

*The Pseudo-Problem of Skepticism**

Brendan Sweetman

Since Descartes, philosophers have taken the problem of skepticism very seriously. So seriously, in fact, that they have devoted vast amounts of time and energy generating countless books, articles and discussions in an attempt to solve the problem. However, these attempts have been entirely futile. No universally agreed upon solution has emerged from over three hundred years of intense discussion; in fact, I suppose it is true to say that these efforts have not even succeeded in producing a solution (or even a hint of a solution) that has come to be accepted by even a few philosophers. What we can say for certain is that the current state of play on what Kant called “the scandal of philosophy” is as follows: many excellent and serious-minded philosophers still take the problem of skepticism just as seriously as Descartes and Kant appeared to take it; no solution to the problem has yet emerged; and many philosophers are genuinely worried by the perceived consequences of this failure. This latter point should not be taken lightly.

The problem of skepticism has now so thoroughly pervaded the practice of modern epistemology that any proposed philosophical account of the nature of human knowledge must be accompanied by the inevitable qualification that the account proposed *might* be true *if only* the obstacle created by the problem of skepticism could be overcome. In short, no philosophical account of knowledge can be accepted as true, such reasoning goes, until we have shown first that knowledge is possible. This has been one approach. And in the twentieth century this approach has had the unfortunate consequence of moving us several steps nearer to relativism about knowledge.

Another approach, exemplified by Barry Stroud for example, is to suggest that the truth of the main principle at issue in the problem of skepticism—the principle that there is some essential connection between our beliefs and the

* This is a revised version of an article which originally appeared in *Philosophical Inquiry*, Vol. XVI (1994), pp. 32-43.

way the world really is—should be given up unless we can prove that skepticism is not a logical possibility.¹ Stroud obviously takes the problem of skepticism seriously enough to suggest that we should be prepared to suspend the reliability of our everyday beliefs if no solution is forthcoming. And yet another philosopher, Peter Unger, has gone even further and embraced a wholesale skepticism.² Yet it is true to say that most philosophers, however, regard the problem of skepticism as a wasteful academic exercise and dismiss it out of hand. Some are even embarrassed to teach the problem in their philosophy courses to fresh, unsuspecting, and philosophically innocent minds.

I want to suggest in these brief reflections that the problem of skepticism is a *pseudo-problem*. By “pseudo-problem” I mean that it is not a real problem, and therefore not one that we should take seriously, and expend much time and energy trying to solve, and worry about the consequences if we fail to solve it. Various continental philosophers have also suggested that the problem of skepticism is a pseudo-problem (although they have not used this term). Among them, Martin Heidegger and Gabriel Marcel spring to mind. Although I think these philosophers are basically right in their defense of this particular claim, I shall not adopt or discuss their approach here. My approach shall be rather to discuss the problem on its own terms, from within as it were, and attempt to show that there are some features about the *formulation* of the problem of skepticism itself which enable us to conclude that the most rational judgement for a philosopher to make with regard to the problem is that it is a pseudo-problem. In what follows I assume a basic familiarity with the problem of skepticism, which is perhaps the most well known philosophical problem of all.

Descartes adopted three arguments to motivate his program of methodic doubt: (i) the argument from illusion; (ii) the dream argument; and (iii) the evil genius argument.³ The third step in his program of methodic doubt, the evil genius argument, is necessary because the first two steps do not provide sufficient warrant for him to doubt *all* his beliefs. It will be recalled that the truths of arithmetic and geometry, and the basics of corporeal nature (extension, shape, size, number, etc.) escape the first two stages of Cartesian doubt. So in order to doubt *all* of his beliefs, it is necessary for Descartes to introduce an evil genius who might be deceiving him in *all* his beliefs. As he puts it in the *Meditations*, “I shall then suppose, not that God who is supremely good and the fountain of

¹ See Barry Stroud, *The Significance of Philosophical Skepticism* (New York: O.U.P., 1984).

² See Peter Unger, *Ignorance: A Case for Skepticism* (New York: O.U.P., 1975).

³ See René Descartes, *Meditations in The Philosophical Works of Descartes*, trans. E. S. Haldane and G. R. T. Ross (London: Cambridge University Press, 1979), pp. 144-149.

truth, but some evil genius not less powerful than deceitful, has employed his whole energies in deceiving me.”⁴ The evil genius is introduced by Descartes as a *logical possibility*. That is to say, such a being *could* exist, and *might* be deceiving him in all his beliefs. However implausible this may be, there is no *contradiction* involved in asserting the existence of such a being. Of course, Descartes does not for a moment think that there really *is* an evil genius. Rather, the evil genius is for Descartes a convenient device to enable him to generate universal doubt before he moves on to his *solution* to the problem of skepticism. With this universal skepticism in place, Descartes’s task then is to illustrate that the evil genius is not in fact a logical possibility at all.

Of course, it is the move from step two to step three which has made many people suspicious of the problem of skepticism. It is particularly *this move* in the argument for skepticism that many find hardest to take seriously (and not just non-philosophers and students, but even most philosophers). Suspicions are raised because there seems to be no good reason to adopt step three at all; it looks as if Descartes is creating a problem where one does not exist. It is true that steps one and two give us pause to think twice about the truth or falsity of some of our beliefs (at least in certain circumstances), and they do seem to lend *some* support (however slight) to the claim that there may be reasons to doubt the principle upon which all our knowledge is based, i.e., the principle that there is an essential connection between our beliefs and the way the world really is. But it is precisely in the move to the third step that the *plausibility* of the overall argument seems to break down. I want to pursue this issue further, for I think that it contains the key to the whole business. However, I want to shift the discussion away from the evil genius argument itself, to its modern-day equivalent, the brain-in-a-vat argument. It seems to me that consideration of this modern-day re-formulation of the evil genius argument will enable us to gain an insight into the problem of skepticism which discussions of the evil genius argument do not yield.

Before proceeding any further, I wish to modify the standard brain-in-a-vat argument just a little, but in a way which will, I think, present the issues raised by the problem of skepticism in a clearer, sharper way. The standard brain-in-a-vat argument poses the question: how do I know that I am not simply a brain-in-a-vat being fed all the experiences, beliefs, etc., that I currently have? In order to accept the principle that our beliefs and experiences reflect the way the world really is, the argument goes, we would have to eliminate this possibility. Now this argument is the modern-day equivalent of Descartes’s evil genius argument, for like that argument it places us in a position of universal doubt. It will be

⁴ *Ibid.*, p. 148.

recalled that Descartes's first move in response to the evil genius argument was the realization that he was at least *doubting*—that no matter how much he doubted or was deceived, *someone* must be doubting, and being deceived, etc. From this he concluded that at least he *could know* for certain that *he existed* (even if he could not yet know his nature).

Let me now turn to my modification of the brain-in-the-vat story. For this we have to move forward several centuries into the future. Suppose at this time there is a graduate student in philosophy, Malcolm, who is writing his doctoral dissertation on the problem of skepticism; more particularly, on the brain-in-a-vat version of the problem. He is near the final stages of his work and is deeply involved in the issues. Malcolm is also quite an unscrupulous fellow for he has a part-time job working for an even more unscrupulous scientist called Dr Frankenstein. Dr Frankenstein is conducting research into the workings of the human brain. More specifically he is working on the theory that the human brain can be removed from the skull of the human being and can be kept alive in a vat. Further, he has perfected a technique for feeding the brain, by means of computer programs, various experiences, beliefs, etc., when it is in the vat which are exactly identical to those the brain would have if it were functioning normally in the body of a human being.

But this Frankenstein is an extraordinarily clever fellow and he has perfected another powerful tool to assist in his research. He has developed a special logic program, Logic II, which can be rigged up, in conjunction with the experience programs, to the brain-in-the-vat. This logic program enables Frankenstein to govern totally the rationality of the brain. He is now in a position where he can choose to either feed the brain-in-the-vat various experiences, etc., and allow the brain's own logical system (call it Logic I)⁵ to respond to the experiences, or he can feed the brain various experiences, etc., and program Logic II to operate on them. In this second instance, Frankenstein is in complete control of the rationality of the brain, whereas in the first instance, the brain is allowed to think and reason independently about the various experiences being fed to it. (The traditional brain-in-a-vat scenario seems to presuppose this rational independence for the brain.) Logic II is a new system of logic, remarkably different from ours. One of its most interesting features is that it enables the brain to validly arrive at what it thinks are true conclusions from premises that are, in fact, false. The conclusions present themselves to the brain-in-the-vat with the same power, sense of certainty, etc. as true conclusions would which were arrived at in ordinary experience.

⁵ Nothing in my discussion in this paper presupposes any particular solution to the body/mind problem. I talk mainly about the "brain" rather than the "mind" in a vat simply for the sake of convenience.

Needless to say such research has been outlawed. Frankenstein is operating a totally underground, clandestine operation. With Malcolm's help, he has had to kidnap people and steal their brains in order to further his research. Malcolm, of course, was drawn to this research initially because of his dissertation topic. He was not in fact alone. For another graduate student in his Department, John, was also working on the same problem, and also took a part-time job in the laboratory as an assistant. However, only the previous week, Frankenstein had kidnapped John and stolen his brain, which he has since rigged up in a vat. He had several reasons for doing this. Firstly, he needed the brain of a very intelligent person. A philosopher seemed ideal. Secondly, he knew that Malcolm was working on the brain-in-a-vat problem and he thought the kidnapping of John might be an interesting case for Malcolm to study, and might help him in his reflections on the topic of his dissertation. Malcolm was none too happy about John being treated in this way, but being rather unscrupulous he went along with it, albeit reluctantly. Although, in this instance he refused to personally assist in the kidnapping.

Now when driving home from the laboratory one night, Malcolm is suddenly gripped by a terrible problem. He wonders if *he* too has been kidnapped and his brain placed in a vat, and if he is being fed all of his present experiences. He knows that Frankenstein has the facilities, as well as both the skill and the technical knowledge, to very easily carry out this action. After all, he has seen him do it (has even assisted him) to many others. Malcolm thought about the problem further and realized that there was another possibility. Frankenstein might not have kidnapped him (he might not even have kidnapped John, for Malcolm had not assisted with that kidnapping). He might have set up this whole situation knowing that Malcolm would eventually turn up this problem for himself, and hoping that it would assist Malcolm in some way with his research. Malcolm's problem is now this: he has to try to discover whether or not he has been reduced from a fully-fledged human being to a brain-in-a-vat.

The first point to notice about Malcolm's plight is that he has *excellent reasons* for believing that he *might* be a brain-in-a-vat. Therefore, the problem *really matters* to him. But the real issue is: how is he going to *determine* whether or not he is a brain-in-a-vat? And it is precisely when we pursue this question that we discover that our problem of skepticism is a pseudo-problem. Like Descartes and the evil genius argument, let us grant that Malcolm can initially work out that he is at least a thinking thing, for much the same reason as Descartes's. *Someone* must be doing the thinking, doubting, wondering, etc. Malcolm can be sure at least that he is thinking, yet he cannot be sure whether he is thinking in his body, or in a vat. He cannot be sure if his present experiences are really occurring in the world of everyday reality, or if they are being fed to his brain-

in-the-vat by the nefarious Frankenstein. What are his options for solving the dilemma?

Obviously Malcolm has to consider the various possibilities logically and see if he can eliminate any or all of them. The first possibility Malcolm considers is that he might be in a vat and rigged to Logic II. However, if this is the case then he cannot even trust his inferences, and the problem will be *impossible* to solve. He will never be able to decide if any solution arrived at is true or only apparently true (but really false). He knows that, if he is rigged to Logic II, false conclusions can present themselves to him with a power, sense of certainty, etc., indistinguishable to that which accompanies true conclusions in the world of ordinary experience. So his first problem is how to decide if his brain is operating according to Logic I (the brain's own logical system) or Logic II. If he could establish that his thinking is regulated by Logic II, then obviously he will be unable to solve his central problem of whether he is a brain-in-a-vat. If, however, he can establish that his thinking is regulated by Logic I, he *still* has the problem of whether or not he is a brain-in-a-vat.

Is there any way for Malcolm to decide whether his thinking is regulated by Logic II or Logic I? The short answer is no, for arriving at a conclusion which *he knows to be true* necessarily requires that he know in advance that he is rigged to Logic I. But this is precisely what he does not know. However, let us consider a second possibility. Suppose for the sake of argument that Malcolm *can* establish that his thinking is regulated by Logic I. How will this fact assist him with his main problem of trying to determine whether or not he is a brain-in-a-vat? If Frankenstein is utilizing Logic I, how much more can Malcolm deduce? Only that he has *good reason* to believe that his brain might have been placed in a vat. I submit that *he cannot go any further*.

One way he might try to solve his problem is to consider whether vat-ideas are the same as real ideas. But it will be impossible to answer this question for the obvious reason that the scientist could be feeding him all the information upon which he bases the premises and conclusions of his arguments, and he has good reason to believe that this might actually be happening. For example, suppose he recalls his idea of his own home, and reasons that no-one could have *exactly* this idea of his home but him. It has a meaning for him that it does not have for anybody else. Further, any vat-idea of his home will not, he reasons, reflect this unique meaning, since its unique meaning for Malcolm is not available to the scientist. He might then rely on his memories to check his current ideas, and if he can discern a difference between them and his memories, then perhaps he can conclude that he is in a vat (and if he cannot discern a difference, that he is not in a vat). It should be clear that this approach will not succeed. For *any* ideas or memories or experiences upon which Malcolm bases any premises,

conclusions, or reflections whatsoever may be fed to him by Frankenstein. Recall that Frankenstein can govern totally Malcolm's mental content (including memory, imagination, emotion, etc.). It is consistent with all Malcolm knows (or thinks he knows) that he might simply be the subject of complete manipulation by Frankenstein. And indeed, if Malcolm admits the problem at all, he is admitting that he *cannot* tell if his ideas are real ideas or vat-ideas. (Similarly, *whatever conclusions* Descartes reasons to after his affirmation of the *cogito*, indeed even the affirmation of the *cogito* itself, might be the result of manipulation by the evil genius, an unpalatable conclusion which Descartes conveniently overlooked).

Once we shift from the evil genius argument to the brain-in-a-vat argument, we can see that the brain-in-a-vat problem *cannot* be solved. *For any piece of reasoning, memory, imagination, experience, etc. upon which Malcolm bases his philosophical reflection might be fed to him by the scientist.* He cannot even trust his logical inferences. (This latter point was overlooked by Descartes too. He did not subject logical inferences to his methodic doubt; for example, his inference from the fact that he is thinking, to the fact that he is. *Why* he did not push his skepticism this far is, I think, quite clear. If he had doubted logical inferences, he could not have proceeded at all with the problem of skepticism.) The shift to the brain-in-a-vat argument brings out this crucial point very clearly: if one takes seriously the possibility that one might be a brain-in-a-vat, then there is no possible way to prove that one is not a brain-in-a-vat. This is simply because any piece of reasoning by which one claims to have solved the problem might be the result of manipulation by the scientist. The possibility, however, must be taken seriously by Malcolm because he has *good reason* to believe that he might be a brain-in-a-vat.⁶

⁶ This discussion indicates that in a society which had developed brain-in-a-vat technology, it would be impossible for any individual to work out whether or not their brain had been placed in a vat. Whether or not one should accept this unpalatable conclusion seems to depend on whether one believes it will ever be possible to exactly reproduce human experiences in the brain-in-a-vat so that one could not tell the difference between vat experiences and real experiences. Malcolm, for example, could visit the laboratory, and kill Frankenstein, in an attempt to find out whether or not his brain had been placed in a vat! But even in doing this, and in going to jail for what appeared to be a very long time, he would be unable to tell if these things were really happening to him, or if they were simply being fed to him by Frankenstein, including the illusion of a lengthy jail term! How far could Malcolm proceed in this fashion before he could determine whether or not his brain had been placed in a vat? The answer to this question appears to depend on how likely we think it is that such technology could ever be developed. This scenario is yet another indication of both the absurdity of *our* problem of skepticism, and of the unsolvability of the problem.

Let us consider again the issue of Malcolm's excellent reasons for thinking that he may be a brain-in-a-vat. Does he have excellent reasons? Well, of course he does, because he knows from his own personal, direct experiences that it can be done, he has seen it done, and has even assisted at the task! This point illustrates that Malcolm should only consider the problem as real *if he knows* that it is possible to put brains in vats. Suppose for a moment that when travelling home on the bus Malcolm speculates that he might *not* know that brain-in-a-vat technology exists; suppose he doubts that he has seen it done, and that he has assisted in the task. That is, suppose, he speculates, that Frankenstein has simply manipulated his brain (in-a-vat) to make him believe that he has taken part in such experiments. This is an interesting thought experiment because if Malcolm were to come to accept this latter scenario he would not be able to solve the problem at all, but also, and just as important, *he would have no good reasons to take the problem seriously in the first place!* This is because if he comes to believe that he has never assisted with and has never even seen brain-in-a-vat experiments, then he has no reason to take the brain-in-a vat possibility seriously.

Now where does all of this leave us today with *our* problem of skepticism? The example above illustrates that it is only because Malcolm has excellent reasons *to begin with* that he should take the problem seriously, and it also shows that *we* should not take the problem seriously because we have no reasons to do so. The obvious lesson the above discussion has for me is that *I have no good reason* to believe that I might be a brain-in-a-vat. Further, I have excellent reasons to believe that I am not a brain-in-a-vat. Therefore, I should not consider that I might be a brain-in-a-vat. Malcolm *has* good reason to believe that he might be a brain-in-a-vat because he knows it is possible to put brains in vats, has seen it done, and has assisted in the task. But *we know* at this point in our history that it is not possible to put brains in vats. But does this really solve the problem of skepticism for us?

The skeptic will no doubt still insist that all I have said so far does not disprove the *possibility* that I might be a brain-in-a-vat all along (say on some other planet). After all, the skeptic will insist, perhaps *they* (the inhabitants of the far-off planet) have the technology, and I really am a brain-in-a-vat. It is *this* possibility, the skeptic might add, which I have to eliminate. However, my above remarks do make two significant points which present problems for this line of reasoning. First, I have no *good reason* to believe that this hypothesis might be true, therefore, it is very reasonable to dismiss it. We should go further and add that it would be *irrational* not to dismiss it. Second, if I *do* take this possibility seriously, there is no possible way to solve it, as we have just seen. Therefore, I must exercise great care before I decide to take it seriously because the consequences for human knowledge are so great. I must consider what *reasons*

there are to take it seriously. When I consider this issue, however, I quickly discover that there are none.

What the skeptic is really urging is that skepticism is a *logical possibility*, and until we prove that it is not a logical possibility we cannot consistently claim to have knowledge of the external world. A few brief remarks about the notion of logical possibility are in order here. When we say that something is logically possible we mean only that no contradiction is involved in its assertion.⁷ For example, it is logically possible for a man to jump one million feet into the air. This means that it is not logically contradictory to assert this statement. However, it is logically impossible to have a square circle. It is logically contradictory to assert the statement "a square circle." It is this sense of logical possibility that the skeptic has in mind when he argues that skepticism is logically possible.

It is obvious, however, that there is a significant difference between something being *logically possible* and being what we might call *practically possible*. It may be logically possible for a man to jump one million feet into the air but it is practically impossible. So the important question for our discussion now is: is it sufficient for us to take skepticism *seriously* that it be logically possible, or must it be practically possible too? Recall that the skeptic holds that the fact that skepticism is logically possible is a *good reason in itself* to doubt the validity and reliability of knowledge.

I wish to make two points in response to the skeptic's position which seriously weaken its force, to my mind. First, the fact that something is logically possible is not a reason, much less a good reason, to believe it or to take it seriously. This point has been obscured by our failure to realize and emphasize that the notion of logical possibility is much too broad to function as a criterion for deciding what possibilities to take seriously in human experience. For all kinds of bizarre, nonsensical and practically impossible scenarios can be described as logically possible. For example, jumping a million feet into the air, the earth is a giant spaceship, the moon is cheese, etc. are all logically possible scenarios. However, we do not regard any of these logical possibilities as practically possible because we do not have *good reasons* to do so. Further, and just as important, we have excellent reasons *not* to do so. We would only regard such things as practically possible if we had good reason, i.e., some empirical evidence, to believe they could occur, or were the case. For example, if the astronauts who visited the moon reported that it appeared to be made mainly of cheese, and if they brought

⁷ See John Hospers, *An Introduction to Philosophical Analysis* (Englewood Cliffs, New Jersey: Prentice Hall, Third Ed. 1988), p. 132ff for a discussion of the notion of logical possibility.

back samples to prove it, we would then have some reasons to believe that the moon is cheese. But to take seriously the belief that the moon is cheese simply because it is not a contradictory belief (i.e., because it is logically possible) would be absurd. The same standard *must* apply to the problem of skepticism. We should only take skepticism seriously if we have *good reason* to doubt our knowledge—if we have good reason to think we might be brains in vats. This is the difference between our situation and Malcolm's, between our problem of skepticism and Malcolm's problem of skepticism.

The second point I wish to make in relation to the claim that skepticism is a logical possibility is related to Stroud's position. The fact that we are apparently unable to eliminate the logical possibility that I might be a brain-in-a-vat does not automatically mean that we should seriously consider giving up on our claim to knowledge. We are, rather, confronted with the following dilemma: *either* we should consider giving up on our claim to knowledge, *or* we should investigate whether there is something wrong with *the problem* we are considering. I am suggesting that it is the problem which must be abandoned.⁸

In the debate concerning skepticism, there are three options: i) we can accept the problem of skepticism as real and important, and also agree that we can never solve the problem, and so we accept complete skepticism—that we do not have, and can never have, knowledge; ii) we can accept the problem of skepticism as real and important, and then solve it, and demonstrate that knowledge is possible; iii) we can argue that there is in fact no problem of skepticism, that we have knowledge, and that we do not have to worry about the hypothesis which suggests that all of our beliefs might be incorrect. My argument shows, I believe, that iii) is the correct response, and that i) and ii) are not serious options.

⁸Traditional realist responses to the question of knowledge are often regarded as avoiding the problem of skepticism, and of not taking seriously the fact that skepticism is a logical possibility. That is to say, any philosophical theory of knowledge which asserts that not only do I believe I am writing on this page now, but that I can know it, and know that I know it (that is, be absolutely certain of it), and that, therefore, I could not be in a vat, is often accused of *not* facing up to the logical possibility that one might be a brain-in-a-vat. However, what is not often brought out in this debate is the fact that realism is precisely the attempt to show that skepticism is *not* a logical possibility at all. That is to say, the realist holds that if I am here now, I cannot (*logically cannot*) be in a vat, or dreaming, or the subject of deception by an evil genius, etc. However reasonable and rooted in common sense, any proposed realist solution to the problem of skepticism has been dismissed by advocates of the problem out of hand. It seems that with *our* problem of skepticism the skeptic is prepared to go as far as it takes to frustrate any offered solution. Apparently, no attempt to refute the notion that skepticism is a *logical possibility* along realist lines will be acceptable. The history of the discussion bears this point out. My approach here is an attempt to develop a slightly different way of rejecting the problem.

In fact, we might label the skepticism that is urged upon us by the recalcitrant skeptic *comprehensive* (or general) skepticism. It is comprehensive because it advocates doubting all our claims to knowledge until we can eliminate the logical possibility that we might be a brain-in-a-vat. However, the fact that we *cannot* solve the problem is in principle a good reason to reject it, since it is a problem of our own *invention*, rather than a problem raised on the basis of evidence. In other words, we have raised an *unsolvable* problem (for no good reasons), and are then engaging in countless futile efforts attempting to solve it! This seems not only unreasonable, but absurd.

The point illustrated by my version of the brain-in-a-vat story is the crucial point here. My conclusion was that we could never solve the problem as stated. I also pointed out that in order to avoid this very problem, Descartes conveniently overlooked (as he went about the task of solving the problem of skepticism) the fact that the evil genius was supposed to place him in a position of *comprehensive* skepticism. Another way of stating this point is that comprehensive skepticism cannot *logically* be defended. For it is a position which would have to be argued to from knowledge of true premises, but knowledge of true premises is precisely what it precludes. As Dallas Willard has pointed out, somebody might still suffer from comprehensive skepticism as an *affliction* for which treatment would be appropriate.⁹ But it cannot be advanced as a *rational* ground for anything.

Let me briefly illustrate these points further by considering the skeptical arguments of Keith Lehrer, which bear directly on several of the points I have raised above.¹⁰ Lehrer argues that the skeptical hypothesis cannot be defeated because we are unable to show that we are ever *completely justified* in accepting as true any particular belief. We might have some good reasons for accepting the belief, but these reasons will always fall short of complete justification, and so we never, according to Lehrer, have *knowledge*. Lehrer attempts to frame and couch his position carefully so as to avoid some of the standard criticisms against skepticism. However, in my view he does not succeed in avoiding these criticisms. I will comment briefly on several of Lehrer's main points in the light of my remarks above.

First, Lehrer attempts to avoid saying that we know that we do not know anything so as to avoid self-contradiction. He talks about "avowing" skepticism instead of, say, proposing skepticism as a theory, and talks about "the contention"

⁹ See Dallas Willard, "Toward a Phenomenology for the Correspondence Theory of Truth", *Discipline filosofiche* (Bologna, Italy) I, 1991, pp. 125-147.

¹⁰ See Keith Lehrer, "Why not Skepticism?", *Philosophical Forum*, 2.3 (1971), pp. 283-298. Reprinted in *The Theory of Knowledge: Classic and Contemporary Readings*, ed. Louis Pojman (Belmont, California: Wadsworth, First Edition 1993), pp. 48-55. All page references are to Pojman.

of his argument rather than “my contention”. All of this is, I submit, a verbal sleight of hand employed in an attempt to avoid what cannot be avoided, i.e., the making of knowledge claims aimed at convincing the reader that you are right! Lehrer phrases his argument to say that “no one knows anything, not even that no one knows anything”¹¹, and claims that this is not a knowledge claim. However, this is a distinction without a difference in the debate about skepticism. For Lehrer is still committed to saying that this statement is *true*, even if he avoids putting it like that (for obvious reasons). If this statement is true, then he does not avoid self-contradiction; if it is false, then his argument for skepticism fails. If he claims that we do not know whether the statement is true or false, then we *know* that the truth-value of this statement is uncertain for us, and so he still faces self-contradiction. This problem cannot be avoided, in my view, by attempts at verbal sleights of hand.

Lehrer then attempts to undermine our confidence in the truth of many of our beliefs, including our ordinary, everyday perceptual beliefs. I will concentrate only on perceptual beliefs here. Let me take as an example the belief that I am sitting here now reading Lehrer’s article. How does Lehrer undermine the truth of this belief? He does so by presenting us with a modified brain-in-a-vat argument. He says that the skeptical hypothesis might be that there are a group of creatures in another galaxy called Googols whose intellectual capacity is vastly superior to ours and who amuse themselves by manipulating our brain waves so that most of our beliefs are incorrect (though they are very nearly correct, so that we can continue to operate). The Googols themselves do not know anything for they are also being manipulated by other beings on some other planet, and so on. Now I hope I have shown above that this is not a reasonable argument, and not one which we should take seriously. To present this hypothesis as a reason for why I should doubt that I am sitting here now reading Lehrer’s article is, to put it charitably, extremely far-fetched. Lehrer argues that if we cannot refute this hypothesis, we are simply ruling it out by fiat, and accepting the truth of perceptual beliefs without any justification. However, both of these claims are disingenuous. We are not ruling out his Googol theory by fiat, we are ruling it out because he simply invented it, and has no evidence whatsoever to back it up, or convince us that we should take it seriously. So our ruling it out is very reasonable, and ruling it in would be irrational. Also, I hold that my belief that I am sitting here now reading Lehrer’s article is a true belief because it is based on evidence, the evidence of the senses, of the reliability of sense knowledge, or the absence of any potential defeaters, etc. If the skeptic claims that we do not know that such evidence is reliable, he must explain why.

¹¹ *Ibid.*, p. 49.

If the only way he can do this is to introduce the Googol theory, then he is simply begging the question.

Lehrer talks as if, when I am considering the truth of my belief that I am sitting here now reading his article, the hypothesis that the Googols might be manipulating me is just as plausible as the belief that I am sitting here now, and that if I rule in favor of one, I am being dogmatic. My claim is that the view that I am sitting here now is totally rational, and the view that Googols might be manipulating me is a complete non-starter, and must be abandoned unless evidence for its plausibility is produced. Lehrer's claim that "we do not know that the skeptical hypothesis is false, and thus we do not know that anything else is true"¹² can be rewritten as "we believe that the skeptical hypothesis might be true, and thus we believe that all our knowledge claims might actually be false". Now this way of putting the matter clearly shows that he must present his reasons for believing in the plausibility of the skeptical hypothesis, otherwise the hypothesis can be rejected. If he tries to avoid the request for evidence, or to argue that all evidence is unreliable, then he is already assuming that the skeptical hypothesis is plausible, and so is begging the question.¹³

Like many skeptics, Lehrer finally tries to avoid the unpalatable consequences of his position. He argues that "we need not mourn the passing of knowledge as a great loss."¹⁴ This is because in scientific enquiry, for example, all contentions should be subject to question and must be defended on demand, and because we can still carry on practical affairs by settling for probability rather than truth in our beliefs. However, there are huge difficulties with these two claims. On the first claim, science would be impossible if we doubted *all* our beliefs, including our perceptual beliefs. And of course Lehrer seems to have forgotten that we cannot ultimately defend any of these beliefs on demand because of his Googol hypothesis. Therefore, no scientific claim, let alone theory, can be established as true on his view. Second, the problem with probability is not (as he argues) that probability is based on observation statements, which must be known to be *true* (which therefore defeats skepticism), a contention he replies to by reassuring us that our observation statements need only be probable too. The real problem with probability is that *we cannot accurately determine the*

¹² *Ibid.*, p. 55.

¹³ The plausibility of the skeptical hypothesis advanced by Lehrer consists solely in its being logically possible. However, it is also logically possible that we have knowledge. Lehrer does not seem to take this latter point as an adequate reason to deny skepticism. But if this fact is not an adequate reason to deny skepticism, why should the fact that skepticism is a logical possibility be an adequate reason to deny that we have knowledge? (I owe this point to Doug Geivett).

¹⁴ Keith Lehrer, "Why not Skepticism?", p. 55.

probability of any claim because of his Googol hypothesis. For any time we judge a belief to be more probable than another it might be the case that we are being manipulated by the Googols. So, on his view, we can only function in our practical affairs either by living a huge lie, believing our beliefs and theories to be true when we can never know this, or by acknowledging every time we entertain a belief, or make a claim, or present a theory, that it might not be true or even probable (which means, in short, in my view, that we cannot manage our practical affairs!). Alternatively, of course, as I have suggested, we could abandon the Googol theory.

It is necessary not to reject knowledge, but the problem of knowledge. An analogous argument will help to make this concluding point clear. Suppose we hypothesize that the earth is a giant spaceship being operated from within, and that human beings have been placed on the surface and given the conditions for life by those who control the ship, all for some unknown reason. What reasons do we have for believing that this might be true? At present, we have none. Therefore, we should dismiss this hypothesis as not worthy of consideration. It would be simply ridiculous to stubbornly insist, despite the complete absence of evidence, that nevertheless the earth *might* still be a giant spaceship, and until we show that it is not, we cannot be sure that our lives are not being controlled by somebody from within the earth. But this is just what we do with the problem of skepticism. The great scandal of philosophy, as Heidegger observed, is not that we have been unable to solve the problem of skepticism, but that we have taken the problem seriously.¹⁵

I have attempted to illustrate that if we *do* take the problem of skepticism seriously we will never be able to solve it. This is an excellent reason in itself in favor of dismissing the problem since it is a problem of our own invention, and not one raised on the basis of evidence. I have also argued that the claim that skepticism is a logical possibility is not a good reason to take the problem seriously, and that there are no other good reasons to take the problem seriously. For reasons such as these, I advocate that we should dismiss the problem of skepticism as a pseudo-problem. Despite all of these arguments, one may yet hear the recalcitrant skeptic *still* insist that nevertheless skepticism remains a logical possibility. It is at this point that I suggest that the skeptic is *not arguing rationally* for the fact that we should take the problem of skepticism seriously, but is in fact dangerously close to suffering from skepticism as an affliction! In cases like this treatment would be appropriate, exactly what kind I cannot say, though my students have suggested several kinds over the years, all of them quite unpleasant!

¹⁵ See Martin Heidegger, *Being and Time*, trans. J. Macquarrie and E. Robinson (San Francisco: Harper and Row, 1962), p. 249.