

I

KNOWLEDGE AND SCEPTICISM

R. NOZICK

You think you are seeing these words, but could you not be hallucinating or dreaming or having your brain stimulated to give you the experience of seeing these marks on paper although no such thing is before you? More extremely, could you not be floating in a tank while super-psychologists stimulate your brain electrochemically to produce exactly the same experiences as you are now having, or even to produce the whole sequence of experiences you have had in your lifetime thus far? If one of these other things was happening, your experience would be exactly the same as it now is. So how can you know none of them is happening? Yet if you do not know these possibilities don't hold, how can you know you are reading this book now? If you do not know you haven't always been floating in the tank at the mercy of the psychologists, how can you know anything—what your name is, who your parents were, where you come from?

The sceptic argues that we do not know what we think we do. Even when he leaves us unconverted, he leaves us confused. Granting that we do know, how *can* we? Given these other possibilities he poses, how is knowledge possible? In answering this question, we do not seek to convince the sceptic, but rather to formulate hypotheses about knowledge and our connection to facts that show how knowledge can exist even given the sceptic's possibilities. These hypotheses must reconcile our belief that we know things with our belief that the sceptical possibilities are logical possibilities.

The sceptical possibilities, and the threats they pose to our knowledge, depend upon our knowing things (if we do) mediately, through or by way of something else. Our thinking or believing that some fact *p* holds is connected somehow to the fact that *p*, but is not itself identical with that fact. Intermediate links establish the

Abridged by permission of the publishers and the author from *Philosophical Explanations*, Cambridge, Mass.: Harvard University Press/Oxford University Press. Copyright © 1981 by Robert Nozick.

connection. This leaves room for the possibility of these intermediate stages holding and producing our belief that p , without the fact that p being at the other end. The intermediate stages arise in a completely different manner, one not involving the fact that p although giving rise to the appearance that p holds true.

Are the sceptic's possibilities indeed logically possible? Imagine reading a science fiction story in which someone is raised from birth floating in a tank with psychologists stimulating his brain. The story could go on to tell of the person's reactions when he is brought out of the tank, of how the psychologists convince him of what had been happening to him, or how they fail to do so. This story is coherent, there is nothing self-contradictory or otherwise impossible about it. Nor is there anything incoherent in imagining that you are now in this situation, at a time before being taken out of the tank. To ease the transition out, to prepare the way, perhaps the psychologists will give the person in the tank thoughts of whether floating in the tank is possible, or the experience of reading a book that discusses this possibility, even one that discusses their easing his transition. (Free will presents no insuperable problem for this possibility. Perhaps the psychologists caused all your experiences of choice, including the feeling of freely choosing; or perhaps you do freely choose to act while they, cutting the effector circuit, continue the scenario from there.)

Some philosophers have attempted to demonstrate there is no such coherent possibility of this sort.¹ However, for any reasoning that purports to show this sceptical possibility cannot occur, we can imagine the psychologists of our science fiction story feeding it to their tank-subject, along with the (inaccurate) feeling that the reasoning is cogent. So how much trust can be placed in the apparent cogency of an argument to show the sceptical possibility isn't coherent?

The sceptic's possibility is a logically coherent one, in tension with the existence of (almost all) knowledge; so we seek a hypothesis to explain how, even given the sceptic's possibilities, knowledge is possible. We may worry that such explanatory hypotheses are *ad hoc*; but this worry will lessen if they yield other facts as well, fit in with other things we believe, and so forth. Indeed, the theory of knowledge that follows was not developed in order to explain how knowledge is possible. Rather, the motivation was external to epistemology; only after the account of knowledge was

¹ See Hilary Putnam, *Reason, Truth and History* (Cambridge, 1981), ch. 1.

developed for another purpose did I notice its consequences for scepticism, for understanding how knowledge is possible. So whatever other defects the explanation might have, it can hardly be called *ad hoc*.

I. KNOWLEDGE

Our task is to formulate further conditions to go alongside

- (1) p is true
- (2) S believes that p .

We would like each condition to be necessary for knowledge, so any case that fails to satisfy it will not be an instance of knowledge. Furthermore, we would like the conditions to be jointly sufficient for knowledge, so any case that satisfies all of them will be an instance of knowledge. We first shall formulate conditions that seem to handle ordinary cases correctly, classifying as knowledge cases which are knowledge, and as non-knowledge cases which are not; then we shall check to see how these conditions handle some difficult cases discussed in the literature.

One plausible suggestion is causal, something like: the fact that p (partially) causes S to believe that p , that is, (2) because (1). But this provides an inhospitable environment for mathematical and ethical knowledge; also there are well-known difficulties in specifying the type of causal connection. If someone floating in a tank oblivious to everything around him is given (by direct electrical and chemical stimulation of the brain) the belief that he is floating in a tank with his brain being stimulated, then even though that fact is part of the cause of his belief, still he does not know that it is true.

Let us consider a different third condition:

- (3) If p were not true, S would not believe that p .

Throughout this work, let us write the subjunctive 'if-then' by an arrow, and the negation of a sentence by prefacing 'not-' to it. The above condition thus is rewritten as:

- (3) not- p — not- $(S$ believes that $p)$.

This subjunctive condition is not unrelated to the causal condition. Often when the fact that p (partially) causes someone to believe that p , the fact also will be causally necessary for his having the belief—without the cause, the effect would not occur. In that case, the subjunctive condition (3) also will be satisfied. Yet this condition is not equivalent to the causal condition. For the causal

condition will be satisfied in cases of causal overdetermination, where either two sufficient causes of the effect actually operate, or a back-up cause (of the same effect) would operate if the first one didn't; whereas the subjunctive condition need not hold for these cases.² When the two conditions do agree, causality indicates knowledge because it acts in a manner that makes the subjunctive (3) true.

The subjunctive condition (3) serves to exclude cases of the sort first described by Edward Gettier, such as the following. Two other people are in my office and I am justified on the basis of much evidence in believing the first owns a Ford car; though he (now) does not, the second person (a stranger to me) owns one. I believe truly and justifiably that someone (or other) in my office owns a Ford car, but I do not know someone does. Concluded Gettier, knowledge is not simply justified true belief.

The following subjunctive, which specifies condition (3) for this Gettier case, is not satisfied: if no one in my office owned a Ford car, I wouldn't believe that someone did. The situation that would obtain if no one in my office owned a Ford is one where the stranger does not (or where he is not in the office); and in that situation I still would believe, as before, that someone in my office does own a Ford, namely, the first person. So the subjunctive condition (3) excludes this Gettier case as a case of knowledge.

The subjunctive condition is powerful and intuitive, not so easy to satisfy, yet not so powerful as to rule out everything as an instance of knowledge. A subjunctive conditional 'if p were true, q would be true', $p \rightarrow q$, does not say that p entails q or that it is logically impossible that p yet not- q . It says that in the situation that would obtain if p were true, q also would be true. This point is brought out especially clearly in recent 'possible-worlds' accounts of subjunctives: the subjunctive is true when (roughly) in all those worlds in which p holds true that are closest to the actual world, q also is true. (Examine those worlds in which p holds true closest to the actual world, and see if q holds true in all these.) Whether or not q is true in p worlds that are still farther away from the actual world is irrelevant to the truth of the subjunctive. I do not mean to endorse any particular possible-worlds account of subjunctives, nor am I committed to this type of account.³ I sometimes shall use it, though, when it illustrates points in an especially clear way.

² I should note here that I assume bivalence throughout this chapter, and consider only statements that are true if and only if their negations are false.

³ See Robert Stalnaker, 'A Theory of Conditionals', in N. Rescher, ed., *Studies*

The subjunctive condition (3) also handles nicely cases that cause difficulties for the view that you know that p when you can rule out the relevant alternatives to p in the context. For, as Gail Stine writes, 'what makes an alternative relevant in one context and not another? . . . if on the basis of visual appearances obtained under optimum conditions while driving through the countryside Henry identifies an object as a barn, normally we say that Henry knows that it is a barn. Let us suppose, however, that unknown to Henry, the region is full of expertly made papier-mâché facsimiles of barns. In that case, we would not say that Henry knows that the object is a barn, unless he has evidence against it being a papier-mâché facsimile, which is now a relevant alternative. So much is clear, but what if no such facsimiles exist in Henry's surroundings, although they once did? Are either of these circumstances sufficient to make the hypothesis (that it's a papier-mâché object) relevant? Probably not, but the situation is not so clear.'⁴ Let p be the statement that the object in the field is a (real) barn, and q the one that the object in the field is a papier-mâché barn. When papier-mâché barns are scattered through the area, if p were false, q would be true or might be. Since in this case (we are supposing) the person still would believe p , the subjunctive

(3) not- $p \rightarrow$ not- $(S$ believes that $p)$

is not satisfied, and so he doesn't know that p . However, when papier-mâché barns are or were scattered around another country, even if p were false q wouldn't be true, and so (for all we have been told) the person may well know that p . A hypothesis q contrary to p clearly is relevant when if p weren't true, q would be true; when not- $p \rightarrow q$. It clearly is irrelevant when if p weren't true, q also would not be true; when not- $p \rightarrow$ not- q . The remaining possibility is that neither of these opposed subjunctives holds; q might (or might not) be true if p weren't true. In this case, q also will be relevant, according to an account of knowledge incorporating condition (3) and treating subjunctives along the lines sketched above. Thus, condition (3) handles cases that befuddle the 'relevant alternatives' account; though that account can adopt the above subjunctive

in *Logical Theory* (Oxford 1968); David Lewis, *Counterfactuals* (Cambridge 1973); and Jonathan Bennett's critical review of Lewis, 'Counterfactuals and Possible Worlds', *Canadian Journal of Philosophy*, 4/2 (Dec. 1974), 381-402. Our purposes require, for the most part, no more than an intuitive understanding of subjunctives.

⁴ G. C. Stine, 'Skepticism, Relevant Alternatives and Deductive Closure', *Philosophical Studies*, 29 (1976), 252, who attributes the example to Carl Ginet.

criterion for when an alternative is relevant, it then becomes merely an alternate and longer way of stating condition (3).

Despite the power and intuitive force of the condition that if p weren't true the person would not believe it, this condition does not (in conjunction with the first two conditions) rule out every problem case. There remains, for example, the case of the person in the tank who is brought to believe, by direct electrical and chemical stimulation of his brain, that he is in the tank and is being brought to believe things in this way; he does not know this is true. However, the subjunctive condition is satisfied: if he weren't floating in the tank, he wouldn't believe he was.

The person in the tank does not know he is there, because his belief is not sensitive to the truth. Although it is caused by the fact that is its content, it is not sensitive to that fact. The operators of the tank could have produced any belief, including the false belief that he wasn't in the tank; if they had, he would have believed that. Perfect sensitivity would involve beliefs and facts varying together. We already have one portion of that variation, subjunctively at least: if p were false he wouldn't believe it. This sensitivity as specified by a subjunctive does not have the belief vary with the truth or falsity of p in all possible situations, merely in the ones that would or might obtain if p were false.

The subjunctive condition

(3) not- p — not- $(S$ believes that $p)$

tells us only half the story about how his belief is sensitive to the truth-value of p . It tells us how his belief state is sensitive to p 's falsity, but not how it is sensitive to p 's truth; it tells us what his belief state would be if p were false, but not what it would be if p were true.

To be sure, conditions (1) and (2) tell us that p is true and he does believe it, but it does not follow that his believing p is sensitive to p 's being true. This additional sensitivity is given to us by a further subjunctive: if p were true, he would believe it.

(4) p — S believes that p .

Not only is p true and S believes it, but if it were true he would believe it. Compare: not only was the photon emitted and did it go to the left, but (it was then true that): if it were emitted it would go to the left. The truth of antecedent and consequent is not alone sufficient for the truth of a subjunctive; (4) says more than (1) and (2). Thus, we presuppose some (or another) suitable account of

subjunctives. According to the suggestion tentatively made above, (4) holds true if not only does he actually truly believe p , but in the 'close' worlds where p is true, he also believes it. He believes that p for some distance out in the p neighbourhood of the actual world; similarly, condition (3) speaks not of the whole, not- p neighbourhood of the actual world, but only of the first portion of it. (If, as is likely, these explanations do not help, please use your own intuitive understanding of the subjunctives (3) and (4).)

The person in the tank does not satisfy the subjunctive condition (4). Imagine as actual a world in which he is in the tank and is stimulated to believe he is, and consider what subjunctives are true in that world. It is not true of him there that if he were in the tank he would believe it; for in the close world (or situation) to his own where he is in the tank but they don't give him the belief that he is (much less instil the belief that he isn't) he doesn't believe he is in the tank. Of the person actually in the tank and believing it, it is not true to make the further statement that if he were in the tank he would believe it—so he does not know he is in the tank.

The subjunctive condition (4) also handles a case presented by Gilbert Harman.⁵ The dictator of a country is killed; in their first edition, newspapers print the story, but later all the country's newspapers and other media deny the story, falsely. Everyone who encounters the denial believes it (or does not know what to believe and so suspends judgement). Only one person in the country fails to hear any denial and he continues to believe the truth. He satisfies conditions (1)–(3) (and the causal condition about belief) yet we are reluctant to say he knows the truth. The reason is that if he had heard the denials, he too would have believed them, just like everyone else. His belief is not sensitively tuned to the truth, he doesn't satisfy the condition that if it were true he would believe it. Condition (4) is not satisfied.

There is a pleasing symmetry about how this account of knowledge relates conditions (3) and (4), and connects them to the first two conditions. The account has the following form.

- (1)
- (2)
- (3) not-1 — not-2
- (4) 1 — 2

⁵ Gilbert Harman, *Thought* (Princeton; 1973), ch. 9, 142–54.

I am not inclined, however, to make too much of this symmetry, for I found also that with other conditions experimented with as a possible fourth condition there was some way to construe the resulting third and fourth conditions as symmetrical answers to some symmetrical looking questions, so that they appeared to arise in parallel fashion from similar questions about the components of true belief.

Symmetry, it seems, is a feature of a mode of presentation, not of the contents presented. A uniform transformation of symmetrical statements can leave the results non-symmetrical. But if symmetry attaches to mode of presentation, how can it possibly be a deep feature of, for instance, laws of nature that they exhibit symmetry? (One of my favourite examples of symmetry is due to Groucho Marx. On his radio programme he spoofed a commercial, and ended, 'And if you are not completely satisfied, return the unused portion of our product and we will return the unused portion of your money.') Still, to present our subject symmetrically makes the connection of knowledge to true belief especially perspicuous. It seems to me that a symmetrical formulation is a sign of our understanding, rather than a mark of truth. If we cannot understand an asymmetry as arising from an underlying symmetry through the operation of a particular factor, we will not understand why that asymmetry exists in that direction. (But do we also need to understand why the underlying asymmetrical factor holds instead of its opposite?)

A person knows that p when he not only does truly believe it, but also would truly believe it and wouldn't falsely believe it. He not only actually has a true belief, he subjunctively has one. It is true that p and he believes it; if it weren't true he wouldn't believe it, and if it were true he would believe it. To know that p is to be someone who would believe it if it were true, and who wouldn't believe it if it were false.

It will be useful to have a term for this situation when a person's belief is thus subjunctively connected to the fact. Let us say of a person who believes that p , which is true, that when (3) and (4) hold, his belief *tracks* the truth that p . To know is to have a belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it.

[A section which introduces some refinements is omitted here.]

D. SCEPTICISM

The sceptic 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 sceptic: 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 scepticism is unreasonable, since it is more likely that the sceptic'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-sceptical way. Even when these counter-arguments satisfy their inventors, they fail to satisfy others, as is shown by the persistent attempts against scepticism. The continuing felt need to refute scepticism, and the difficulty in doing so, attests to the power of the sceptic's position, the depth of his worries.

An account of knowledge should illuminate sceptical 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 sceptic 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 spectre of scepticism to rest without first hearing what it shall unfold.

Our goal is not, however, to refute scepticism, to prove it is wrong or even to argue that it is wrong. We have elsewhere 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 sceptic 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 sceptic, 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 possibilities the sceptic 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.⁶

Sceptical Possibilities

The sceptic 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 sceptic 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 sceptic's possible situation shows that (3) is false, and if (3) is a necessary condition for knowledge, then the sceptic's possible situation shows that there isn't knowledge.

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

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

This conditional, with the same antecedent as (3) and the contradictory consequent, is incompatible with the truth of (3). If (3) is true, then R is not. However, R is stronger than the sceptic 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

⁶ From the perspective of explanation rather than proof, the extensive philosophical discussion, deriving from Charles S. Peirce, of whether the sceptic'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.

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 sceptic 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 sceptic's possibility appears to show that (3) is false.

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; (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. 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 sceptic'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. 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 sceptic 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 sceptic 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 sceptical hypotheses.

C: not- p — SK does not obtain.

Any sceptical situation SK which satisfies condition C is ruled out. For a sceptical 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 — SK doesn't obtain).

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

The sceptic 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 .

Sceptical Results

According to our account of knowledge, S knows that the sceptic'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 sceptic 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).

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 sceptic, 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 sceptic asserts we do not know his possibilities don't obtain, and he is right. Attempts to avoid scepticism by claiming we do know these things are bound to fail. The sceptic'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. Nor has the sceptic 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 sceptic 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 sceptic 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 sceptic 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 non-actual situations doxically identical to the actual one. The sceptic, however, describes worlds doxically identical to the actual world in which almost everything believed is false.⁷

⁷ I say almost everything, because there still could be some true beliefs such as '1

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 sceptic's possibilities show. To be shown this is non-trivial, especially when we recall that we do not know the sceptic'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.

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

The sceptic 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 (non-trivial) 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 sceptic 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. 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

exist. More limited sceptical 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.

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 scepticism.

Yet we must grant that it appears that if the sceptic 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 sceptic'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.

Non-closure

In taking the 'short step', the sceptic 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 sceptic 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 sceptical 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 sceptic has argued and we have agreed) and so (he argues) you don't know the first. Another intuitive way of putting the sceptic'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.

No doubt, it is possible to argue over the details of principle P, to point out it is incorrect as 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 → q)*. Other similar principles face their own difficulties: for example, the principle that *K(p → q) → (Kp → Kq)* fails if *Kp* stops *p → q* from being true, that is, if *Kp → not-(p → q)*; the principle that *K(p → q) → K(Kp → Kq)* faces difficulties if *Kp* makes the person forget that *(p → 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 → q) & KKp → Kq*; this involves strengthening considerably the antecedent of P and so is not useful for the sceptic'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*' entails *q*' 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 sceptical 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 sceptic 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 sceptic's argument.

Principle P is wrong, however, and not merely in detail. Knowledge is not closed under known logical implication. *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 from 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.⁸ 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.

⁸ Here again I assume a limited readership, and ignore possibilities such as those described in James Blish, *Cities in Flight* (New York, 1982).

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 neighbouring city or town. If q were false, I would be floating 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. It is clear that any account that includes as a necessary condition for knowledge the subjunctive condition (3), not- p —not- $(S$ believes that $p)$, will have the consequence that knowledge is not closed under known logical implication.

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 non-trivial 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.

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 sceptic 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 sceptic 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 sceptic 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.⁹

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 sceptic is right that we don't track some particular truths—the ones stating that his sceptical 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 scepticism contains writers who endorse these sceptical arguments (or similar narrower ones), but confess their inability to maintain their sceptical beliefs at times when they are not focusing explicitly on the reasoning that led them to sceptical 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

⁹ 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*, 67, (1970), 1007–23.) Furthermore, Dretske presented a subjunctive condition for knowledge (in his 'Conclusive Reasons', *Australasian Journal of Philosophy*, 49, (1971), 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. The independent statement and delineation of the position here I hope will make clear its many merits.

After Goldman's paper on a causal theory of knowledge, in *Journal of Philosophy*, 64, (1967), an idea then already 'in the air', it required no great leap to consider subjunctive conditions. Some 2 months after the first version of this chapter was written, Goldman himself published a paper on knowledge utilizing counterfactuals ('Discrimination and Perceptual Knowledge', Essay II in this collection), also talking of relevant possibilities (without using the counterfactuals to identify which possibilities are relevant); and R. Shope has called my attention to a paper of L. S. Carrier ('An Analysis of Empirical Knowledge', *Southern Journal of Philosophy*, 9, (1971), 3–11) that also used subjunctive conditions including our condition (3). Armstrong's reliability view of knowledge (*Belief, Truth and Knowledge*, Cambridge, 1973, 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.

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 sceptics of various sorts have had such difficulties in sticking to their far-reaching sceptical conclusions 'outside the study', or even inside it when they are not thinking specifically about sceptical arguments and possibilities SK.

The sceptic's arguments do show (but show only) that we don't know the sceptic'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 sceptic'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 sceptic tracks such facts too, it is not surprising that when he focuses on them, on his relationship to such facts, the sceptic 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 sceptical belief and make it salient. However, this sceptical triumph is evanescent, it vanishes when his attention turns to other facts. Only by fixating on the sceptical possibilities SK can he maintain his sceptical virtue; otherwise, unsurprisingly, he is forced to confess to sins of credulity.