed myself to the view that they were not known to be true or that it was not absolutely certain they were. But on the other hand, even if I am right in saying that these consequences do *not* follow from the mere fact that they were contingent, it, of course, does not follow from this that I *did* know them to be true, when I asserted them, or that they were absolutely certain. The questions whether, when I first said that I was standing up, I did know that I was, and whether, therefore, it was absolutely certain that I was, still remain completely open.

(2) A second respect, in addition to the fact that they were all of them contingent, in which all those seven propositions resembled one another, was this: In the case of every one of them part at least of what I was asserting, in asserting it, was something from which nothing whatever about the state or condition of my own mind followed—something from which no psychological proposition whatever about myself followed. Every one of them asserted something which might have been true, no matter what the condition of my mind had been either at that moment or in the past. For instance, that I was then inside a room is something which might have been true, even if at the time I had been asleep and in a dreamless sleep, and no matter what my character or disposition or mental abilities might have been: from that fact alone no psychological proposition whatever about myself followed. And the same is true of part at least of what I asserted in each of the other six propositions. I am going to refer to this common feature of all those seven propositions, by saying that they were all of them propositions which implied the existence of *an external world*—that is to say, of a world *external to my mind*. These phrases "external world" and "external to my mind" have often been used in philosophy; and I think that the way in which I am

now proposing to use them is in harmony with the way in which they generally (though not always) have been used. It is indeed not obvious that my assertion that I was standing up implied the existence of anything external to *my body*; but it has generally been clear that those who spoke of a world *external* to any given *individual*, meant by that a world external to that individual's *mind*, and that they were using the expression "external to a mind" in some metaphorical sense such that my body *must* be external to my mind. Accordingly a proposition which implies the existence of my body does, for that reason alone, with this use of terminology, imply the existence of a world *external to my mind*; and I think that the reason why it is said to do so is because from the existence of my body at a given time nothing whatever logically follows as to the state or condition of my mind at that time. I think, therefore, that I am not saying anything that will be misleading to those familiar with philosophical terminology, if I say, for the reason given, that each of those seven assertions implied the existence of something external to my mind; and that hence, if I did know any one of them to be true, when I asserted it, the existence of an external world was at that time absolutely certain. If, on the other hand, as some philosophers have maintained, the existence of an external world is never absolutely certain, then it follows that I cannot have known any one of these seven propositions to be true.

(3) A third characteristic which was common to all those seven propositions was one which I am going to express by saying that I had for each of them, at the time when I made it, *the evidence of my senses*. I do not mean by this that the evidence of my senses was the *only* evidence I had for them: I do not think it was. What I mean is that, at the time when I made each, I was seeing or hearing or feeling things (or, if that will make my meaning clearer, "having visual, auditory, tactile or organic sensations"), or a combination of these, such that to see or hear or feel those things *was* to have evidence (not necessarily *conclusive* evidence)

for part at least of what I asserted when I asserted the proposition in question. In other words, in all seven cases, what I said was at least partly *based on* "the then present evidence of my senses."

(4) Fourth and finally, I think that all those seven assertions shared in common the following characteristic. Consider the class of all propositions which resemble them in the second respect I mentioned, namely, that they imply the existence of something external to the mind of the person who makes them. It has been and still is held by many philosophers that no proposition which has this peculiarity is ever known to be true—is ever quite certain. And what I think is true of those seven propositions with which I began this lecture is this: namely, that, if I did not know them to be true when I made them, then those philosophers are right. That is to say, if those propositions were not certain, then nothing of the kind is ever certain: if *they* were not certain, then no proposition which implies the existence of anything external to the mind of the person who makes it is ever certain. Take any one of the seven you like: the case for saying that I *knew* that one to be true when I made it is as strong as the case ever is for saying of any proposition which implies the existence of something external to the mind of the person who makes it, that *that* person knows it to be true.

This, it will be seen, is not a matter of logic. Obviously it is logically possible, for instance, that it should have been false then that I knew I was standing up and yet should be true now that I know I am standing up. And similarly in the other cases. But though this is logically possible—though the proposition "I know that I am standing up now, but I did not know then that I was" is certainly not self-contradictory—yet it seems to me that it is certainly false. If I didn't know then that I was standing up, then certainly I know nothing of the sort now, and never have known anything of the sort; and, not only so, but nobody else ever has. And similarly, conversely (though this also is not a matter of logic), if I did know then that I

was standing up then I certainly also know that I am standing up now, and have in the past constantly known things of the sort; and, not only so, but millions of other people have constantly known things of the sort: we all of us constantly do. In other words, those seven propositions of mine seem to be as good test-cases as could have been chosen (as good as, but also no better than thousands of others) for deciding between what seems to me to be the only real (though far from the only logically possible) alternatives—namely the alternative that none of us ever knows for certain of the existence of anything external to his own mind, and the alternative that all of us—millions of us—constantly do. And it was because they seemed to me to be as good test-cases as could be chosen for deciding this that I chose them.

But can we decide between these two alternatives?

I feel that the discussion of this question is frightfully difficult; and I feel sure that better and more decisive things could be said about it than I shall be able to say. All that I can do is to discuss, and that very inadequately, just one of the types of argument which have sometimes been alleged to show that nobody ever has known for certain anything about a world external to his mind.

Suppose I say now: "I know for certain that I am standing up; it is absolutely certain that I am; there is not the smallest chance that I am not." Many philosophers would say: "You are wrong; you do not know that you are standing up; it is *not* absolutely certain that you are; there is *some* chance, though perhaps only a very small one, that you are not." And one argument which has been used as an argument in favour of saying this, is an argument in the course of which the philosopher who used it would assert: "You do not know for certain that you are not dreaming; it is not absolutely certain that you are not; there is *some* chance, though perhaps only a very small one, that you are." And from this, that I do not know for certain that I am not dreaming, it is supposed to follow that I do not know for certain that I am standing up. It is argued: If it

is not certain that you are not dreaming, then it is not certain that you are standing up. And that *if* I don't know that I'm not dreaming, I also don't know that I'm not sitting down, I don't feel at all inclined to dispute. From the hypothesis that I am dreaming, it would, I think, certainly follow that I don't *know* that I am standing up; though I have never seen the matter argued, and though it is not at all clear to me how it is to be proved that it would follow. But, on the other hand, from the hypothesis that I am dreaming, it certainly would not follow that I am *not* standing up; for it is certainly logically possible that a man should be fast asleep and dreaming, while he is standing up and not lying down. It is therefore logically possible that I should both be standing up and also at the same time dreaming that I am; just as the story, about a well-known Duke of Devonshire, that he once dreamt that he was speaking in the House of Lords and, when he woke up, found that he *was* speaking in the House of Lords, is certainly logically possible. And if, as is commonly assumed, when I am dreaming that I am standing up it may also be correct to say that I am *thinking* that I am standing up, then it follows that the hypothesis that I am now dreaming is quite consistent with the hypothesis that I am both thinking that I am standing up and also actually standing up. And hence, if, as seems to me to be certainly the case and as this argument assumes, from the hypothesis that I am now dreaming it *would* follow that I don't know that I am standing up, there follows a point which is of great importance with regard to our use of the word "knowledge," and therefore also of the word "certainty"—a point which has been made quite conclusively more than once by Russell, namely that from the conjunction of the two facts that a man thinks that a given proposition *p* is true, and that *p* is in fact true, it does *not* follow that the man in question *knows* that *p* is true: in order that I may be justified in saying that I know that I am standing up, something more is required than the mere conjunction of the two facts that I both think I am and actually am—as

Russell has expressed it, true belief is not identical with knowledge; and I think we may further add that even from the conjunction of the two facts that I feel certain that I am and that I actually am it would not follow that I know that I am, nor therefore that it is certain that I am. As regards the argument drawn from the fact that a man who dreams that he is standing up and happens at the moment actually to be standing up will nevertheless not *know* that he is standing up, it should indeed be noted that from the fact that a man is dreaming that he is standing up, it certainly does not *follow* that he *thinks* he is standing up; since it does sometimes happen in a dream that we *think* that it is a dream, and a man who thought this certainly might, although he was dreaming that he was standing up, yet *think* that he was not, although he could not *know* that he was not. It is not therefore the case, as might be hastily assumed, that, if I dream that I am standing up at a time when I am in fact lying down, I am necessarily *deceived*: I should be deceived only if I thought I was standing when I wasn't; and I may dream that I am, without thinking that I am. It certainly does, however, often happen that we do dream that so-and-so is the case, without at the time thinking that we are only dreaming; and in such cases, I think we may perhaps be said to *think* that what we dream is the case *is* the case, and to be deceived if it is not the case; and therefore also, in such cases, if what we dream to be the case happens also to *be* the case, we may be said to be thinking truly that it is the case, although we certainly do not *know* that it is.

I agree, therefore, with that part of this argument which asserts that if I don't know now that I'm not dreaming, it follows that I don't *know* that I am standing up, even if I both actually am and think that I am. But this first part of the argument is a consideration which cuts both ways. For, if it is true, it follows that it is also true that if I *do* know that I am standing up, then I do know that I am not dreaming. I can therefore just as well argue: since I do know that I'm standing up, it follows that I do know that I'm not

dreaming; as my opponent can argue: since you don't know that you're not dreaming, it follows that you don't know that you're not standing up. The one argument is just as good as the other, unless my opponent can give better reasons for asserting that I don't know that I'm not dreaming, than I can give for asserting that I do know that I am standing up.

What reasons can be given for saying that I don't know for certain that I'm not at this moment dreaming?

I do not think that I have ever seen clearly stated any argument which is supposed to show this. But I am going to try to state, as clearly as I can, the premisses and the reasonings from them, which I think have led so many philosophers to suppose that I really cannot now know for certain that I am not dreaming.

I said, you may remember, in talking of the seven assertions with which I opened this lecture, that I had "the evidence of my senses" for them, though I also said that I didn't think this was the only evidence I had for them, nor that this by itself was necessarily conclusive evidence. Now if I had *then* "the evidence of my senses" in favour of the proposition that I was standing up, I certainly have *now* the evidence of my senses in favour of the proposition that I *am* standing up, even though this may not be all the evidence that I have, and may not be conclusive. But have I, in fact, the evidence of my senses *at all* in favour of this proposition? One thing seems to me to be quite clear about our use of this phrase, namely, that, if a man at a given time is only dreaming that he is standing up, then it follows that he has *not* at that time the evidence of his senses in favour of that proposition: to say "Jones last night was *only* dreaming that he was standing up, and yet all the time he had the evidence of his senses that he was" is to say something self-contradictory. But those philosophers who say it is possible that I am now dreaming, certainly mean to say also that it is possible that I am *only dreaming* that I am standing up; and this view, we now see, entails that it is possible that I have *not* the

evidence of my senses that I am. If, therefore, they are right, it follows that it is not certain even that I have the evidence of my senses that I am; it follows that it is not certain that I have *the evidence of my senses* for anything at all. If, therefore, I were to say now, that I certainly have the evidence of my senses in favour of the proposition that I am standing up, even if it's not certain that I am standing up, I should be begging the very question now at issue. For if it is not certain that I am not dreaming, it is not certain that I even have the evidence of my senses that I am standing up.

But, now, even if it is not certain that I have at this moment the evidence of my senses for anything at all, it is quite certain that I *either* have the evidence of my senses that I am standing up *or* have an experience which is *very like* having the evidence of my senses that I am standing up. *If* I am dreaming, this experience consists in having dream-images which are at least very like the sensations I should be having if I were awake and had the sensations, the having of which would constitute "having the evidence of my senses" that I am standing up. Let us use the expression "sensory experience," in such a way that this experience which I certainly am having will be a "sensory experience," whether or not it merely consists in the having of dream-images. If we use the expression "sensory experience" in this way, we can say, I think, that, if it is not certain that I am not dreaming now, then it is not certain that *all* the sensory experiences I am now having are not mere dream-images.

What then are the premisses and the reasonings which would lead so many philosophers to think that all the sensory experiences I am having now *may* be mere dream-images—that I do not know for certain that they are not?

So far as I can see, one premiss which they would certainly use would be 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." This seems a very harmless premiss, and I am quite willing to

admit that it is true. But I think 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. All the philosophers I have ever met or heard of certainly did know that dreams have occurred: we all know that dreams *have* occurred. 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. By using this premiss he implies that he himself knows that dreams have occurred; while, if his conclusion is true, it follows that he himself does not know that he is not dreaming, and therefore does not know that he is not only dreaming that dreams have occurred.

However, I admit that the premiss is true. Let us now try to see by what sort of reasoning it might be thought that we could get from it to the conclusion.

I do not see how we can get forward in that direction at all, unless we first take the following huge step, unless we say namely: since there have been dream-images similar in important respects to some of the sensory experiences I am now having, it is logically possible, therefore, that there should be dream-images *exactly like all* the sensory experiences I am now having, and logically

possible, therefore, that all the sensory experiences I am now having *are* mere dream-images. And it might be thought that the validity of this step could be supported to some extent by appeal to matters of fact, though only, of course, at the cost of the same sort of inconsistency which I have just pointed out. It might be said, for instance, that some people have had dream-images which were *exactly like* sensory experiences which they had when they were awake, and that therefore it must be logically possible to have a dream-image exactly like a sensory experience which is not a dream-image. And then it may be said: If it is logically possible for some dream-images to be exactly like sensory experiences which are not dream-images, surely it must be logically possible for *all* the dream-images occurring in a dream at a given time to be exactly like sensory experiences which are not dream-images, and logically possible also for all the sensory experiences which a man has at a given time when he is awake to be exactly like all the dream-images which he himself or another man had in a dream at another time.

Now I cannot see my way to deny that it is logically possible that all the sensory experiences I am having now should be mere dream-images. And if this is logically possible, and if further the sensory experiences I am having now were the only experiences I am having, I do not see how I could possibly know for certain that I am not dreaming.

But the conjunction of my memories of the immediate past with these sensory experiences *may* be sufficient to enable me to

know that I am not dreaming. I say it *may* be. But what if our sceptical philosopher says: It is *not* sufficient; and offers as an argument to prove that it is not, this: 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. 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." And I don't think I've ever seen an argument expressly directed to show that it is not.

One final point should be made clear. It is certainly logically possible that I *should have* been dreaming now; I *might* have been dreaming now; and therefore the proposition that I *am* dreaming now is not self-contradictory. But what I am in doubt of is whether it is logically possible that I should *both* be having all the sensory experiences and the memories that I have and *yet* be dreaming. The conjunction of the proposition that I have these sense experiences and memories with the proposition that I am dreaming does seem to me to be very likely self-contradictory.