

12. This is the best summary I can give of what is wrong and right in Prichard's claim that one can determine by reflection whether one knows something or merely believes it. A good part of the ideas in this essay were provoked by conversations with Wittgenstein. A brief and rough account of those talks is to

be found in my *Ludwig Wittgenstein: A Memoir* (New York: Oxford University Press, 1958), pp. 87–92. Jaakko Hintikka provides an acute treatment of the topic of "knowing that one knows," with special reference to Prichard's claim. See his *Knowledge and Belief* (Ithaca: Cornell University Press, 1962), ch. 5.

Skepticism

ROBERT NOZICK

The skeptic about knowledge argues that we know very little or nothing of what we think we know, or at any rate that this position is no less reasonable than the belief in knowledge.* The history of philosophy exhibits a number of different attempts to refute the skeptic: to prove him wrong or show that in arguing against knowledge he presupposes there is some and so refutes himself. Others attempt to show that accepting skepticism is unreasonable, since it is more likely that the skeptic's extreme conclusion is false than that all of his premisses are true, or simply because reasonableness of belief just means proceeding in an anti-skeptical way. Even when these counterarguments satisfy their inventors, they fail to satisfy others, as is shown by the persistent attempts against skepticism.¹ The continuing felt need to refute skepticism, and the difficulty in doing so, attests to the power of the skeptic's position, the depth of his worries.

An account of knowledge should illumi-

nate skeptical arguments and show wherein lies their force. If the account leads us to reject these arguments, this had better not happen too easily or too glibly. To think the skeptic overlooks something obvious, to attribute to him a simple mistake or confusion or fallacy, is to refuse to acknowledge the power of his position and the grip it can have upon us. We thereby cheat ourselves of the opportunity to reap his insights and to gain self-knowledge in understanding why his arguments lure us so. Moreover, in fact, we cannot lay the specter of skepticism to rest without first hearing what it shall unfold.

Our goal is not, however, to refute skepticism, to prove it is wrong or even to argue that it is wrong. In the Introduction we distinguished between philosophy that attempts to prove, and philosophy that attempts to explain how something is possible. Our task here is to explain how knowledge is possible, given what the skeptic says that we do accept (for example, that it is logically possible that we are dreaming or are floating in the tank). In doing this, we need not convince the skeptic, and we may introduce explanatory hypotheses that he would reject. What is important for our task of explanation and understanding is that *we* find those hypotheses acceptable or plausible, and that they show us how the existence of knowledge fits together with the logical

Reprinted from *Philosophical Explanations* by Robert Nozick (Cambridge, Mass.: Harvard University Press), with permission of the author and the publishers, copyright 1981 by Robert Nozick.

*Nozick's conditions for knowledge are to be found on pages 26–37 of this volume.

possibilities the skeptic points to, so that these are reconciled within our own belief system. These hypotheses are to explain to ourselves how knowledge is possible, not to prove to someone else that knowledge is possible.²

SKEPTICAL POSSIBILITIES

The skeptic often refers to possibilities in which a person would believe something even though it was false: really, the person is cleverly deceived by others, perhaps by an evil demon, or the person is dreaming or he is floating in a tank near Alpha Centauri with his brain being stimulated. In each case, the *p* he believes is false, and he believes it even though it is false.

How do these possibilities adduced by the skeptic show that someone does not know that *p*? Suppose that someone is you; how do these possibilities count against your knowing that *p*? One way might be the following. (I shall consider other ways later.) If there is a possible situation where *p* is false yet you believe that *p*, then in that situation you believe that *p* even though it is false. So it appears you do not satisfy condition 3 for knowledge.

(3) If *p* were false, S wouldn't believe that *p*.

For a situation has been described in which you do believe that *p* even though *p* is false. How then can it also be true that if *p* were false, you wouldn't believe it? If the skeptic's possible situation shows that 3 is false, and if 3 is a necessary condition for knowledge, then the skeptic's possible situation shows that there isn't knowledge.

So construed, the skeptic's argument plays on condition 3; it aims to show that condition 3 is not satisfied. The skeptic may seem to be putting forth

R: Even if *p* were false, S still would believe *p*.³

This conditional, with the same antecedent as 3 and the contradictory consequent, is in-

compatible with the truth of 3. If 3 is true, then R is not. However, R is stronger than the skeptic needs in order to show 3 is false. For 3 is false when if *p* were false, S might believe that *p*. This last conditional is weaker than R, and is merely 3's denial:

T: not-[not-*p* → not-(S believes that *p*)].

Whereas R does not simply deny 3, it asserts an opposing subjunctive of its own. Perhaps the possibility the skeptic adduces is not enough to show that R is true, but it appears at least to establish the weaker T; since this T denies 3, the skeptic's possibility appears to show that 3 is false.⁴

However, the truth of 3 is not incompatible with the existence of a possible situation where the person believes *p* though it is false. The subjunctive

(3) not-*p* → not-(S believes *p*)

does not talk of all possible situations in which *p* is false (in which not-*p* is true). It does not say that in all possible situations where not-*p* holds, S doesn't believe *p*. To say there is no possible situation in which not-*p* yet S believes *p*, would be to say that not-*p* entails not-(S believes *p*), or logically implies it. But subjunctive conditionals differ from entailments; the subjunctive 3 is not a statement of entailment. So the existence of a possible situation in which *p* is false yet S believes *p* does not show that 3 is false;⁵ 3 can be true even though there is a possible situation where not-*p* and S believes that *p*.

What the subjunctive 3 speaks of is the situation that would hold if *p* were false. Not every possible situation in which *p* is false is the situation that would hold if *p* were false. To fall into possible worlds talk, the subjunctive 3 speaks of the not-*p* world that is closest to the actual world, or of those not-*p* worlds that are closest to the actual world, or more strongly (according to my suggestion) of the not-*p* neighborhood of the actual world. And it is of this or these not-*p* worlds that it says (in them) S does not believe that *p*. What

happens in yet other more distant not- p worlds is no concern of the subjunctive 3.

The skeptic's possibilities (let us refer to them as SK), of the person's being deceived by a demon or dreaming or floating in a tank, count against the subjunctive

(3) if p were false then S wouldn't believe that p

only if (one of) these possibilities would or might obtain if p were false; only if one of these possibilities is in the not- p neighborhood of the actual world. Condition 3 says: if p were false, S still would not believe p . And this can hold even though there is some situation SK described by the skeptic in which p is false and S believes p . If p were false S still would not believe p , even though there is a situation SK in which p is false and S does believe p , provided that this situation SK wouldn't obtain if p were false. If the skeptic describes a situation SK which would not hold even if p were false then this situation SK doesn't show that 3 is false and so does not (in this way at least) undercut knowledge. Condition C acts to rule out skeptical hypotheses.

C: not- $p \rightarrow$ SK does not obtain.

Any skeptical situation SK which satisfies condition C is ruled out. For a skeptical situation SK to show that we don't know that p , it must fail to satisfy C which excludes it; instead it must be a situation that might obtain if p did not, and so satisfy C's denial:

not-(not- $p \rightarrow$ SK doesn't obtain).

Although the skeptic's imagined situations appear to show that 3 is false, they do not; they satisfy condition C and so are excluded.

The skeptic might go on to ask whether we know that his imagined situations SK are excluded by condition C, whether we know that if p were false SK would not obtain. However, typically he asks something stronger: do we know that his imagined situation SK does not actually obtain? Do we know that we are not being deceived by a

demon, dreaming, or floating in a tank? And if we do not know this, how can we know that p ? Thus we are led to the second way his imagined situations might show that we do not know that p .

SKEPTICAL RESULTS

According to our account of knowledge, S knows that the skeptic's situation SK doesn't hold if and only if

- (1) SK doesn't hold
- (2) S believes that SK doesn't hold
- (3) If SK were to hold, S would not believe that SK doesn't hold
- (4) If SK were not to hold, S would believe it does not.

Let us focus on the third of these conditions. The skeptic has carefully chosen his situations SK so that if they held we (still) would believe they did not. We would believe we weren't dreaming, weren't being deceived, and so on, even if we were. He has chosen situations SK such that if SK were to hold, S would (still) believe that SK doesn't hold—and this is incompatible with the truth of 3.⁶

Since condition 3 is a necessary condition for knowledge, it follows that we do not know that SK doesn't hold. If it were true that an evil demon was deceiving us, if we were having a particular dream, if we were floating in a tank with our brains stimulated in a specified way, we would still believe we were not. So, we do not know we're not being deceived by an evil demon, we do not know we're not in that tank, and we do not know we're not having that dream. So says the skeptic, and so says our account. And also so we say—don't we? For how could we know we are not being deceived that way, dreaming that dream? If those things *were* happening to us, everything would seem the same to us. There is no way we can know it is not happening for there is no way we could tell if it were happening; and if it were happening we would believe exactly what we do now—in particular, we still would believe

that it was not. For this reason, we feel, and correctly, that we don't know—how could we?—that it is not happening to us. It is a virtue of our account that it yields, and explains, this result.

The skeptic asserts we do not know his possibilities don't obtain, and he is right. Attempts to avoid skepticism by claiming we do know these things are bound to fail. The skeptic's possibilities make us uneasy because, as we deeply realize, we do not know they don't obtain; it is not surprising that attempts to show we do know these things leave us suspicious, strike us even as bad faith.⁷ Nor has the skeptic merely pointed out something obvious and trivial. It comes as a surprise to realize that we do not know his possibilities don't obtain. It is startling, shocking. For we would have thought, before the skeptic got us to focus on it, that we did know those things, that we did know we were not being deceived by a demon, or dreaming that dream, or stimulated that way in that tank. The skeptic has pointed out that we do not know things we would have confidently said we knew. And if we don't know these things, what can we know? So much for the supposed obviousness of what the skeptic tells us.

Let us say that a situation (or world) is doxically identical for S to the actual situation when if S were in that situation, he would have exactly the beliefs (*doxa*) he actually does have. More generally, two situations are doxically identical for S if and only if he would have exactly the same beliefs in them. It might be merely a curiosity to be told there are nonactual situations doxically identical to the actual one. The skeptic, however, describes worlds doxically identical to the actual world in which almost everything believed is false.⁸

Such worlds are possible because we know mediately, not directly. This leaves room for a divergence between our beliefs and the truth. It is as though we possessed only two-dimensional plane projections of three-dimensional objects. Different three-dimensional objects, oriented appropriately, have the same two-dimensional plane projection.

Similarly, different situations or worlds will lead to our having the very same beliefs. What is surprising is how very different the doxically identical world can be—different enough for almost everything believed in it to be false. Whether or not the mere fact that knowledge is mediated always makes room for such a very different doxically identical world, it does so in our case, as the skeptic's possibilities show. To be shown this is nontrivial, especially when we recall that we do not know the skeptic's possibility doesn't obtain: we do not know that we are not living in a doxically identical world wherein almost everything we believe is false.⁹

What more could the skeptic ask for or hope to show? Even readers who sympathized with my desire not to dismiss the skeptic too quickly may feel this has gone too far, that we have not merely acknowledged the force of the skeptic's position but have succumbed to it.

The skeptic maintains that we know almost none of what we think we know. He has shown, much to our initial surprise, that we do not know his (nontrivial) possibility SK doesn't obtain. Thus, he has shown of one thing we thought we knew, that we didn't and don't. To the conclusion that we know almost nothing, it appears but a short step. For if we do not know we are not dreaming or being deceived by a demon or floating in a tank, then how can I know, for example, that I am sitting before a page writing with a pen, and how can you know that you are reading a page of a book?

However, although our account of knowledge agrees with the skeptic in saying that we do not know that not-SK, it places no formidable barriers before my knowing that I am writing on a page with a pen. It is true that I am, I believe I am, if I weren't I wouldn't believe I was, and if I were, I would believe it. (I leave out the reference to method.) Also, it is true that you are reading a page (please, don't stop now!), you believe you are, if you weren't reading a page you wouldn't believe you were, and if you were reading a page you would believe you were.

So according to the account, I do know that I am writing on a page with a pen, and you do know that you are reading a page. The account does not lead to any general skepticism.

Yet we must grant that it appears that if the skeptic is right that we don't know we are not dreaming or being deceived or floating in the tank, then it cannot be that I know I am writing with a pen or that you know you are reading a page. So we must scrutinize with special care the skeptic's "short step" to the conclusion that we don't know these things, for either this step cannot be taken or our account of knowledge is incoherent.

NONCLOSURE

In taking the "short step", the skeptic assumes that if S knows that p and he knows that ' p entails q ' then he also knows that q . In the terminology of the logicians, the skeptic assumes that knowledge is closed under known logical implication; that the operation of moving from something known to something else known to be entailed by it does not take us outside of the (closed) area of knowledge. He intends, of course, to work things backwards, arguing that since the person does not know that q , assuming (at least for the purposes of argument) that he does know that p entails q , it follows that he does not know that p . For if he did know that p , he would also know that q , which he doesn't.

The details of different skeptical arguments vary in their structure, but each one will assume some variant of the principle that knowledge is closed under known logical implication. If we abbreviate "knowledge that p " by " Kp " and abbreviate "entails" by the fishhook sign " \rightarrow ", we can write this principle of closure as the subjunctive principle

P: $K(p \rightarrow q) \ \& \ Kp \rightarrow Kq$.

If a person were to know that p entails q and he were to know that p then he would know that q . The statement that q follows by modus

ponens from the other two stated as known in the antecedent of the subjunctive principle P; this principle counts on the person to draw the inference to q .

You know that your being in a tank on Alpha Centauri entails your not being in place X where you are. (I assume here a limited readership.) And you know also the contrapositive, that your being at place X entails that you are not then in a tank on Alpha Centauri. If you knew you were at X you would know you're not in a tank (of a specified sort) at Alpha Centauri. But you do not know this last fact (the skeptic has argued and we have agreed) and so (he argues) you don't know the first. Another intuitive way of putting the skeptic's argument is as follows. If you know that two statements are incompatible and you know the first is true then you know the denial of the second. You know that your being at X and your being in a tank on Alpha Centauri are incompatible; so if you knew you were at X you would know you were not in the (specified) tank on Alpha Centauri. Since you do not know the second, you don't know the first.¹⁰

No doubt, it is possible to argue over the details of principle P, to point out it is incorrect at it stands. Perhaps, though Kp , the person does not know that he knows that p (that is, not- KKp) and so does not draw the inference to q . Or perhaps he doesn't draw the inference because not- $KK(p \rightarrow q)$. Other similar principles face their own difficulties: for example, the principle that $K(p \rightarrow q) \rightarrow (Kp \rightarrow Kq)$ fails if Kp stops $p \rightarrow q$ from being true, that is, if $Kp \rightarrow \text{not-}(p \rightarrow q)$; the principle that $K(p \rightarrow q) \rightarrow K(Kp \rightarrow Kq)$ faces difficulties if Kp makes the person forget that $(p \rightarrow q)$ and so he fails to draw the inference to q . We seem forced to pile K upon K until we reach something like $KK(p \rightarrow q) \ \& \ KKp \rightarrow Kq$; this involves strengthening considerably the antecedent of P and so is not useful for the skeptic's argument that p is not known. (From a principle altered thus, it would follow at best that it is not known that p is known.)

We would be ill-advised, however, to

quibble over the details of P. Although these details are difficult to get straight, it will continue to appear that something like P is correct. If S knows that ' p entails q ' and he knows that p and knows that ' p and p entails q ' (shades of the Lewis Carroll puzzle we discuss below!) and he does draw the inference to q from all this and believes q via the process of drawing this inference, then will he not know that q ? And what is wrong with simplifying this mass of detail by writing merely principle P, provided we apply it only to cases where the mass of detail holds, as it surely does in the skeptical cases under consideration? For example, I do realize that my being in the Van Leer Foundation Building in Jerusalem entails that I am not in a tank on Alpha Centauri; I am capable of drawing inferences now; I do believe I am not in a tank on Alpha Centauri (though not solely via this inference, surely); and so forth. Won't this satisfy the correctly detailed principle, and shouldn't it follow that I know I am not (in that tank) on Alpha Centauri? The skeptic agrees it should follow; so he concludes from the fact that I don't know I am not floating in the tank on Alpha Centauri that I don't know I am in Jerusalem. Uncovering difficulties in the details of particular formulations of P will not weaken the principle's intuitive appeal; such quibbling will seem at best like a wasp attacking a steamroller, at worst like an effort in bad faith to avoid being pulled along by the skeptic's argument.

Principle P is wrong, however, and not merely in detail. Knowledge is not closed under known logical implication.¹¹ S knows that p when S has a true belief that p , and S wouldn't have a false belief that p (condition 3) and S would have a true belief that p (condition 4). Neither of these latter two conditions is closed under known logical implication.

Let us begin with condition

- (3) if p were false, S wouldn't believe that p .

When S knows that p , his belief that p is contingent on the truth of p , contingent in

the way the subjunctive condition 3 describes. Now it might be that p entails q (and S knows this), that S's belief that p is subjunctively contingent on the truth of p , that S believes q , yet his belief that q is not subjunctively dependent on the truth of q , in that it (or he) does not satisfy:

- (3') if q were false, S wouldn't believe that q .

For 3' talks of what S would believe if q were false, and this may be a very different situation than the one that would hold if p were false, even though p entails q . That you were born in a certain city entails that you were born on earth.¹² Yet contemplating what (actually) would be the situation if you were not born in that city is very different from contemplating what situation would hold if you weren't born on earth. Just as those possibilities are very different, so what is believed in them may be very different. When p entails q (and not the other way around) p will be a stronger statement than q , and so not- q (which is the antecedent of 3') will be a stronger statement than not- p (which is the antecedent of 3). There is no reason to assume you will have the same beliefs in these two cases, under these suppositions of differing strengths.

There is no reason to assume the (closest) not- p world and the (closest) not- q world are doxically identical for you, and no reason to assume, even though p entails q , that your beliefs in one of these worlds would be a (proper) subset of your beliefs in the other.

Consider now the two statements:

- p = I am awake and sitting on a chair in Jerusalem;
 q = I am not floating in a tank on Alpha Centauri being stimulated by electrochemical means to believe that p .

The first one entails the second: p entails q . Also, I know that p entails q ; and I know that p . If p were false, I would be standing or lying down in the same city, or perhaps sleeping there, or perhaps in a neighboring city or town. If q were false, I would be float-

ing in a tank on Alpha Centauri. Clearly these are very different situations, leading to great differences in what I then would believe. If p were false, if I weren't awake and sitting on a chair in Jerusalem, I would not believe that p . Yet if q were false, if I was floating in a tank on Alpha Centauri, I would believe that q , that I was not in the tank, and indeed, in that case, I would still believe that p . According to our account of knowledge, I know that p yet I do not know that q , even though (I know) p entails q .

This failure of knowledge to be closed under known logical implication stems from the fact that condition 3 is not closed under known logical implication; condition 3 can hold of one statement believed while not of another known to be entailed by the first.¹³ It is clear that any account that includes as a necessary condition for knowledge the subjunctive condition 3, not- $p \rightarrow$ not-(S believes that p), will have the consequence that knowledge is not closed under known logical implication.¹⁴

When p entails q and you believe each of them, if you do not have a false belief that p (since p is true) then you do not have a false belief that q . However, if you are to know something not only don't you have a false belief about it, but also you wouldn't have a false belief about it. Yet, we have seen how it may be that p entails q and you believe each and you wouldn't have a false belief that p yet you might have a false belief that q (that is, it is not the case that you wouldn't have one). Knowledge is not closed under the known logical implication because "wouldn't have a false belief that" is not closed under known logical implication.

If knowledge were the same as (simply) true belief then it would be closed under known logical implication (provided the implied statements were believed). Knowledge is not simply true belief, however; additional conditions are needed. These further conditions will make knowledge open under known logical implication, even when the entailed statement is believed, when at least one of the further conditions itself is open. Knowledge stays closed (only) if all of the

additional conditions are closed. I lack a general nontrivial characterization of those conditions that are closed under known logical implication; possessing such an illuminating characterization, one might attempt to prove that no additional conditions of that sort could provide an adequate analysis of knowledge.

Still, we can say the following. A belief that p is knowledge that p only if it somehow varies with the truth of p . The causal condition for knowledge specified that the belief was "produced by" the fact, but that condition did not provide the right sort of varying with the fact. The subjunctive conditions 3 and 4 are our attempt to specify that varying. But however an account spells this out, it will hold that whether a belief that p is knowledge partly depends on what goes on with the belief in some situations when p is false. An account that says nothing about what is believed in any situation when p is false cannot give us any mode of varying with the fact.

Because what is preserved under logical implication is truth, any condition that is preserved under known logical implication is most likely to speak only of what happens when p , and q , are true, without speaking at all of what happens when either one is false. Such a condition is incapable of providing "varies with"; so adding only such conditions to true belief cannot yield an adequate account of knowledge.¹⁵

A belief's somehow varying with the truth of what is believed is not closed under known logical implication. Since knowledge that p involves such variation, knowledge also is not closed under known logical implication. The skeptic cannot easily deny that knowledge involves such variation, for his argument that we don't know that we're not floating in that tank, for example, uses the fact that knowledge does involve variation. ("If you were floating in the tank you would still think you weren't, so you don't know that you're not.") Yet, though one part of his argument uses that fact that knowledge involves such variation, another part of his argument presupposes that knowledge does

not involve any such variation. This latter is the part that depends upon knowledge being closed under known logical implication, as when the skeptic argues that since you don't know that not-SK, you don't know you are not floating in the tank, then you also don't know, for example, that you are now reading a book. That closure can hold only if the variation does not. The skeptic cannot be right both times. According to our view he is right when he holds that knowledge involves such variation and so concludes that we don't know, for example, that we are not floating in that tank; but he is wrong when he assumes knowledge is closed under known logical implication and concludes that we know hardly anything.¹⁶

Knowledge is a real factual relation, subjunctively specifiable, whose structure admits our standing in this relation, tracking to p without standing in it to some q which we know p to entail. Any relation embodying some variation of belief with the fact, with the truth (value), will exhibit this structural feature. The skeptic is right that we don't track some particular truths—the ones stating that his skeptical possibilities SK don't hold—but wrong that we don't stand in the real knowledge-relation of tracking to many other truths, including ones that entail these first-mentioned truths we believe but don't know.

The literature on skepticism contains writers who endorse these skeptical arguments (or similar narrower ones), but confess their inability to maintain their skeptical beliefs at times when they are not focusing explicitly on the reasoning that led them to skeptical conclusions. The most notable example of this is Hume:

I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another . . . Most fortunately it happens that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium, either by relaxing this bent of mind, or by some avocation, and lively impression of my senses, which obliterate all these chimeras. I dine, I play a game of backgammon, I

converse, and am merry with my friends; and when after three or four hours' amusement, I would return to these speculations, they appear so cold, and strained, and ridiculous, that I cannot find in my heart to enter into them any farther. (*A Treatise of Human Nature*, Book I, Part IV, section VII).

The great subverter of Pyrrhonism or the excessive principles of skepticism is action, and employment, and the occupations of common life. These principles may flourish and triumph in the schools; where it is, indeed, difficult, if not impossible, to refute them. But as soon as they leave the shade, and by the presence of the real objects, which actuate our passions and sentiments, are put in opposition to the more powerful principles of our nature, they vanish like smoke, and leave the most determined skeptic in the same condition as other mortals . . . And though a Pyrrhonian may throw himself or others into a momentary amazement and confusion by his profound reasonings; the first and most trivial event in life will put to flight all his doubts and scruples, and leave him the same, in every point of action and speculation, with the philosophers of every other sect, or with those who never concerned themselves in any philosophical researches. When he awakes from his dream, he will be the first to join in the laugh against himself, and to confess that all his objections are mere amusement. (*An Enquiry Concerning Human Understanding*, Section XII, Part II)

The theory of knowledge we have presented explains why skeptics of various sorts have had such difficulties in sticking to their far-reaching skeptical conclusions "outside the study", or even inside it when they are not thinking specifically about skeptical arguments and possibilities SK.

The skeptic's arguments do show (but show only) that we don't know the skeptic's possibilities SK do not hold; and he is right that we don't track the fact that SK does not hold. (If it were to hold, we would still think it didn't.) However, the skeptic's arguments don't show we do not know other facts (including facts that entail not-SK) for we do track these other facts (and knowledge is not closed under known logical entailment.) Since we do track these other facts—you, for example, the fact that you are reading a book; I, the fact that I am writing on a

page—and the skeptic tracks such facts too, it is not surprising that when he focuses on them, on his relationship to such facts, the skeptic finds it hard to remember or maintain his view that he does not know those facts. Only by shifting his attention back to his relationship to the (different) fact that not-SK, which relationship is not tracking, can he revive his skeptical belief and make it salient. However, this skeptical triumph is evanescent, it vanishes when his attention turns to other facts. Only by fixating on the skeptical possibilities SK can he maintain his skeptical virtue; otherwise, unsurprisingly, he is forced to confess to sins of credulity.

SKEPTICISM AND THE CONDITIONS FOR KNOWLEDGE

We have considered how the skeptic's argument from the skeptical possibilities SK plays off condition 3: if p weren't true S wouldn't believe that p . His argument gains its power by utilizing this condition ("but even if SK held, you still would believe it didn't, so you do not know it doesn't"); the deep intuitive force of the argument indicates that condition 3 (or something very much like it) is a necessary condition for knowledge. Similarly, are there any skeptical arguments or moves that play off condition 4: if p were true then S would believe that p (and wouldn't believe that not- p)? If condition 3 specifies how belief somehow should vary with the truth of what is believed, condition 4 specifies how belief shouldn't vary when the truth of what is believed does not vary. Condition 3 is a variation condition, condition 4 is an adherence condition. Both conditions together capture the notion that S (who actually truly believes p) would have a true belief that p . He wouldn't have a false belief that p if p weren't true (condition 3), and he would have a true belief that p if p were true (condition 4). Just as the skeptic argued earlier that the belief wouldn't vary when it should, he also can argue that it would vary

when it shouldn't, concluding both times that we don't have knowledge.

We would expect skeptical arguments playing off condition 4 to be less powerful and compelling than ones playing off 3. Condition 3 requires that we wouldn't falsely believe p , and we can be led to worry not only whether we might but whether we do. While condition 4 requires that we would truly believe p (and wouldn't falsely believe not- p), and though we might worry whether we might violate this, we need have no fear that we are—for we know we are believing p and are not believing not- p . Skeptical arguments playing off condition 4, unlike those with 3, cannot make us wonder also whether we violate the condition's indicative version.

Condition 4 is an adherence condition, so the relevant doubts concern how securely you are tied to the truth. For many (most?) of the things p you believe, if a group of people came and deceitfully told you not- p , you would believe them and stop believing p . (Relevant experiments frequently have been done by social psychologists.) So do you really know p ? If physicists told you that Newton's theory turns out to have been correct after all, wouldn't (or mightn't) you believe them? So do you really know Newtonian theory is false?

But, as before, the mere possibility of its being true while you do not believe it is not sufficient to show you don't actually know it. That possibility must be one that might arise. Call this possibility of p 's being true while you don't believe it: sk . (Lowercase " sk " is p 's being true and your not believing it, while capital SK is p 's being false and your believing p .) Possibility sk need not concern us when: if p were true, sk wouldn't hold; $p \rightarrow$ not- sk ; sk is false throughout the first part of the p neighborhood of the actual world. It is fortunate for my knowing that p that there wouldn't be people who trick me, just as it is fortunate for my knowing I am in Emerson Hall that whatever would occur if I weren't there does not include people tricking or hypnotizing me into believing I am there.

Suppose I present a certain argument to

someone who believes (truly) that p , and he is convinced by it and comes to believe not- p . Look how easily he can be moved from believing p to believing not- p . Suppose it happens that I do not present the argument to him, so he does not start to believe not- p , and he continues to believe p . Does he know that p ? Is it merely the case that his knowledge is insecure, or does such instability show it is not knowledge after all?

A skeptic might argue that for almost each p we (think we) know, there is an argument or happening that would get us to believe not- p even though p was true. We reply to this skeptic as before—the fact that some possible argument or happening would get us to believe not- p when p doesn't show that it is false that 4: if p were true then S would believe p and S wouldn't believe not- p . To show the falsity of 4, the skeptic would have to refer to something that might occur if p were true; if it wouldn't hold if p were true, what he refers to is irrelevant.

Among the arguments that get people to stop believing things are the skeptic's arguments themselves. These arguments often puzzle people, sometimes they get people to stop believing they know that p . They do not know that they know. Should we describe this as a case of people who first know that they know but who, after hearing the skeptic's arguments, no longer know that they know because they no longer believe that they know (and knowledge entails belief)? Our present view is that such people did not know that they know that p , even before hearing the skeptic. For their previous belief that they knew that p would vary when it shouldn't, so it violates condition 4. Similarly, some people who never have heard the skeptic's arguments would (if they heard them) become convinced that they don't know that p . It is pleasant to grant the skeptic a partial victory after all, one gained by the plausibility of his arguments, not their cogency. Because of the skeptical arguments, some people would falsely believe they don't know that p , and these people do not know they know it. The existence of skeptical arguments makes one type of skeptical conclu-

sion (that we don't know we know things) true of some people—those the shoe fits have been wearing it.

Meno claimed he could speak eloquently about virtue until Socrates, torpedolike, began to question him. He did not know what virtue was, for Socrates' questions uncovered Meno's previously existing confusions. Even if it had been a sophist's questions that bewildered Meno, getting him to believe the opposite, what he previously had would not have been knowledge. Knowledge should be made of sterner stuff.¹⁷

Thus, some skeptical arguments play off condition 3, others off condition 4. In addition to these conditions, our (full) account of knowledge formulates a condition about outweighing to cover the situation when multiple methods, not all satisfying 3 and 4, give rise to the belief. Do any skeptical arguments play off this outweighing condition? Here, presumably, would fit various attempts at unmasking the dominant sources of our belief as methods that do not track: faith, prejudice, self-interest, class-interest, deep psychological motives. The outweighing view involves subjunctives, but does anything here correspond to the skeptic's focusing upon a possibility that is so far out that it wouldn't occur, even if p were false? Perhaps the following is comparable. Recall that it was not necessary for the tracking method to win out against the combined opposed weight of all other methods; the person's belief merely had to vary with the verdict of the tracking method when the recommendations of every other way used to arrive at belief were held fixed. (It was only Case III in the chart that needed to be examined.) Any actual split in the verdict of nontracking methods will be welcome support. The skeptic should not load the other methods against what tracking recommends, any more than they actually are; to suppose more counts as too far out.

Some skeptical arguments play off condition 3, some off condition 4, some (perhaps) off the outweighing condition when multiple methods are involved. Still other skeptical arguments play off the methods

themselves, off the fact that knowledge is gained via methods or ways of believing. In the situations when we are aware of what methods we are using, do we know we are using those methods? To decide whether we know this, according to condition 3 we must consider what we would believe if we weren't using the methods. Would we then still believe we were? If so, condition 3 is violated, and so we did not actually know we were using the methods.

Along this pathway lies trouble. For if we weren't using that method, the very method we use to track various facts—a situation we have to contemplate in applying condition 3—who knows what we would believe about what methods we are using? That method *M* we are using to track various facts may be the very method via which we believe that we are using method *M*. This is likely if (and only if) *M* is described widely and deeply enough, for example, as the sum total of our (rational or effective) methods. But then, how are we to treat the question of what we would believe if we weren't using that method *M*, a question condition 3 pushes at us in order to decide if we know we are using *M*? "If I weren't using *M*, would I still believe I was?" What methods of believing am I left by this question? After all, condition 3 when fully formulated says: not-*p* and *S*, via *M*, comes to a belief about the truth of *p* → not-(*S* believes that *p*). And the method *M* of condition 3 is the very one said to be actually utilized, in condition 2: *S* believes, via *M*, that *p*.

Yet now we face the situation where *S* believes of himself that he is applying method *M*, via an application of method *M* itself;¹⁸ moreover, in this situation the statement *p*, which we are trying to decide whether *S* knows, is: *S* is using method *M*. The result of substituting this *p* in the full condition 3 is: If *S* weren't using method *M*, and *S*, via using *M*, were to decide about the truth of 'S is using method *M*' then *S* would not believe 'S is using method *M*'. But the antecedent of this subjunctive is supposing both that *S* is not using method *M* (this supposition is the not-*p* of the antecedent of

condition 3) and that *S* is using method *M* (he uses this method in 3 to decide whether or not *p*, since that is the method via which, in condition 2, he actually believes *p*). We have no coherent way to understand this.¹⁹

Yet if we cannot simply include the use of method *M* in determining what *S* would believe if he were not using *M*, neither can we simply suppose (for the purposes of condition 3) that *S* is using some other method to arrive at a belief about this matter. We saw earlier, in considering a range of examples, the great importance of holding the method fixed in deciding questions about knowledge. Recall the grandmother who sees her grandson visit her and so believes he is healthy and ambulatory; yet if he weren't ambulatory, other relatives would tell her he was fine to spare her anxiety and upset. She sees her grandson walking; does she know he is ambulatory? According to condition 3 we must ask what she would believe if he weren't ambulatory. If the method via which she believes is not held fixed, the answer will be wrong. True, if he weren't ambulatory, she would then believe he was (via hearing about him from other relatives). But the relevant question is: what would she believe if he weren't ambulatory and (as before) she saw him and spoke to him. Thus, to reach the correct answer about her knowledge, the method must be held fixed—that is one of the reasons why we introduced explicit reference to the method or way of believing.

How then are we to treat the question of whether the person knows he is using method *M*, when he believes he is via that very method *M*? If he knows he is, then his belief that he is tracks the fact that he is, and varies with that fact. To determine whether it so varies, we must look to the question of what he would believe if *p* were false, that is, if he weren't using method *M*. How are we to understand this question? It seems we must hold fixed the method *M* via which he believes, in order to reach the correct answer about knowledge (as is shown by the case of the grandmother), and that we cannot hold the method *M* fixed, for then we have the (apparently) incoherent supposition that he

is applying the method to the situation where he is not using it, in order to determine whether or not he is—and this supposes that he both is and isn't using the method.

This problem does not arise when we know via another method that we are using some particular method; it arises only for our knowledge of our use of our deepest methods, though not for shallower specifications of these methods in specific instances. Still, what should we say about our knowledge of these deepest methods or of the conditions in which we apply them? Do you know you are rational, do you know you are sane? If you were irrational or insane, mightn't you think you were rational and sane? Yes, but not by applying methods under (fixed) conditions of rationality and sanity. We cannot conclude simply that condition 3 is not satisfied so you don't know you are rational or sane; for that condition is not satisfied only when the method is allowed to vary. It would be best to be able coherently to discover whether or not that method is being used. I can use M to discover whether you are using M (if you weren't, I wouldn't believe, via M, that you were), or whether I was using M in the past (if I hadn't been, I wouldn't now believe, via M, that I had been). The difficulty is to make sense of saying that M, if currently used, would detect that it was not being used (if it weren't). And while I do not think this simply is incoherent, neither is it pellucidly clear.²⁰

Questions about knowing one is rational or sane need not depend on varying the method used. If what we have to go on as we apply methods is the appearance of rationality and sanity, then mightn't we appear sane and rational to ourselves even if we are not? So how do we know we are? We do have more to go on than how we appear to ourselves; there also is the agreement with others. Let us leave aside the possibility that all those others also might be insane and irrational, or be engaged in a plot to convince me (falsely) that I was rational and sane. Neither of these is what (actually) would or might occur if I weren't rational or

same. Might an insane and irrational person also be mistaken about whether others are agreeing with him, though, interpreting their disagreement as concord? If a person were insane or irrational in this way then others would appear (to him) to agree with him, and so he would appear sane and rational to himself. Things would appear qualitatively indistinguishable to him from the situation where he rationally and sanely judges the world. There appears to be no shift in method here, at least insofar as how using the method is experienced internally by the user. Do you know, then, that you are not in that particular skeptical situation SK? Perhaps not, but (as before) from our not knowing that particular not-SK it does not follow that we don't know other things, including that we are being sane and rational in particular situations in particular ways. For if we weren't, we wouldn't believe we were; if we weren't then sane and rational in those particular ways, what would or might obtain is not this skeptic's possibility SK. These points emerge even more clearly if we consider positions skeptical not about (almost) all knowledge in general, but about particular kinds of knowledge.

NOTES

1. There is an immense amount of literature concerning skepticism. See, for example, Sextus Empiricus, *Writings* (4 vols., Loeb Classical Library, Harvard University Press, Cambridge); Richard Popkin, *History of Skepticism from Erasmus to Descartes* (rev. ed., Humanities Press, New York, 1964); Arne Naess, *Skepticism* (Humanities Press, New York, 1968); René Descartes, *Meditations on First Philosophy* (Liberal Arts Press, New York, 1960); G. E. Moore, "Proof of an External World", "A Defense of Common Sense", "Certainty" [Pp. 257–271 in this volume. Ed.] and "Four Forms of Skepticism" in his *Philosophical Papers* (Allen and Unwin, London, 1959); J. L. Austin, "Other Minds" in his *Philosophical Papers* (Oxford University Press, 1961); Ludwig Wittgenstein, *On Certainty* (Basil Blackwell, Oxford, 1969); Keith Lehrer, "Why Not Skepticism?" (in Swain and Pappas, eds., *Essays on Knowledge and Justification* (Cornell University Press, 1978), pp. 346–363) [Pp. 227–237 in this volume.]; Peter Unger, *Ignorance* (Oxford University Press, 1975), pp. 7–24 [Chapter 3 reprinted in this volume. Ed.]; Michael Slote, *Reason and Skepticism* (Allen and Unwin, London, 1970); Roderick

Firth, "The Anatomy of Certainty", *Philosophical Review*, Vol. 76, 1967, pp. 3-27; Thompson Clarke, "The Legacy of Skepticism", *Journal of Philosophy*, Vol. 69, 1972, pp. 754-769; Stanley Cavell, *The Claim of Reason* (Oxford University Press, 1979).

2. From the perspective of explanation rather than proof, the extensive philosophical discussion, deriving from Charles S. Peirce, of whether the skeptic's doubts are real is beside the point. The problem of explaining how knowledge is possible would remain the same, even if no one ever claimed to doubt that there was knowledge.

3. Subjunctives with actually false antecedents and actually true consequents have been termed by Goodman *semi-factuals*. R is the semi-factual: $\text{not-}p \rightarrow S \text{ believes } p$.

4. Should one weaken condition 3, so that the account of knowledge merely denies the opposed subjunctive R? That would give us: $\text{not-}(\text{not-}p \rightarrow S \text{ believes } p)$. This holds when 3 does not, in situations where if p were false, S might believe p , and also might not believe it. The extra strength of 3 is needed to exclude these as situations of knowledge.

5. Though it does show the falsity of the corresponding entailment, " $\text{not-}p$ entails $\text{not-}(S \text{ believes that } p)$ ".

6. If a person is to know that SK doesn't hold, then condition 3 for knowledge must be satisfied (with "SK doesn't hold" substituted for p). Thus, we get

(3) $\text{not-}(\text{SK doesn't hold}) \rightarrow \text{not-}(S \text{ believes that SK doesn't hold})$.

Simplifying the antecedent, we have

(3) $\text{SK holds} \rightarrow \text{not-}(S \text{ believes that SK doesn't hold})$.

The skeptic has chosen a situation SK such that the following is true of it:

$\text{SK holds} \rightarrow S \text{ believes that SK doesn't hold}$.

Having the same antecedent as 3 and a contradictory consequent, this is incompatible with 3. Thus, condition 3 is not satisfied by the person's belief that SK does not hold.

7. Descartes presumably would refute the tank hypothesis as he did the demon hypothesis, through a proof of the existence of a good God who would not allow anyone, demon or psychologist, permanently to deceive us. The philosophical literature has concentrated on the question of whether Descartes can prove this (without begging the question against the demon hypothesis). The literature has not discussed whether even a successful proof of the existence of a good God can help Descartes to conclude he is not almost always mistaken. Might not a good God have his own reasons for deceiving us; might he not deceive us temporarily—a period which includes all of our life thus far (but not an afterlife)? To the question of why God did not create us so that we never would make any errors, Descartes answers that the motives of God are inscrutable to us. Do we know that such an inscrutable God could not

be motivated to allow another powerful "demon" to deceive and dominate us?

Alternatively, could not such a good God be motivated to deceive itself temporarily, even if not another? (Compare the various Indian doctrines designed to explain our ignorance of our own true nature, that is, Atman-Brahman's or, on another theory, the purusha's nature.) Whether from playfulness or whatever motive, such a good God would temporarily deceive itself, perhaps even into thinking it is a human being living in a material realm. Can we know, via Descartes' argument, that this is not our situation? And so forth.

These possibilities, and others similar, are so obvious that some other explanation, I mean the single-minded desire to refute skepticism, must be given for why they are not noticed and discussed.

Similarly, one could rescutinize the *cogito* argument. Can "I think" only be produced by something that exists? Suppose Shakespeare had written for Hamlet the line, "I think, therefore I am", or a fiction is written in which a character named Descartes says this, or suppose a character in a dream of mine says this; does it follow that they exist? Can someone use the *cogito* argument to prove he himself is not a fictional or dream character? Descartes asked how he could know he wasn't dreaming; he also should have asked how he could know he wasn't dreaming; he also should have asked how he could know he wasn't dreamed. See further my fable "Fiction," *Ploughshares*, Vol. 6, no. 3, Oct. 1980.

8. I say almost everything, because there still could be some true beliefs such as "I exist." More limited skeptical possibilities present worlds doxically identical to the actual world in which almost every belief of a certain sort is false, for example, about the past, or about other people's mental states. See the discussion below in the section on narrower skepticisms.

9. Let w_1, \dots, w_n be worlds doxically identical to the actual world for S. He doesn't know he is not in w_1 , he doesn't know he is not in w_2, \dots ; does it follow that he doesn't know he is in the actual world w_A or in one very much like it (in its truths)? Not if the situation he would be in if the actual world w_A did not obtain wasn't one of the doxically identical worlds; if the world that then would obtain would show its difference from the actual one w_A , he then would not believe he was in w_A .

However, probably there are some worlds not very different from the actual world (in that they have mostly the same truths) and even doxically identical to it, which might obtain if w_A did not. In that case, S would not know he was in w_A specified in all its glory. But if we take the disjunction of these harmless worlds (insofar as drastic skeptical conclusions go) doxically identical with w_A , then S will know that the disjunction holds. For if it didn't, he would notice that.

10. This argument proceeds from the fact that floating in the tank is incompatible with being at X. Another form of the skeptic's argument, one we shall consider later, proceeds from the fact that floating in the tank is incompatible with knowing you are at X (or almost anything else).

11. Note that I am not denying that $Kp \& K(p \rightarrow q) \rightarrow$ Believes q .

12. Here again I assume a limited readership, and ignore possibilities such as those described in James Blish, *Cities in Flight* (Winchester, Mass.: Faber and Faber, 1968).

13. Thus, the following is not a deductively valid form of inference.

$p \rightarrow q$ (and S knows this)
 $\text{not-}p \rightarrow \text{not-}(S \text{ believes that } p)$
 Therefore, $\text{not-}q \rightarrow \text{not-}(S \text{ believes that } q)$.

Furthermore, the example in the text shows that even the following is not a deductively valid form of inference.

$p \rightarrow q$ (and S knows this)
 $\text{not-}p \rightarrow \text{not-}(S \text{ believes that } p)$
 Therefore, $\text{not-}q \rightarrow \text{not-}(S \text{ believes that } p)$.

Nor is this one deductively valid:

$p \rightarrow q$
 $\text{not-}q \rightarrow r$
 Therefore, $\text{not-}p \rightarrow r$.

14. Does this same consequence of nonclosure under known logical implication follow as well from condition 4: $p \rightarrow S$ believes that p ? When p is not actually true, condition 4 can hold of p yet not of a q known to be entailed by p . For example, let p be the (false) statement that I am in Antarctica, and let q be the disjunction of p with some other appropriate statement; for example, let q be the statement that I am in Antarctica or I lost some object yesterday though I have not yet realized it. If p were true I would know it, p entails q , yet if q were true I wouldn't know it, for the way it would be true would be by my losing some object without yet realizing it, and if that happened I would not know it.

This example to show that condition 4 is not closed under known logical implication depends on the (actual) falsity of p . I do not think there is any suitable example to show this in the case where p is true, leaving aside the trivial situation when the person simply does not infer the entailed statement q .

15. Suppose some component of the condition, call it C' , also speaks of some cases when p is false, and when q is false; might it then provide "varies with", even though C' is preserved under known logical implication, and is transmitted from p to q when p entails q and is known to entail q ? If this condition C' speaks of some cases where $\text{not-}p$ and of some cases where $\text{not-}q$, then C' will be preserved under known logical implication if, when those cases of $\text{not-}p$ satisfy it, and p entails q , then also those cases of $\text{not-}q$ satisfy it. Thus, C' seems to speak of something as preserved from some cases of $\text{not-}p$ to some cases of $\text{not-}q$, which is preservation in the reverse direction to the entailment involving these, from $\text{not-}q$ to $\text{not-}p$. Thus, a condition that is preserved under known logical implication and that also provides some measure of "varies with" must contain a compo-

nent condition saying that something interesting (other than falsity) is preserved in the direction opposite to the logical implication (for some cases); and moreover, that component itself must be preserved in the direction of the logical implication because the condition including it is. It would be interesting to see such a condition set out.

16. Reading an earlier draft of this chapter, friends pointed out to me that Fred Dretske already had defended the view that knowledge (as one among many epistemic concepts) is not closed under known logical implication. (See his "Epistemic Operators", *Journal of Philosophy*, Vol. 67, 1970, pp. 1007–1023.) Furthermore, Dretske presented subjunctive condition for knowledge (in his "Conclusive Reason", *Australasian Journal of Philosophy*, Vol. 49, 1971, pp. 1–22), holding that S knows that p on the basis of reasons R only if: R would not be the case unless p were the case. Here Dretske ties the evidence subjunctively to the fact, and the belief based on the evidence subjunctively to the fact through the evidence. (Our account of knowledge has not yet introduced or discussed evidence or reasons at all. While this condition corresponds to our condition 3, he has nothing corresponding to 4.) So Dretske has hold of both pieces of our account, subjunctive and nonclosure, and he even connects them in a passing footnote (*Journal of Philosophy*, Vol. 67, p. 1019, n.4), noticing that any account of knowledge that relies on a subjunctive conditional will not be closed under known logical implication. Dretske also has the notion of a relevant alternative as "one that might have been realized in the existing circumstances if the actual state of affairs had not materialized" (p. 1021), and he briefly applies all this to the topic of skepticism (pp. 1015–1016), holding that the skeptic is right about some things but not about others.

It grieves me somewhat to discover that Dretske also had all this, and was there first. It raises the question, also, of why these views have not yet had the proper impact. Dretske makes his points in the midst of much other material, some of it less insightful. The independent statement and delineation of the position here, without the background noise, I hope will make clear its many merits.

After Goldman's paper on a causal theory of knowledge (in *Journal of Philosophy*, Vol. 64, 1967) [Pp. 6–16 in this volume.], an idea then already "in the air", it required no great leap to consider subjunctive conditions. Some two months after the first version of this chapter was written, Goldman himself published a paper on knowledge utilizing counterfactuals ("Discrimination and Perceptual Knowledge", *Journal of Philosophy*, Vol. 78, 1976, pp. 771–791), also talking of relevant possibilities (without using the counterfactuals to identify which possibilities are relevant); and Shope's survey article has called my attention to a paper of L. S. Carrier ("An Analysis of Empirical Knowledge", *Southern Journal of Philosophy*, Vol. 9, 1971, pp. 3–11) that also used subjunctive conditions including our condition 3. Armstrong's reliability view of knowledge (*Belief, Truth and Knowledge*, pp. 166, 169) involved a

lawlike connection between the belief that p and the state of affairs that makes it true. Clearly, the idea is one whose time has come.

17. Is it a consequence of our view that of two people who know p , each believing he knows p and satisfying condition 3 for knowing he knows p , one may know he knows and the other not, because (although identical in all other respects) the second might encounter skeptical arguments while the first somehow lives hermetically sealed from the merest brush with them?

18. Our task now is not to wonder whether it is legitimate to use M to reach a belief that M is being used. What, after all, is the alternate? Presumably, an infinite regress of methods, or a circle, or reaching a

method which is used but either is not believed to be used, or is believed to be though not via any method or way of believing.

19. Similar questions arise about our knowledge of other statements such that if they were false, we would not be using the methods via which we know they are true, for example, "there are eyes", "I am alive", "I am sentient", perhaps "I sometimes am tracking something".

20. Should we say for these cases discussed in the text that condition 3 does not apply, so that, as in the previous case of necessary truths, the whole weight of tracking devolves upon condition 4? The issue then simply turns on whether in similar situations where the person uses method M , he also would believe he does.

The Significance of Scepticism

BARRY STROUD

One of the topics announced for this symposium is the contrast between two different approaches or tendencies in philosophical studies of the foundations of science. On the one hand there are those who would abandon the quest for a general justification of empirical knowledge in favour of a purely naturalistic study of the procedures actually employed by scientists and other knowing subjects. On the other hand there are those who take seriously the challenge of philosophical scepticism and, seeing that it cannot be met by a straightforward Cartesian or "foundationalist" theory of knowledge, resort to so-called "transcendental arguments" to show that certain concepts or principles enjoy a privileged status in our thought because without them no human knowledge or experience would be possible at all.

These two alternatives are not exhaustive, and it is certainly impossible to say something definitive on the dispute between

them, but I want to discuss *one* important respect in which it seems to me that the second approach is superior to the first. Whatever one might think about the prospects for the success of transcendental arguments in epistemology, the need that is felt for them, and the consequent search for plausible candidates, seem to issue from a finer appreciation of the potential force of philosophical scepticism and of the considerations that have traditionally been thought to lead to it than is true of the first approach. And that is one important respect in which it seems to me that the second approach is superior, since I think we must somehow come to terms with the threat of philosophical scepticism if we are ever to get the kind of understanding we seek of the nature and possibility of human knowledge. In this paper I will try to make that claim as plausible as I can while remaining at the regrettably lofty level of discussing only general "strategies" or "positions" in the theory of knowledge. And even at that I will only be able to offer some questions to be discussed, or some issues to be raised, and not a set of philosophical theses to challenge, or even to compete with, some of the weighty doctrines