March, 1984

The Epistemic Predicament: Knowledge, Nozickean Tracking, and Scepticism

Steven Luper

Australasian Journal of Philosophy
Vol. 62, No. 1; March 1984

THE EPISTEMIC PREDICAMENT: KNOWLEDGE, NOZICKIAN TRACKING, AND SCEPTICISM*

Steven Luper-Foy

Possibly, you are a disembodied brain kept alive in a vat by scientists who are using electrodes to give you the experiences you now have. If you *were* in this or another of the scenarios described by sceptics, your sensory information would be exactly what it would be if you were not, and although you would have the sensory information you have now, many or most of your beliefs might be false. That you are in this scenario *is* a (logical) possibility, though admitting it seems like epistemic suicide. For once you do, can't the sceptic force you to conclude that you do not know most of what you believe even when you are *not* in the scenario?

The account of knowledge Robert Nozick has offered in *Philosophical Explanations*[1] allows him a striking reply to at least one sceptical argument that makes this link, namely, (every) one that assumes any version of the following principle: If a person *S* knows that *p* and believes that *q* by deducing it from *p*, then *S* knows that *q*. Given Nozick's tracking analysis, no version of this Principle of Closure is true. The sceptical argument must be rejected because all versions of the Principle of Closure must be rejected.

But whether Nozick's reply is acceptable depends on whether (something like) Condition (3) of his analysis is necessary. I will suggest that it is not, though the case against it is not strong. Moreover, once Nozick's analysis is replaced with a more adequate one that does not encorporate (3), we end up with an analysis that *supports* the Principle of Closure, so that Nozick's reply fails. Nonetheless, the sceptic's argument remains unsound. While the sceptic is safe in assuming the Principle of Closure, a further assumption required by the sceptic is false, namely, that if *S* were accidently correct in believing that *p* in one set of circumstances C, and if *S* were to arrive at that same true belief through the same causal sequence on the basis of the same evidence in distinct circumstances C', then *S* would be accidently correct in believing that *p* in C'.

* I wish to thank Robert Nozick for the many helpful comments and criticisms he has provided in conversations about the material in this paper. This paper would not have been possible without his generous help. I also want to thank Roderick Firth and Robert Shope for their comments, as well as my colleagues at Trinity University, especially Curtis Brown, Peter French and Jack Murphy.

[1] Cambridge: Harvard University Press, 1981, Section 3.

I. NOZICK'S THEORY

1. *The Tracking Account*

Though the root idea developed by Nozick is simple, refinements in it are necessary. I will introduce them in two stages.

The first version of Nozick's account, Version I, holds not only that

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

are necessary for S to know that p, but also that the following two conditions are necessary, while the four together are sufficient: if p were false, S would not believe that p; that is, using an arrow to abbreviate the subjunctive conditional,

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

and if p were true, S would believe that p; that is,

(4) $p \rightarrow S$ believes that p.

Condition (3) of Version I makes easy work of the typical Gettier case.[2] Consider the case in which S has (adequate) evidence that Mr Nogot, who is in S's office, owns a Ford, and on these grounds S believes that p: Someone in S's office owns a Ford. In fact, Mr Nogot does not own a Ford. Yet p is true because Mr Havit, who is also in S's office, owns a Ford. Nozick's third condition is violated. If p were false, either Mr Havit would not own a Ford or he would not be in S's office, and S would (or might) still believe that p on the same grounds as before.

Nonetheless, Version I is too strong an account of knowledge. Condition (3) says that if p were false S would not believe that p on any grounds or via *any method at all*; if p were false, for instance, not only would the evidence S actually relied on not deceive S, but no other evidence that might arise would deceive S either. This is counterintuitive. It is all right that S would still believe that p via some method even if p were false, *so long as S* would not believe it via the method S *actually* believes that p through.

Nozick offers the following illustration.[3] Suppose that a grandmother, watching her grandson cavort, believes that p: he is not dead. Presumably she knows that p. But suppose also that if he *were* dead, concerned relatives would deceive her into believing he is well (telling her lies, etc.). She still knows, but now Condition (3) is violated: even if the boy were dead, the grandmother might still believe that he is not dead.

Such cases[4] lead Nozick to offer Version II of his analysis in which he makes explicit reference to the 'methods' via which S believes that p:

[2] E. Gettier, 'Is Justified True Belief Knowledge?', *Analysis*, 1963, pp. 121-123.

[3] *Philosophical Explanations*, p. 179.

[4] Nozick is led (see p. 180) to complicate his analysis by another type of case. (3) rules out knowledge in any situation such that although S's belief via some method M_1 meets (3) and (4), S's belief via another method does not; nonetheless S's belief via M_1 outweighs S's belief via the others. In order to deal with cases of this sort, Nozick adds the following condition:

(1) p is true
(2) S believes that p via method M
(3) not-p & (S believes that p via M or S believes that not-p via M) $\rightarrow$ not-(S believes that p via M)
(4) p and (S believes that p via M or S believes that not-p via M) $\rightarrow$ S believes that p via M.[5]

Whether Version II can handle the Grandmother Case and various other examples, however, will need discussing.

2. *Problems with Tracking*

Conditions (3) and (4) function properly only when 'methods' are restricted to (what Nozick might call) two-sided methods.[6] Two-sided methods are ones capable of recommending the belief that not-p as well as the belief that p. We might, e.g., understand the method used in the Grandmother Case as follows:

(a) If: (S believes that) she has percepts of Type P,
Infer: p
(b) If: (S believes that) she has percepts of Type P',
Infer: not-p
(c) Otherwise, infer nothing about the truth value of p.

However, some methods, such as one consisting of just (a) and (c) above, are one-sided, i.e., incapable of recommending the belief that not-p.

The problem that arises when we apply (3) to a one-sided method M is that (3) can never be met. Its antecedent requires that S come to believe that p or that not-p via M. But in order to do so when M is one-sided, S must come to believe that p via M. And then the consequent of (3) cannot be satisfied. Thus if the grandmother believes the boy is alive via a one-sided method, then (counterintuitively) she does not know.

The opposite problem arises when we apply (4) to a one-sided method M. In order for such a method to satisfy the antecedent of (4), S *must* come to believe that p via M, so that the consequent of (4) is automatically satisfied.

A new Version III in which (3) and (4) of Version II are altered can ward off these difficulties (but others rush in):

---

(5) All other methods (besides M) via which S believes that p and which fail to satisfy (1)-(4) are outweighed by M (or at least, M is not outweighed by them).

But are these really cases of ignorance? Why not say that S knows so long as S has one *good* method? Some people who believe things on the basis of excellent evidence also give greater credence to gurus of one sort or another.

Another reason why it is best to drop (5) is because it involves a 'conditional fallacy' of the type discussed by R. Shope (in 'The Conditional Fallacy in Contemporary Philosophy', *Journal of Philosophy*, 75 (1978), pp. 397-413). Perhaps when M recommends the belief that not-p and M' the belief that p, S (in the circumstances at hand) follows M for the wrong reasons or, for the wrong reasons, fails to follow M.

[5] See his p. 179.

[6] See *Philosophical Explanations*, p. 183.

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

(4) p & S applies M $\rightarrow$ S believes that p via M.[7]

The grandmother meets (3) (as well as (4)) regardless of whether her method is the one- or the two-sided version suggested above. Neither would lead her to believe that p if it were false; rather, she would believe that p via the lies of her relatives if it were false.[8]

The notion of a 'method' of belief still needs clarification. Suppose that although S took a drug that produces hallucinations of hammers, S now sees a real hammer and believes that p: There is a hammer in front of S. Intuitively, S does not know that p. But whether Nozick's analysis has the favoured result depends on how we interpret S's method of belief.

We could describe the method M via which S believes as a causal theorist might. It is the chain of events usually involved in visual perception, including the stimulation of S's retinas by light from the hammer, nerve impulses, brain activity, etc. Then, counterintuitively, S knows that p according to Nozick's conditions. Condition (3) in particular is met since if there were no hammer, S would not believe there was *via M*; rather S would believe it, if at all, via some sequence of events involving hallucinated percepts.

If, on the other hand, we give the method a phenomenalist interpretation, then S does not know that p. Such a method M' would differ from M in excluding the events leading up to the production in S of percepts, and that makes all the difference. It is specified solely in terms of S's percepts, S's belief(s), and relations among these two. M' might, for example, be the following:

If: (S believes that) S has percepts of Type P
Infer: There is a hammer in front of S.

S's belief via M' fails Condition (3): if there were hammer, S might well still believe there was via (hallucinated) percepts of Type P.

Though a phenomenalist interpretation of methods would allow Nozick to deal properly with the Hammer Case, other cases are more intractable. Knowledge-yielding mechanisms that do not *involve* percepts cannot be interpreted phenomenally.[9]

[7] A problem with III (4) that remains, however, is that it cannot handle methods that occasionally fail to indicate anything, but are always correct when they do indicate something, e.g., when M is, roughly, 'seek a proof of p', then S might know p via M although (4) is not met (since not-(S applies M $\rightarrow$ S believes that p via M) assuming that (S applies M but M indicates nothing $\rightarrow$ not-S believes that p via M)). Replacing (4) with

(4*) p & (M indicates that p to S or M indicates that not-p to S) $\rightarrow$ S believes that p via M

handles this problem, but reintroduces the problem of one-sided methods.

[8] Unless the relatives would deceive her by leading her to have the same type of percepts mentioned in (a). Where she would be deceived via the same method through which she originally arrived at her belief, Nozick would say she fails to know.

[9] Roderick Firth mentioned to me some years ago that a story by D. H. Lawrence in which a young boy is mysteriously able to predict the outcomes of horse races by riding a rocking horse can be construed as a case of a knower who is not justified in believing. Suitably elaborated, this might also be a case of nonperceptual knowledge.

An example is a reading device for blind children that works by directly producing beliefs in their brains. Percepts are not involved. Or imagine that in the year 2070 learning devices accelerate the learning process by producing beliefs directly. Couldn't people acquire knowledge in these ways?[10]

Such cases do not involve percepts in terms of which methods of belief can be specified. Consequently not all methods of arriving at knowledge can be interpreted phenomenally. Nor, as we have seen, can Nozick simply interpret all methods as causal sequences. Clearly he needs to restrict the contents of methods to percepts on some occasions, such as the Hammer Case, and not on others. When should the contents of methods be restricted to percepts?

Cases such as the Learning Device Case are examples of noninferential (and nonperceptual) knowledge. Hence what Nozick might do is to interpret methods phenomenally, as rules of inference involving percepts and beliefs, in all cases of inferential belief (as we did in the Grandmother Case), and causally in all cases of noninferential belief.[11] That he intends something like this is, moreover, indicated by his claim [12] that any two methods M and M' are identical if and only if the percepts involved in them are identical. This claim indicates that when methods involve percepts, they are to be identified and specified only in terms of percepts and relations among these and beliefs. (It does not imply that methods with components other than percepts and beliefs are identical to methods that do not have those nonperceptual components but which involve the same percepts and beliefs.) Another passage indicates that not all methods need involve percepts: '*Usually*, a method will have a final upshot in experience on which the belief is based. . . .'[13] These methods are, on my suggestion, ones involved in noninferential knowledge.

Given this interpretation, Nozick's is a powerful analysis, but still too weak. It cannot handle all cases of noninferential ignorance, including one that, as Nozick says, is ruled out by Version I of his analysis. Suppose that *S* is a brain in a vat on a distant planet and that scientists are giving *S* beliefs by direct stimulation. They choose which scenarios to give *S* by whim.[14] Thus for one week they might lead *S* to believe *S* is living on a tropical island,

---

Others have noticed cases similar to Firth's. See P. Unger's infallible but unjustified gypsy example in 'Experience and Factual Knowledge', *The Journal of Philosophy* 64.5 (1967), pp. 152-173.

[10] The inventors of the devices in these examples have percepts, and so we may grant that the *child* need not, so long as *someone* had percepts (and is suitably related to the child). But that would be beside the point, which is just that the child need not have percepts in order to know.

[11] In cases of noninferential belief, the clause '*S* believes that *p* via M' of (3) and (4) becomes 'M causes *S* to believe *p*', while the clause '*S* applies M' becomes 'M occurs'.

[12] *Philosophical Explanations*, pp. 184-185.

[13] Ibid., my emphasis.

[14] The mere fact that scientists are controlling what *S* believes does not make this an instance of ignorance. If the scientists give *S* accurate beliefs only, then, intuitively, *S* knows what *S* believes. (Think of the scientists and their apparatus as artificial sensory devices.)

etc. One day the scientists lead *S* to believe that *S in in* this vat situation. If the method via which *S* believes that *p* is the (nonphenomenalist) causal sequence the scientists initiate in order to get *S* to believe that *p*, then Nozick's conditions are met. Condition (3) is met since if *S* were *not* in such a vat, *S* would not believe *S* was, via *any* method. (4) is met since if *S* were to believe anything via that sequence it would be that *p*. True, there is no guarantee that the scientists *will* initiate the sequence. That depends on their mood. But *if* they did, *S* would believe that *p*, and that is all (4) requires.[15]

Tracking is not only possible in some cases of *non*inferential ignorance, such as this Vat Case, but also in some cases of *in*ferential ignorance. Suppose that Flip believes that *o*: The lights in the other room are on. After tossing a coin, he believes that *o* via the following irrational method: if (*S* believes that) *S* has percepts indicating that the result of *S*'s toss is heads, *o* is true; if (*S* believes that) *S* has percepts indicating that the result is tails, *o* is false. So far Flip's belief presents no threat to Nozick's account. Even if *o* were false, the result of his coin toss might, in ordinary circumstances, be heads, so that he will not know that *o*.[16]

But suppose Flip is not in ordinary circumstances. Though he is unaware of it, Jessica is watching his coin tosses, and is flipping on the lights whenever the result is heads, and off whenever the result is tails. Now he *does* meet Nozick's conditions without knowing. Condition (3) is met because if the lights were out, Flip would believe they were out, not on, via the coin toss method. Condition (4) is met since if the lights were on, he would believe that they were on via the coin toss method. (Rather than reject this example on the grounds that Jessica might be unreliable, replace her with an ultra reliable machine that is able to sense the result of Flip's tosses and turn on and off the lights.)[17]

As the Jessica Case brings out, Nozick's conditions allow knowledge via methods that are completely unreliable and irrational. This is particularly counterintuitive if our belief is arrived at by applying a rule of inference that

[15] Nozick's analysis does not properly handle necessary truths. As he points out (p. 186), Condition (4) stands alone as a condition on our knowledge of necessary truths since (3) will be met vacuously. But according to (4), *S* knows any necessary truth *n* which *S* believes via a method that, if it indicated anything at all, would indicate that *n* is true. This is plausible but there are counterexamples. Suppose *S*'s (one-sided) method is: If John is not both fat and not fat, then *n* is true, where *n* is the statement, $\sqrt{4} = 2$. Condition (4) is met, but we cannot know that $\sqrt{4} = 2$ on the basis of the fact that John is not both fat and not fat.

Nozick's distinction between M and *S*'s possibly fallible way of applying M does not help if we suppose that to apply M is either to move (successfully) from the belief that John is not both fat and not fat to the belief that *p*, or to no belief. To reach the belief that not-*p* is to apply some other method. If we do not make this supposition, then reaching the wrong recommendation would count as following a method. (And even if that were acceptable, we can always specify our example so that *S would* not reach the wrong recommendation.)

[16] A causal interpretation of *S*'s method will not alter *S*'s ignorance: even if *o* were false, the causal sequence leading to *S*'s belief that *o* might still occur.

Incidentally, a simple causal condition would eliminate this example, but very similar examples would elude (1)-(4) plus the causal condition as well as the justification condition.

[17] Nor is there a causal sequence responsible for *S*'s belief via which *S* fails Nozick's conditions. Jessica's actions are not part of the causal sequence leading to *S*'s belief; however, they would be performed in all of the relevant situations which might arise if *p* were false.

is irrational and unreliable but which would yield the correct results in the circumstances at hand. If we come to know something in circumstances that make it very difficult for us to make a mistake, then Nozick's conditions do not require that our methods be reliable. In recalcitrant circumstances, our method must be very reliable in order for us to track with it. But in ones that conspire to render Nozick's conditions true, the method can be as treacherous as you like.[18]

If we come to a true belief via a completely unreliable method such as by tossing a coin, our belief is correct only coincidentally even if we happen to be in cooperative circumstances. The reason is that as far as the method is concerned, it is a sheer coincidence that we are in cooperative conditions. Therefore, even when we are in (abnormally) cooperative ones, the processes through which we know must be at least as reliable as is required to eliminate the possibility of accidentally correct belief in normal circumstances.[19] (One way to attain this reliability, of course, is to *find out* that we are in cooperative circumstances.)

II. A CAUSAL INDICATOR ANALYSIS OF EMPIRICAL KNOWLEDGE

In view of the Jessica case, an account is needed that sets out the minimum level of reliability required of the processes through which we arrive at empirical knowledge. Such an account I will now attempt to provide, but first I wish to discuss other requirements the analysis must meet.

I should note that a somewhat similar case, C (on p. 190), is discussed by Nozick who agrees that it meets his conditions (even if supplemented with a justification and causal condition, I might add), yet only counterintuitively could be called a case of knowledge. In it, *S* believes that there is a vase in the box in front of *S*, by looking at an illuminated hologram of a vase. The box alternately displays the hologram and a real vase, but contains a special lever that is able to sense vases, so that if the hologram were displayed, a vase would be in the box. Nozick defends his position by saying that *S* can know things via methods about which *S* has false beliefs. The idea is that case C strikes us as an instance of ignorance only because in it *S* has used a method about which *S* believes falsehoods.

Surely it is true that we can know things through processes about which we believe falsehoods. But this fact cannot help show that *S* really does know in the Jessica case since there *S* has no false beliefs about *S*'s methods.

[18] An example that tends to show that Nozick's account is too strong is as follows (I have forgotten who originated this example): Surely *S* knows that *p*: *S* does not believe everything. But (3) seems to be violated:

(*) *S* believes everything $\rightarrow$ not-(*S* believes (*S* does not believe everything)).

For the left side of the conditional says that *S* believes everything, while the right says *S* does not. (Be careful not to export the negation from within the context of belief to without in order to get the true '*S* believes everything $\rightarrow$ *S* believes that *S* believes everything'.)

It is not possible for *S* to believe everything (as Jeffery Robbins has reminded me), however. So (*) may well be true vacuously through having an impossible antecedent.

[19] P. Unger treats the condition that it not be accidental that *S* is right about *p*'s being true as a complete analysis of knowledge in his article 'The Analysis of Factual Knowledge', *The Journal of Philosophy*, 65 (1968), pp. 157-170.

Be careful to distinguish the claim (1) that it is not accidental that *S* is right about *p*'s being true from the claim (2) that it is not accidental that *p* is true.

1. *Normal vs. Rigged Circumstances*

Consider the following two situations. In Situation (I), Sue Seesit arrives at the belief that *t*: A table is in front of her. She does so by the normal visual process, and visual stimulation is the only information relevant to her belief. That is, she has not touched the table nor asked a friend about the table, etc. Moreover, Sue's circumstances are entirely ordinary. No scientists are projecting table holograms into the room, nor has anybody sneaked in a papier-mâché table, etc.

In situation (II), as in (I), Sue again believes that *t* through the visual process. However, her pal has built a trap door under the table so that every so often the table drops out of sight into the basement, and when it does, a device projects a table hologram where the real table would have been. Still, Sue has not spotted the hologram. In fact, her mental condition is precisely as it would be in Situation (I).

In Situation (II), Sue does not know that *t*. She is ignorant in (II) simply because her belief, though correct, is correct accidentally. The visual process is reliable, but in rigged circumstances, ones laced with holograms and the like, it is not reliable enough to rule out the possibility of accidently correct belief. In these circumstances, extraordinary reliability is necessary. Thus cases such as (II) show two things: first, no matter how reliable a belief process is, it almost always can be used in circumstances that are so rigged that in them our belief is merely accidentally correct at best. Second, when a process is used in rigged circumstances, it must be especially reliable in order to yield knowledge.

Since a knowledge-yielding process must be extremely reliable in rigged circumstances, we seem forced to conclude that it must *always* be extremely reliable, no matter *what* our circumstances are, [20] and that the evidence we have for our beliefs must always be very strong. We seem forced to conclude that an analysis of knowledge can eliminate the possibility of accident only if it requires that our known beliefs be ones arrived at through very reliable processes. This conclusion is mistaken, however. As cases such as Situation (I) show, our knowledge-yielding processes often *need not be* very reliable. For in (I), it is no coincidence that Sue's belief is correct. Under normal circumstances such as those in (I), vision is reliable enough to eliminate the possibility of accident (and thus to yield knowledge).

The contrast between cases such as (I) and (II)[21] reveals a third point: the degree of reliability required of our knowledge-yielding processes is a function of the degree to which our circumstances are rigged. I submit that the

[20] In order to move from the premise that there are circumstances in which our knowledge-yielding processes must be extremely reliable (in order to eliminate accident) to the conclusion that our knowledge-yielding processes must *always* be extremely reliable (in order to eliminate accident), one must assume Principle CAB below.

[21] If you think Situation (II) is a case of knowledge, replace it with a version just like it except that Sue actually sees the hologram. If we construe the visual processes as beginning with light of a certain configuration striking the retina, then the same process is used in both situations, even though in one Sue knows while in the other she does not.

minimum level is set in unrigged circumstances: so long as a process is reliable enough to eliminate accident here, it can yield knowledge here. Occasionally, however, the level must exceed the minimum, but only when we are in *rigged* circumstances, and then it need only be great enough to ensure that our beliefs are not accidentally correct.[22] The Causal Indicator Analysis I will now sketch clarifies the sliding scale of reliability required of our knowledge-yielding processes. I begin with the analysis of noninferential knowledge.

2. *Noninferential Knowledge*

A belief that *S* has not inferred from another is nonaccidentally correct and thus noninferentially known if and only if a causal chain that meets certain reliability conditions joins it to an event or state of affairs that guarantees its truth. That is my analysis of noninferential knowledge in a nutshell. But let me spell it out carefully.

In order for *S* to know noninferentially that *p*, there must be a sequence of events F that meets, among others, the following condition:

(1) F causes (or sustains) *S*'s belief that *p*.

Notice that I have not required that F be the *only* cause of *S*'s belief. If *S*'s belief is causally overdetermined, then *S* may believe that *p* by hypnosis as well as by the visual process, e.g., and either causal sequence may meet the conditions I will impose.

A second condition that the sequence F must meet requires that F join *S*'s belief that *p* to an event that guarantees the belief's truth. Such an event must be one of F's *stages*, i.e., one of the (possibly momentary) time-sections of F. An approximation to what we need is the familiar causal condition, i.e., the requirement that the event (possibly a state of affairs) described by *p* be a stage of F (presumably the first stage, so that we can ignore all goings-on prior to *p*'s being true). On second thought, however, there are other events that guarantee the truth of *S*'s belief. Any event that is causally sufficient for *p*'s truth will do so, as will any event that would cause someone to believe that *p* only if *p* were true.[23] The latter suggests a way to put our second condition:

[22] Most versions of the Defeasibility Analysis of knowledge can be seen along these lines. The justification condition requires that *S*'s evidence have a minimum level of reliability. And (any version of) the restriction forbidding that *S*'s justification be susceptible to defeat by any truth requires that *S*'s evidence be especially reliable in *rigged* conditions, but no more than minimally reliable in unrigged ones. For rigged conditions will involve truths that defeat *S*'s justification, while unrigged ones will not.

The defeasibility approach, however, cannot handle noninferential knowledge, knowledge of beliefs that are not justified (as I argue in 'Defeasibilism and Unjustified Belief', unpublished). P. Klein in his important book *Certainty* (Minneapolis: University of Minnesota Press, 1981), p. 150, e.g., claims that beliefs that I have not inferred from other beliefs can be justifiably believed on the basis of themselves. But then we can have (self) justified (and other things equal, *known*) beliefs against which we have overwhelming evidence since, short of not-*p*, it is not possible to defeat the evidence *p* for *p*.

[23] M. Swain suggests a condition that involves the notion of a causally sufficient state of affairs in 'Knowledge, Causality, and Justification', reprinted in *Essays on Knowledge and Justification* (Ithaca: Cornell University Press, 1978), pp. 87-100.

(2) ($\exists$E) [E is a stage of F & (S') (E is a stage of a sequence that causes S' to believe that p) $\rightarrow$ p)].[24]

Condition (2) can be read as saying that one of F's stages must be an *infallible indicator* that p. This formulation exploits the fact that some events can be said to 'indicate' the occurrence of others. For example, a bulb's glowing *indicates* that current is flowing through the lamp cord. Similarly, the stimulation of S's retinas by light of a certain configuration, one event in the chain leading to S's belief that the lamp is on, indicates that the lamp is on. (1) together with (2) imply that p is true, hence the truth condition would be redundant.

Following Nozick, I will interpret the subjunctive conditional $p \rightarrow q$ in a stronger way than is usual. It is not automatically true when p and q are both true in the actual world. Rather, it is true if and only if q holds throughout the p-worlds that are near to the actual world, i.e., throughout the *p-neighbourhood* of the actual world.[25]

A belief that p via a sequence F that includes the fact that p meets (1) and (2) and is arrived at in an infallible way. But such a sequence F need not be reliable in all of the ways necessary for knowledge. For S's *access* to the fact that p is via *other* stages of F, and one or more of these may be the type of event that tends to occur even though no state of affairs of the type which p describes holds. For example, imagine a Situation (III) that is similar to the (rigged) Situation (II), but now Sue Seesit *sees the hologram* rather than the table. And her pal has added one more device to her room, namely a table detecting device which accurately displays a hologram when (and only when) a table is present.[26] Since the fact that t causes Sue to believe that t by causing the hologram to be displayed, Condition (2) is met in this case. But she does not know that t.

The trouble with (III) is that some of the stages of F involved in Sue's access to her infallible indicator are unreliable indicators that t, namely, those involved in Sue's seeing the hologram. Clearly, we must require that *all* of the stages of F be reliable indicators that t in some sense.

Let us try to clarify that sense with a requirement on each of F's concrete stages. Suppose we require of each of these *token* events that if it were to occur, p would hold:

(a) (E) [E is a stage of F only if (E occurs $\rightarrow$ p)].

E and *F* range over *token* events, not event *types*, so (a) requires only that

[24] Gimmicky ways of describing the event E so as to render (2) vacuously true can be dealt with by restricting the allowable descriptions of E to its intrinsic properties (see Condition (3) below).

[25] See D. Lewis, *Counterfactuals* (Cambridge: Harvard University Press, 1973), especially Chapter 1. See also R. Stalnaker, 'A Theory of Conditionals', N. Rescher, ed., *Studies in Logical Theory* (Oxford: Basil Blackwell, 1968), reprinted in E. Sosa, ed., *Causation and Conditionals* (London: Oxford University Press, 1975), pp. 165-179. On the notion of a neighbourhood of possible worlds, see R. Nozick's suggested emendation of Lewis' theory in *Philosophical Explanations*, footnote 8 on p. 174 (pp. 680-681).

[26] Versions of these table cases are discussed by R. Shope in *The Analysis of Knowing: A Decade of Research* (Princeton: Princeton University Press, 1983).

p hold throughout the closest worlds to the actual world in which each *token* stage of F occurs. Let us call the sort of reliability required of each event E by (a) *token-reliability*. (a) fails to handle Situation (III). Even the projection of a hologram is a token-reliable indicator that *t* *if* it occurs in certain circumstances. Those described in (III) are an example. For in all worlds close to the actual world in which that particular event of a hologram's being projected occurs, it is produced by the accurate projector that would have displayed it only if *t* were true.

A second tack (that also fails) is to require of each stage E of F that, on *most* occasions, in the close worlds to the actual world in which that *type* of event occurs, *p* would hold. But this will not do. First, the condition is easily met by gimmicky types of events. Consider the Situation (III) event that includes the projection of a hologram. Suppose we understand its type to be this: *caused by a table in front of Sue.* Then any event of the same type would occur only if *t* were true. A second problem involves the Situation (I) event of *light (configured to give the appearance of a table) striking a pair of retinas,* which is a perfectly acceptable stage in Sue's coming to know that *t* in (I). On most occasions, even if this type of event were to occur, *t* would be false. For *others* frequently see tables (in front of *them*) with the help of the same type of event even though *no* table is in front of Sue.

The first problem can be handled if we type F's stages in terms of their *intrinsic* properties. An intrinsic property *e* is, roughly, one such that possibly something *x* has *e* although no contingent thing wholly distinct from *x* exists.[27] Thus the property of *caused by a table in front of Sue* is not an intrinsic property of the event of a *hologram's being projected.* Moreover, no intrinsic property that event does have makes it the type of event that, on most occasions, would occur only if *t* were true.

The second problem can be handled if, instead of the fact that *p*, we speak of a state of affairs of the same type as *p* (or 'a *p*-type situation' for short). True, on most occasions it is the case that even if light configured to give the appearance of a table strikes retinas, *t* would be false. Nonetheless, usually that type of event would occur only if a state of affairs of the same *type* as *t* were to hold, namely, one in which some table is in front of some onlooker. As an approximation, then, we can say that a stage E of F is a reliable indicator that *p* if and only if on most occasions in the close worlds in which an event of the same type as E occurs, a *p*-type situation holds (where E is typed in terms of one or more of its intrinsic properties).

Situation (III) is now excluded since the causal sequence F leading to Sue's belief that *t* includes the projection of a hologram. Consider any of this event's intrinsic properties. On most occasions, in the close worlds in which an event occurs that shares one or more of these, no *t*-type situation holds.

[27] This definition of 'intrinsic' is offered (and rejected) by D. Lewis in 'Extrinsic Properties', forthcoming in *Philosophical Studies*. 'Loneliness' cannot be a property of a stage in a series, hence the fact that it is intrinsic by our definition does no harm. (Contingent objects are ones that, unlike numbers, do not exist in all possible worlds.)

In order to make the notion of a 'p-type situation' more precise, let me introduce a device that converts the statement p which S believes at time t via sequence F into one in which variables have been substituted for all indexicals that designate S and the time of S's belief. Let $p(S,t)/(S',t')$ be the result of replacing the indexical expressions S and t that occur in p with the variables S' and t', respectively. E.g., where p states that the table now beside Fred is near me, we convert to: The table beside Fred at t' is near S'.[28]

The next step is to relativise reliable indicators to a (more or less wide) community of belief-holders that includes S. To say that E is a reliable indicator that p is to say *roughly* that if a majority of the members of such a community relied on an E-type event in order to decide whether a p-type situation holds, they would be correct a majority of the time. I can make this notion more precise using the indexical removing device. E is a reliable indicator that p for S at t if and only if:

> For most persons (or for most members of S's doxastic community) S', the following condition holds at most times t':
> $(\exists e)$ {e is an intrinsic property of E & [in the close worlds in which $(\exists E')$ (E' has e & E' is a stage of a sequence that, at t', causes S' to believe that $p(S,t)/(S',t')$), $p(S,t)/(S',t')$ holds]}.[29]

We are now in a position to state our third condition for knowledge:

(3) (E) (E is a stage of F only if E is a reliable indicator that p).[30]

Conditions (2) and (3) operate to set a ground level on the reliability of the (entire) cause F of S's belief that p. However, we have seen that this minimum is often not enough. Recall the first two situations we described above in which Sue is seeing a table and believing that t: There is a table in front of her. Intuitively, we would say that Sue knows that t in (I) but not in (II). Unfortunately, our conditions do not back our intuitions. For in both situations the fact that t is itself part of a chain of events F that causes her to believe that t, so that Condition (2) is met. And (3) is met because the sequence F involved in the visual process, used in both cases, contains nothing but reliable indicators.

The problem with Conditions (2) and (3) that is exploited by (II) is this. Let T be the fact that t, the fact that there is a table in front of Sue. Consider the stages of F that occur after T. In circumstances similar to those Sue is

[28] All rigid designators remain in p. In order to know that the table is beside Fred, e.g., it is not enough that I be able to tell (in the sense specified by our conditions) when there are tables near just anyone. I must be able to tell when *Fred* has a table beside him.

[29] For *any* event E there will be certain intrinsic properties that it alone is capable of having, namely, properties e' such that, for any event E', necessarily E' has e' only if E' and E are one and the same. Call these E's *unique* properties.

There may be a way to show that our third condition is too weak because it allows us to type events in terms of their unique properties. (In some cases, e.g., the token-reliability of an event may become sufficient for it to be a reliable indicator.) If so, a clause must be conjoined within (3) that restricts the properties e by virtue of which an event can be a reliable indicator: not-[(E') $\Box$ (E' has e only if E' = E)].

[30] I am indebted to Curtis Brown for helpful discussions about this condition.

in, this *T*-exclusive portion tends to indicate that *t* is true even if it is false. Having our retinas stimulated by light that gives the *appearance* that a table is present is a reliable guide to the presence of real tables, but in conditions abounding with holograms of tables, it will tend to mislead us. In such circumstances, the stages of F must be *even more* reliable than usual. We must supply a restriction requiring that all parts of F be reliable enough to rule out the possibility of accidently correct belief in the circumstances at hand.

Fortunately, we already know of a restriction that handles (II). We can apply (a), and require that each stage E of F other than T have *token-reliability* with respect to *t*'s truth, i.e., if E were to occur (in circumstances like *S*'s), *t* would hold. Notice that a stage's being an infallible indicator that *p* entails that if the stage occurred (in *S*'s circumstances), *p* would hold. Hence our requirement can be stated as follows:

(4) (E) [E is a stage of F only if (E occurs $\rightarrow$ *p*)].[31]

When E meets (4), we can say that E is a *conclusive indicator* that *p*.

Condition (4) rules out knowledge in Situation (II) since every stage of the T-exclusive sequence of events that constitutes seeming to see a table might occur in the circumstances there described even if not-*t*. This sequence would occur if Sue were looking at a hologram while the table was removed from the room. Moreover, Condition (4) is met in Situation (I), where Sue sees the table in normal conditions. Under those conditions, if any stage of F occurred, then *t* would be true.

3. *Inferential Knowledge*

S noninferentially knows that *p*, we have said, when and only when *S*'s belief that *p* is caused (or sustained) by a sequence of events each stage of which is both a reliable and conclusive indicator that *p*, and some stage of which is an infallible indicator that *p*. Our reliability and infallibility restrictions ensure that *S*'s belief will not be accidently correct in unrigged, cooperative circumstances. They set a ground level on the reliability of the causal indicators leading to *S*'s belief that *p*. And the conclusive indicator restriction makes sure that if the circumstances are uncooperative, then the causal indicators are exceptionally reliable, to an extent depending on just how uncooperative the circumstances are.

Our task now is to construct an account of inferential knowledge which, like our account of noninferential knowledge, eliminates the possibility of accidentally correct belief yet is sensitive to *S*'s circumstances. To do this is *not* to require that *S*'s evidence *e* for believing something *h* be so reliable that it would be able to eliminate the possibility of accidentally correct belief in any possible set of circumstances. Rather, it is to require that a sequence F via which *S* believes that *h* consist of proper causal indicators that *h*, ones that meet our conditions (2)-(4) for noninferentially knowing that *h* via F.

[31] Condition (4) is similar to the conclusive reasons restriction that Fred Dretske offered in 'Conclusive Reasons', reprinted in *Essays on Knowledge and Justification*, *op. cit.*, pp. 41-60.

Now, in order to meet these conditions, *S*'s evidence *e* must confer a minimal degree of warrant on *h* for *S*, and either *S*'s actual circumstances are cooperative, or else *S*'s evidence is especially good. Thus the possibility of accident is not eliminated by a condition requiring that *S* have strong evidence, but rather by ones that apply balanced pressure on all three of the factors involved in knowledge: belief, evidence and the circumstances.

Let us apply (1)-(4) to an example.[32] Suppose that Gene Geology sees solidified lava spread about in a certain way and on the basis of this belief and other beliefs (*e*), he believes that long ago a volcano erupted nearby (*h*). The fact that *h* (call it H) itself is part of the chain leading to L, the fact that lava is strewn about as Gene found it. If L is part of the chain leading to his belief that *e*, then *H* will be part also, so he will meet Condition (2). Each stage E of the sequence F_e leading to Gene's belief that *e* is such that, for most persons *S*' and times *t*', in the close worlds in which an event of the same type as E leads *S*' to an *h*-type belief, an *h*-type situation holds, so Condition (3) is met. Finally, Condition (4) is met since (in circumstances like his) *h* would hold if each stage E of F_e occurred.

(1)-(4) ensure that *S*'s inferential belief that *h* is not accidentally correct. Moreover, since inferring that *h* via an inference rule R is one stage of F, then our conditions require that *S* arrive at the belief that *h* via an inference rule R that is reliable enough to eliminate accident in normal circumstances (3) and, *if S* is in circumstances rigged in a certain way, they require that R be especially reliable (4). Inference rules function essentially by licensing leaps from one sort of belief to another, and one rule R_1 will be a more reliable rule for reaching the belief that *h* than another R_2 primarily because R_1 requires that *S*'s evidence be of greater quality than does R_2. Therefore, by requiring that *S* reach known beliefs by applying especially reliable rules when *S* is in rigged circumstances, our conditions require that the quality of *S*'s evidence increase directly as does the obstinacy of *S*'s circumstances.

However, a rule could meet our conditions and thus eliminate accident yet still be irrational.[33] Imagine a Situation (IV) that is just like the (normal) Situation (I) except that now Sue's belief that *t* is inferential and arrived at by following irrational inference rule R_t: If *S* (believes that *S*) is appeared to table-ly, *or S* seems to see 100 hungry martians eating the Empire State building, then infer that *t*; otherwise, infer nothing. R_t is a very reliable rule, and since Sue is in normal circumstances, Sue's belief that *t* by applying R_t meets our conditions and is not accidentally correct. It is reliable because on most occasions, the close worlds in which people believe that *t* via R_t are worlds in which they are appeared to table-ly and infer that *t*. On few occasions will there be close worlds in which they seem to see even *one* martian eating the Empire State building, so on few occasions will there be close worlds in which they falsely believe that *t* via R_t.

[32] This is A. Goldman's example. See 'A Causal Theory of Knowing', reprinted in *Essays on Knowledge and Justification*, *op. cit.*, p. 72.

[33] This is particularly easy if the rule is one for believing a necessary truth. *Any* rule will be infallible if it has us believe that a necessary truth is true and never has us believe that it is false. Yet such a rule might base our belief on another belief that is completely irrelevant.

Situation (IV) brings out the fact that merely requiring that our beliefs be correct but not accidentally, or (*a fortiori*) merely requiring that they be correct and arrived at in a reliable way, does not guarantee that they are known. Pure versions of reliabilism, whether restricted to the analysis of knowledge or applied to that of justified belief as well, are clearly wrong.[34] Some beliefs that are arrived at reliably and that are not accidentally correct could be irrational, and I presume it to be plain that irrational (as opposed to nonrational) beliefs are now known (though nonrational beliefs such as noninferential ones might be). For this reason we must add a justification condition to the analysis of inferential knowledge:

(5) *S*'s inference of *h* is rational.

(1)-(5) complete the analysis of inferential knowledge.[35]

Although the two analyses I have just provided improve on Nozick's, a troublesome matter remains to be resolved. Notice that Condition (4) is quite similar to the contraposition of Nozick's Condition (3). And like (4), (3) requires that a kind of reliability of our belief-forming processes vary with our circumstances. So offhand a version of Nozick's condition, say

(4') (E) [E is a stage of F only if (not-$p \rightarrow$ not-(E occurs))],

seems just as good as (4). Can anything be said in defence of (4) over (4')? The question is important because, as we will see, (4) allows us to know that we are not in the scenarios sketched by sceptics, while (4') does not. Since the issue of scepticism is affected by whether we choose (4) or (4'), it is best to postpone that choice until after we have clarified the effect the choice will have.

III. WHAT SCEPTICS DON'T KNOW REFUTES THEM

Sceptics usually argue by first concocting a scenario such that: if you were in it, (1) your sensory information would not be any different from the way it would be if you were not, and (2) some (often, depending on the scenario, most) of what you believe as a result of this information, no matter how extensive the information is, might be false even though you have that

[34] For example, see Goldman's relatively pure analysis in 'What is Justified Belief?', *Justification and Knowledge*, ed., G. Pappas (Boston: D. Reidel Company, 1979), pp. 1-25.

[35] A version of a case discussed by D. Armstrong in *Belief, Truth and Knowledge* (New York, London: Cambridge University Press, 1973), pp. 178-179, might appear to present a problem for my account. Suppose that the occurrence of belief chemical K in *S*'s brain is the direct cause of *S*'s belief that chemical K is present in *S*'s brain. Suitably elaborated, this is a case in which my conditions for noninferential knowledge are met. But is it a case of knowledge?

Any hesitation to call it a case of knowledge is due to our belief that we can only arrive at such beliefs as those involving brain chemistry inferentially (the case Armstrong himself discusses is one of inferential knowledge). And in *practice*, this is entirely true, so our knowledge of such matters is entirely inferential. But why deny that if K *could* directly cause us to believe that K is present in our brains, then we could know K is present noninferentially? Arguably, a similar sort of thing occurs when one of our beliefs causes us to believe we have it and therefore to *know* that we have it.

information. (By 'sensory information' I mean the signals received by a brain from a spinal cord.) Sceptics deal in scenarios that range from Descartes' tale of a deceitful demon (and its modern counterpart in which you are a brain in a vat being deceived through direct electrochemical stimulation) to the more modest Gettier case. From the fact that it is possible that you are in one of these scenarios, the sceptic tries to conclude that you do not know those of your beliefs that might be false if you *were* in one. But clearly this move cannot be made without additional assumptions.

I will lay out two versions of the sceptic's argument, including one offered by Nozick, then argue that there are reasons to prefer an alternative to Nozick's, but in any case the sceptic's argument is unsound.

1. *Sceptical Scenarios*

It is important to realise why it is that we arrive at the knowledge of a belief such as *t*: There is a table in front of *S*, using the visual process under normal conditions, such as in Situation (I) above, but not in *rigged* conditions, such as in Situation (II). Why we know that *t* in normal but not rigged conditions is important because the sceptic assumes that we know in normal ones only if we also know in rigged ones.

Contrasting normal and rigged conditions, we notice that whether or not *S* knows that *p* sometimes depends solely on features of *S*'s circumstances rather than on the process (whether a causal sequence and/or a chain of reasoning) through which *S* arrives at *S*'s belief that *p*, and rather than on whether or not *p* is true. Whether or not a sequence that yields knowledge in one set of circumstances C will also yield knowledge in distinct circumstances C' depends on whether or not a true belief resulting from that sequence would be accidentally correct in C'. Just because a true belief would be accidentally correct in C does not entail that it would be accidentally correct in C'. Clearly, the following metaphysical principle, that of the Closure of Accidental Correctness of Belief Under Change of Circumstances (CAB), is false:

CAB: If *S* were to arrive at the belief that *p* in circumstances C through process P and *S*'s belief were accidentally correct, and if in different circumstances C' *S* arrived at the true belief that *p* through the same process, then *S*'s belief that *p* in C' would also be accidentally correct.

It follows from the falsity of CAB that the following principle *CH* is also false:

CH: If *S* were to arrive at the belief that *p* in circumstances C through process P and *S* failed to know that *p*, and if in different circumstances C' *in which p is also true S* arrived at the belief that *p* through the same process P, then *S* would not know that *p* in C' either.

In fact, a stronger principle *CH'*, one that implies CH, is false, namely the result of deleting the italicised clause from CH. CH and CH' must be false since they entail CAB, which is false. I claim that it is CH or CH' that is the flawed assumption behind the sceptic's argument. An analysis of knowledge indeed must impose conditions that eliminate the possibility of accidentally correct belief in any possible set of circumstances in which we know. But the sceptic conflates this demand with the requirement that we arrive at our beliefs through processes that, when used in any possible set of circumstances, eliminate the possibility of accidentally correct belief.

Let the term 'sceptical scenario for *S* relative to *p*' refer to any situation such that if *S* were in it, (1) *S*'s sensory information would not be any different from the way it would be if *S* were not in it, and (2) something that *S* believes as a result of this information, namely *p*, might be false even though *S* has that information. The (or a) sceptic's argument consists of the following premises and conclusion:

(a) There are sceptical scenarios for *S* relative to what *S* believes by processing sensory information.
(b) If *S* were in a sceptical scenario, then *S* would not know anything that *S* believes by processing sensory information.
(c) CH or CH'
(d) Therefore, *S* does not know anything that *S* believes by processing sensory information.

CH or CH' is assumed because the sceptic is basing a conclusion about what *S* would know by processing sensory information in *one* set of circumstances on a claim about what *S* would know through that process in *another* set. The idea is that since *S* would not know much through that process in the sceptical scenario, then *S* would not know much through it in any other circumstances. Which principle sceptics assume will depend on the sceptical scenario. Some scenarios place *S* in a situation in which most of what *S* believes *is* false. Others, such as Gettier cases, place *S* in a situation in which some of *S*'s beliefs *might* be false for all *S*'s sensory information contributes to the matter, even though in fact they are not false. But since CH' implies CH, the sceptic always assumes CH.

It is usually crucial to the sceptic's argument that the process through which we arrive at most of our beliefs be some sort of processing of sensory information. The reason this is usually crucial is because the sceptic usually describes scenarios in which *no one could be* if he or she used other processes. Suppose that we described the process through which you arrive at your beliefs as a causal chain involving light from illuminated objects, retinal stimulation, spinal cord activity and finally the processing of sensory information. It is not possible to use that process in the most common sceptical scenarios. One common sceptical scenario, for example, is one in which you are a brain in a vat and scientists are feeding you sensory information through electrochemical stimulation. Light then could not be stimulating your retinas, etc., simply because you would not have any retinas to stimulate.

Nonetheless it would not *always* be necessary for the sceptic to limit the processes through which we arrive at our beliefs to the processing of sensory information. Even if we described the process as we did above, it is still possible to create scenarios in which the same process would be used though what you would be believing might be false. Thus, for example, Situation (II) could be used by the sceptic as a sceptical scenario to counter your claim to knowledge in (I). In a similar fashion, sceptical scenarios could be provided for each of most of *S* beliefs. The sceptic could conclude that most of what *S* believes *S* fails to know. There is no reason why all of the scenarios have to be compossible.

2. *Semiscepticism*

Must we assume CH or CH' in order to get to the conclusion (d), the claim that *S* does not know anything that *S* believes by processing sensory information? Nozick has provided an explanation of scepticism that gets there without assuming CH or CH'.[36] The crucial assumption sceptics make, a mistaken one according to Nozick, is their assumption of the Principle of the Closure of Knowledge Under Entailment (*CK*):

CK: For any *p* and *q*, if *S* knows that *p* and *S* believes that *q* by deducing that *q* from *S*'s belief that *p*, then *S* knows that *q*.

Now, *simply* adding CK to (a) and (b) does not allow the sceptic to get to (d). Instead of a sceptical conclusion, we could as easily (and as questionably) conclude (as did G. E. Moore in a similar context) that we know a great deal.[37] Pick some belief, say *h*, that *S* would not know in a sceptical scenario. In fact, let *h* be a statement which entails that *S* is *not* in a sceptical scenario. For example, let the scenario, call it *sk*, involve *S*'s being a brain in a vat on a planet near Alpha Centauri, and let *h* be the statement that *S* is at the

[36] Peter Klein offered other sceptical arguments that do not appeal to CH or CH' in his book *Certainty (op. cit.)*. These arguments are no more and, perhaps no less important than the one I have offered. The main difficulty with them is that, like Nozick's, they attack only inferential knowledge, while mine attacks noninferential knowledge as well. Sceptics certainly mean to attack both. A problem with Klein's *reply* to his sceptic's arguments is his claim that sceptics cannot argue for the premise that S does not know not-*sk* without begging the question. In fact, sceptics can simply claim that something like Nozick's Condition (3) is a necessary condition for knowledge (or justified belief), and that the belief that not-*sk* cannot meet it. The reply to the sceptic will then be similar to Nozick's (see below): given (3), the sceptic's Basic Maxim (p. 23) is false.

[37] *Philosophical Papers* (London: George Allen and Unwin, Ltd, 1959), p. 245. The Principle of Closure (of the Preconditions of Knowledge) that Moore himself adopts is as follows:

CPK: [(*S* knows that *p*) & [(*S* knows that *p*) entails *q*)] entails *S* knows that *q*.

From CPK and

(a) (*S* knows that *S* is standing) entails *S* is not dreaming

it follows that

(b) (*S* knows that *S* is standing) entails *S* knows that *S* is not dreaming.

And (b) is the contraposition of Moore's claim that if he does not know he is not dreaming, then he does not know that he is standing.

Incidentally, while CK, (a) and (b) do not rule out antiscepticism, they do not sanction it either. In order for antisceptics to conclude that *S does* know that not-*sk*, they must adopt an analysis of knowledge that makes it clear how such knowledge is possible.

Harding in Cambridge, Massachusetts. Why shouldn't sceptics adopt the *anti*-sceptical argument that since we know that *h*, and since we can believe that not-*sk* by deducing it from *h*, then we know that not-*sk*?

Sceptics cannot refuse to argue in this fashion on the grounds that we do *not* know that not-*sk*. For CK, (a) and (b) do not warrant that claim. In order to warrant it, sceptics would have to assume something like CH or CH'. If sceptics assumed CH, then they could claim that *S* does not know that not-*sk* because *S* would not know not-*sk* in the circumstances described by *sk*.

Another assumption that warrants the claim that we do not know not-*sk* is one that Nozick attributes to the sceptic. According to Nozick, sceptics begin by assuming that Nozick's third condition (call it *N*) is a necessary condition for knowledge:

(a') Where M is the method via which *S* believes that *p*, *S* knows that *p* only if
N: not-*p* $\rightarrow$ not-(*S* believes that *p* via M).

On the strength of (a'), the sceptics conclude that we do not know we are not in a sceptical scenario such as *sk*. In order for *S* to *know* that not-*sk*, N requires that if *S were* in a vat, *S*'s sensory information would not lead *S* to believe that *S* is not in a vat. But *S*'s sensory information *would* lead *S* to believe that *S* is not in a vat.

(b') not-(*S* knows that not-*sk*).

So far the sceptic has committed no error according to Nozick who will agree with the sceptic's claim that we do not know that we are not in sceptical scenarios. Nozick is a *semisceptic*. But the sceptic's assumption of CK is objectionable to Nozick.

(c') CK.

The rest of the sceptic's argument is a reductio against the claim that we know anything by processing sensory information. Take a claim such as *h* that is incompatible with *sk*, so that *h* entails not-*sk*, but which you think you know by processing sensory information. Suppose that

(d') *S* knows that *h*, and *S* believes that not-*sk* by deducing it from *S*'s belief that *h*.

Along with the assumption of CK at (c'), (d') entails that

(e') *S* knows that not-*sk*.

But *S* does *not* know that not-*sk*. That has already been established at (b') on the assumption of (a'). So *S* does not know that *h* after all.

(f') not-(*S* knows that *h*).

According to Nozick, the problem with this sceptical argument is that it assumes CK. If, as Nozick argues, N is necessary for knowledge, CK is *false*. Hence, the sceptic's argument is both inconsistent and based on a false premise.

What is to be made of Nozick's explanation of scepticism? It certainly is true that if *N were* a condition for knowledge, then the falsity of CK is virtually assured. The closure of knowledge under entailment virtually assures the closure of condition N.[38] That is, something like the following inference form must be valid:

(A) for any e and h,
1. not-e → not-(S believes that e via M)
2. S believes that h by deducing h from e (so that e entails h)

---

3. Therefore, not-h → not-(S believes that h by deducing h from e).

But (A) is not valid.

Consider an example.[39] S, looking at a striped animal in a cage marked 'zebra', believes that z: There is one and only one large animal in the cage in front of me, and it is a zebra. The method M used by S is this: If S (believes that S) has percepts of a single, large, striped animal in a cage marked 'zebra', infer z. S tracks z via M. If there were not a caged zebra, the cage would be empty or would contain some other animal, and S would not have the percepts necessary for M to indicate that z:

not-z → not-(M indicates that z).

Now suppose that S deduces from z the belief that not-m: There is not a mule in the cage that is cleverly disguised to look like a zebra. S does not track not-m. It is not the case that

m → not-(S believes not-m by inferring not-m from z)

since it is not the case that

m → not-(M indicates that z).

If there were a disguised animal, M would still indicate it is a zebra and S might still infer that it is not a mule.[40]

[38] In his comments on my paper 'What Sceptics Don't Know Refutes Them', given at the 1983 meeting of the Western Division of the A.P.A., Peter Klein pointed out that the conditions for knowledge could act jointly to sustain CK even though one of them is not closed. Let me offer a trivial example. The following analysis of knowledge as *contracking* sustains closure even though tracking does not: S contracks p iff:

S tracks p *and* S tracks every logical consequence of p.

[39] Borrowed from F. Dretske, 'Epistemic Operators', *The Journal of Philosophy* 67 (1970), pp. 1007-1023.

[40] It is noteworthy that (A) is invalid regardless of whether or not we give methods of belief phenomenal interpretations, *so long as* the cause of S's belief on which M operates does not include E, the fact that e, itself.

If the fact that e *were* a required part of what M operates on, and if 3 of (A) is rewritten 'not-h → not-(S believes that h by deducing h from *e which is believed via M*)', (A) would be valid since

(C) For any e and h,
1. e entails h
2. not-e → not-(E occurs)

---

3. Therefore, not-h → not-(E occurs)

is valid. The analysis of knowledge as tracking sustains closure when methods are interpreted as causal sequences that include the facts to be known. (A causal analysis that requires no more than that S's belief that h be caused by the fact that h has no trouble closing knowledge.)

3. *Nonscepticism*

Nozick is entirely correct in saying that CK would be false if N were a condition on knowledge, but he has not succeeded in providing a satisfactory explanation of scepticism and its flaw(s). The problem with Nozick's explanation is that it forces us to deny CK. We must reject the argument (a')-(f') not because it assumes CK, but because it assumes that among the necessary conditions for knowledge is one such as N given which CK is false.

How can I argue that no condition that opens knowledge is a necessary one? *Part* of the strategy is to provide an adequate analysis such as the Causal Indicator Analysis and then show that it closes knowledge. Momentarily I will show this. But first I must point out that this strategy is not enough. For as we noted earlier, the Nozickian counterpart (4') to our Condition (4) seems just as adequate. The two disagree only on various sceptical scenarios (such as the one involved in the Zebra Case). And to appeal to the sceptical cases in order to decide between (4) and (4') is to appeal to cases where we do not have clear intuitions pretheoretically. Moreover, neither lends any support to scepticism, so it is not possible to treat the claim that either does as a reductio against it. ((4') does not support scepticism simply because it fails to sustain CK.) The sceptic's argument fails regardless of our choice, and so we can make it without begging the sceptic's question.

In the final analysis, there is no *strong* reason to insist on (4) over (4') and hence no strong reason to retain CK. Only pragmatic considerations will settle the dispute. Appealing, e.g., to considerations of simplicity, we can adopt (4) and CK on the grounds that doing so allows us straightaway to know the logical consequences of what we know by deducing them. An epistemology containing (4) and CK is *simpler* than one that contains (4').

Keeping in mind that even if all goes well (4) is only slightly preferable to (4'), let me now argue that the Causal Indicator Analysis does sustain closure, and it does so in a manner that is not *ad hoc*. CK cannot be defended by providing an account that sustains closure in an *ad hoc* way simply because we could convert any account that leaves knowledge open into an *ad hoc* one that leaves it closed by disjoining closure to it. Consider Nozick's tracking analysis, for example. Suppose we say that *S distracks p* if and only if

> (*S* tracks *p*) or [*S* tracks *q* & *S* believes that *p* because *S* has deduced it from *q*].

An analysis of knowledge as distracking sustains closure. Moreover, *S* distracks everything *S* tracks and nothing else except what is entailed by what *S* tracks. So it is likely to handle the same problem cases.

Knowledge is closed for the following reason. Consider the oversimplified

---

When the facts to be known are not included, however, closure is not sustained by knowledge as tracking. Suppose in the Zebra example we say that the method *S* uses has *S* believe that *z* given the occurrence of the sequence of events which led *S* to believe that *z*, and we understood that the fact *z* is excluded from the sequence. Although these events would not have occurred if *z* had been false, they still would have occurred if *m* had been true.

version of the Causal Indicator Analysis such that S knows that h if and only if the sequence F through which S believes that h is such that

L: F occurs $\rightarrow h$,

where F is a sequence of events leading to S's belief that h. Analysis L sustains closure because the following is a valid inference form:

(B) For any e and h,
 (1) S knows that e, i.e.,
  (a) There is a sequence of events F leading to S's belief that e
  (b) F occurs $\rightarrow e$.
 (2) And e entails h and S believes that h by deducing that h from S's belief e.
 (3) Therefore, S knows that h, i.e.,
  (c) There is a sequence of events F' leading to S's belief that h
  (d) F' occurs $\rightarrow h$.

Consider (c). If we let F' be the sequence consisting of the sequence F that leads to S's belief that e, plus the sequence consisting of the events involved in S's deducing that h from e, then (c) follows from (a) and (2). Now consider (d). With negligible qualification, it follows from (b) and (2) [(2) includes the fact that e entails h] as a straightforward instance of strengthening the consequence (i.e., for any p, q, and r, $[(r \rightarrow p)$ & $(p$ entails $q)]$ entails $(r \rightarrow q))$.

(I ignore a trivial qualification of CK needed to handle a rare possibility: let g refer to the fact that S believes that h by deducing h from S's belief that e. Since it is possible that $g \rightarrow$ not-h, then (d) does not follow.)

(B) is valid because L is a restriction not just on h, but also on the consequences ch of h. For L requires that ch be true in every possible F occurs-world except those that are far from the actual world. Hence if S's belief that h meets L, and S deduces a consequence ch from h, then ch must be true in every F occurs-world that is not remote.

Of course, that a simplification of my analysis closes knowledge is no guarantee that my complete analysis does likewise. The simplification analyses knowledge in terms of only one of the causal indicators which we have discussed, namely, the infallible indicator. The simplification sustains closure only because infallible indicators sustain closure, that is, only because the following principle holds:

CII: (*ii* is a infallible indicator that e & e entails h) entails *ii* is a infallible indicator that h.

In order to show that my complete analysis closes knowledge, then, I need to show that each of the causal indicators in terms of which I analyse knowledge sustains closure.

Recall that a second required causal indicator is the conclusive indicator. That this indicator sustains closure can be seen when we notice that

CCI: (*ci* is a conclusive indicator that e & e entails h) entails *ci* is a conclusive indicator that h,

like CII, is just a straightforward instance of strengthening the consequence.

That leaves us with reliable indicators. The principle we need is

CRI: (*ri* is a reliable indicator that *h* & *h* entails *p*) entails *ri* is a reliab indicator that *p*.

To see that CRI holds, recall that *ri* is a reliable indicator that *h* only i for most persons *S'* at most times *t'*, if, at *t'*, *S'* relied on an *ri*-type eve in order to decide whether an *h*-type situation holds, usually one *would* hol But since *h* entails *p*, *that* an *h*-type situation holds entails a *p*-type one hol (i.e., for each *S'* at each *t'*, $h(S, t)/(S',t')$ entails $p(S, t)/(S', t')$). Hence a type situation will hold *whenever* an *h*-type one does. So if *S'*, at *t'*, reli on an *ri*-type event in order to decide whether a *p*-type situation holds, usual one would.

The last step is to show that the justification condition also sustains closur if *S* rationally believes that *e*, and if *S* believes that *h* by deducing that from *S*'s belief that *e*, then *S* rationally believes that *h*. To see that th principle is true, notice that when *S* deduces *h* from *S*'s belief that *e, e* available as sufficient grounds for *h*. *S* need not use *S*'s *evidence for e* as *S* (sole) grounds for *h*.[41]

Even if it is granted that analysis L sustains closure, however, we still nee an explanation of how it is *possible* to know such things as that we are n disembodied brains.

Suppose that I came by my belief that not-*sk* (I am not floating in a va by deducing it from my belief that *h* (I am in the Harding). And suppo I know that *h* through a causal sequence F that includes the stimulation my retinas by light from the engraved sign 'Harding' on the face of th building. L requires that the consequences of *h* be true in every possible occurs-world except those that are remote from the actual world. To dedu a consequence such as not-*sk* when we know that *h* is to believe it becau of a sequence F such that for each possible F occurs-world, either th consequent is true in it, or else it is a world that is remote.

Another way to express the sense in which we know that not-*sk* upo deducing it from *h* is this: our account says that given the origin of our beli that not-*sk* (i.e., a sequence including the cause of the belief that *h*), th possibility that *sk* is an *irrelevant alternative* to not-*sk*. Assuming that w know statements whose alternatives are either irrelevant or ruled out by u then we know not-*sk*. Our account would yield the following criteria fo relevant alternatives. First, *A* is an *alternative* to *h* if and only if *A* enta not-*h*. And secondly, *A* will be *relevant* if and only if not-(F occurs → no *A*), i.e., *A* might hold even if F occurred. According to these criteria, *s* is an irrelevant alternative to not-*sk* since F occurs → not-*sk*.[42]

Contrast analyses of knowledge such as Nozick's that leave it open. W can spell out a criterion for when an alternative *A* to *h* is relevant in term of Condition N (as Nozick does).[43] *A* is relevant if and only if not-(not-

[41] I owe this to P. Klein. See *Certainty*, Chapter 2.

[42] See also G. Stine, 'Scepticism, Relevant Alternatives and Deductive Closure', *Philosophic Studies* 29 (1976), pp. 249-261.

[43] *Philosophical Explanations,* p. 175.

$\rightarrow$ not-A). By this criterion, however, the negation of every statement p is a relevant alternative to that statement p. Thus the possibility that I am hovering over your head in an invisible spaceship watching you read is one that you must consider to be a relevant alternative to the fact that I am nowhere near you. (Intuitions about relevant alternatives will probably differ. If it were clear that not-*sk* is an irrelevant alternative to *sk*, then (given our observations about (4) and (4')) it would be clear that we should adopt CK and reject Nozick's semiscepticism.)

The plausibility of my nonsceptical explanation of scepticism is enhanced, I think, by the fact that if knowledge is analysed in a way strong enough for the sceptic's purposes, then CH' is compatible with CK.[44] Both are true if the sceptic analyses knowledge as follows: first we say that S believes that p through a *CAB process* if and only if S's belief that p via a process P is correct but not accidentally and there are no circumstances in which p is true and S's belief that p via P would be accidentally correct. Then we stipulate that S knows that p if and only if S believes that p through a CAB process. Deduction of p from something q that S believes through a CAB process is itself a CAB process. So sceptics would not be guilty of inconsistency if they assumed CH as well as CK, both of which are entailed by the analysis. Moreover, on such an analysis we would know exactly what sceptics claim we know; namely, almost nothing. In order to know something p through a CAB process, we would have to be correct about our belief that p non-accidentally in *every* possible set of circumstances whatever.

I have refuted the sceptic's argument by showing that it contains a false premise, namely, CH or CH'. Even the sceptic must acknowledge this. Otherwise the sceptic will be using the term 'know' incorrectly.

But let me emphasise that I have only refuted the sceptic's *argument*. Some think that in order to refute scepticism it is necessary to show that there *are no* (it is not *possible* for there to be) sceptical scenarios relative to what we believe by processing sensory information. Such a refutation cannot be had since clearly there *are* such sceptical scenarios. Others think that a refutation requires an argument from premises the sceptic will accept to the conclusion that we are not *in* a sceptical scenario. This, too, is not forthcoming simply because the sceptic allows only premises arrived at through processes relative to which there are no sceptical scenarios.

That neither refutation is possible, however, is beside the point, which is that there is not a sound argument from the existence of sceptical scenarios to the conclusion that we do not know what we believe by processing sensory information.

Trinity University, Texas

Received October 1982
Revised June 1983

[44] CK does not entail CH since the Causal Indicator Analysis rejects CH yet sustains CK. To see this, recall the simplified version I. (where F is a process through which S believes that p, if F occurs, p would hold). S knows that p in Situation (I) above but not in Situation (II), so CH is false given I.. As we have seen, however, CK is true.