

1985

What Skeptics Don't Know Refutes Them

Steven Luper

Trinity University, sluper@trinity.edu

Follow this and additional works at: http://digitalcommons.trinity.edu/phil_faculty



Part of the [Philosophy Commons](#)

Repository Citation

Luper, S. (1985). What skeptics don't know refutes them. *Pacific Philosophical Quarterly*, 65, 86-97.

This Article is brought to you for free and open access by the Philosophy Department at Digital Commons @ Trinity. It has been accepted for inclusion in Philosophy Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

WHAT SKEPTICS DON'T KNOW REFUTES THEM¹

BY

STEVEN LUPER-FOY

SKEPTICS usually argue that from the fact that (1) it is *possible* that you are in certain situations that I will call 'skeptical scenarios' where you would not know anything you believe by processing sensory information, it follows that (2) you do not know anything that you believe by processing sensory information no matter *what* circumstances you are in. Skeptical scenarios, which range from Descartes' deceitful demon to the modest Gettier case, are situations such that: if you were in one of them, (a) your sensory information would not be any different from the way it would be if you were not in it, and (b) some (often most) of what you believe as a result of this information might be false, even though you have that information.

I point out the assumption on which the skeptic's argument rests and argue that it is false, namely, CH': If, through process P, S were to arrive at the belief that *p* in circumstances C in which *p* is true and S failed to know *p*, and if in different circumstances C' (in which *p* is also true) S arrived at the belief that *p* through P, then S would not know that *p* in C' either. I also argue that Nozick's important explanation of skepticism in *Philosophical Explanations* (which does not attribute CH' to the skeptic) is not right since it assumes that the following Principle of Closure is false: If S knows that *p* and S believes that *q* by deducing it from *p*, then S knows that *q*.

* * *

Skeptics 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 information. (By 'sensory information' I mean the signals received by a brain from a spinal cord.) Skeptics deal in scenarios that range from Descartes' 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. Clearly it is possible that you are in one of those scenarios. It is also clear that if you *were* in one of them, then you would not know those of your beliefs that might then be

false. On the basis of these two facts, skeptics conclude that you do not know those beliefs.

I will point out the assumption that this skeptical argument requires. Then I will argue that the assumption is false, hence the skeptic's argument is unsound. I will also argue that Robert Nozick's fascinating attempt to explain the problem with skepticism in *Philosophical Explanations*² is not correct.

1. Skeptical Scenarios

Before attempting to locate the flawed assumption of the skeptic's argument, it is important to remember why it is possible for a person *S* to arrive at the knowledge of a belief such as *p*: There is a table in front of *S*, using the visual process under *normal* conditions but not in *rigged* ones. An example of rigged circumstances is a situation in which there is a trap door under *S*'s table so that every so often the table drops out of sight into the basement, and while it is in the basement a hologram of a table is projected where the real table would have been. It is important to understand why we know that *p* in normal but not in rigged circumstances because the skeptic assumes that we know in the former only if we also know in the latter.³

The reason we know in normal circumstances but not in rigged ones certainly need not be that in the one our belief is justified while in the other it is not. It is entirely possible to be justified in believing that *p* via a given process *P* in both normal as well as in rigged circumstances, though in the former we know via *P* while in the latter we do not. Virtually any Gettier-style case illustrated this point. The contrast between normal and rigged circumstances described above reveals 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 that process' inability to render *S*'s belief that *p* justified, 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'*.⁴ 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 a slightly qualified version of CAB that the following principle 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'* *S* arrived at the belief that *p* through the same process *P*, then *S* would not know that *p* in *C'* either.

For ease of reference I designate this principle *CH'* (short for "Principle of Closure of Knowledge Under Change of Circumstances"). In fact, a weaker principle, one implied by *CH'*, is false:

CH: If *S* were to arrive at the true 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.

Both principles *CH* and *CH'* must be false since they entail a modification of *CAB* which is false.

Before I state the weaker version of *CAB* entailed by *CH* and *CH'*, I must comment briefly on the analysis of knowledge. Elsewhere⁵ I have argued that the conditions for noninferential knowledge (i.e., knowledge of noninferential beliefs) are slightly different from the conditions for inferential knowledge. For noninferential knowledge to exist, it suffices that our belief be nonaccidentally correct. However, inferential knowledge requires not just that our belief be nonaccidentally correct but also that it be justified. One might put matters this way: knowledge has two components, a metaphysical and a doxastic component. The metaphysical component is the relationship between a person's circumstances and that person's belief by virtue of which it is the case that in those circumstances the belief is correct but not accidentally. The conditions for noninferential knowledge are completely exhausted by the metaphysical component of knowledge. However, there are additional conditions for inferential knowledge, conditions we might refer to as the doxastic component of knowledge. This component is the relationship between two or more beliefs and our inference rules by virtue of which a person justifiably (or rationally) arrives at the one belief by applying the rules to the other.

Assuming that these analyses of knowledge are correct, the following version of *CAB* is entailed by *CH'*:

CAB': If *S* were to arrive at the belief that *p* in circumstances *C* through process *P* so as to either justifiably or noninferentially believe that *p*, and if *S*'s belief were accidentally correct, and further, if in different circumstances *C'* *S* arrived at the true belief that *p* through the same process (and thus continued to justifiably or noninferentially believe that *p*), then *S*'s belief that *p* in *C'* would also be accidentally correct.

But a contrast of normal with the rigged circumstances described earlier not only shows that *CAB* is false, but also that *CAB'* is false.

That *CH* (and *CH'*) entails *CAB'* can be shown even if no particular construal of 'accidentally correct belief' is assumed.⁶ Suppose, to begin with, that *CAB'* is false, so that it is possible for situation *C'* to be one in which *S* correctly believes that *p* through *P* nonaccidentally and either justifiably or noninferentially, while the situation in *C* is that *S*'s belief that *p* is accidentally correct. Since knowledge entails nonaccidentally correct belief, the belief in *C* is unknown. However, since in *C'* *S* correctly believes that *p* through *P* nonaccidentally and either justifiably or noninferentially, then *S*'s belief that *p* in *C'* is known. Hence *CH* is false.

I claim that it is CH or CH' that is the flawed assumption behind the skeptic's argument. An analysis of knowledge must indeed impose conditions that eliminate the possibility of accidentally correct belief in any possible set of circumstances in which we know. But the skeptic 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 'skeptical scenario for *S* relative to *P*' refer to any situation such that if *S* were in it, (1) *S*'s sensory information would be no 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) skeptic's argument consists of the following premises and conclusions:

- (a) There are skeptical scenarios for *S* relative to what *S* believes by processing sensory information.
- (b) If *S* were in a skeptical 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.

Clearly CH or CH' is assumed because the skeptic 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 skeptical scenario, then *S* would not know much through it in any other circumstances. Which principle skeptics assume will depend on the skeptical 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 most 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 is implied by CH', the skeptic always assumes CH.

It is usually crucial to the skeptic's argument that the process through which we arrive at most of our beliefs be some sort of processing of sensory information. This is because the global scenarios that the skeptic usually describes (ones that are to call almost *all* of our beliefs into doubt) are scenarios in which people *could not be* if they 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 skeptical scenarios. One common skeptical scenario, for example, is one in which you are a bevvatted brain 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 skeptic 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, the rigged conditions described above could be used by the skeptic as a skeptical scenario to counter your claim

to knowledge in normal conditions. In a similar fashion, (Gettier-like) skeptical scenarios could be provided for each of most of *S*'s beliefs. The skeptic could conclude from the possibility of each of these *local* scenarios (ones that are to call few beliefs into doubt) that most of what *S* believes *S* fails to know. There is no reason why all of the scenarios have to be compossible.

2. Semiskepticism

Must we assume CH or CH' in order to get to the conclusion (d)? Recently Nozick has provided an explanation of skepticism that gets there without assuming CH or CH'. According to Nozick, skeptics begin by assuming that one of the conditions (call it *N*) of Nozick's analysis is necessary for knowledge:

(a') Where *M* is the method via which *S* believes that *p*,

S knows that *p* only if

N: not-*p* → not-(*S* believes that *p* via *M*).⁷

The arrow here signifies the subjunctive conditional. On the strength of (a'), skeptics conclude that we do not know we are not in a skeptical scenario. For example, consider the scenario (call it *sk*) in which *S* is a brain in a vat on a distant planet near Alpha Centauri. *S* does not know that not-*sk*. For suppose *S* comes to believe that not-*sk* through sensory information. 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 skeptic has committed no error, according to Nozick who will agree with the skeptic's claim that we do not know that we are not in skeptical scenarios. Nozick is a *semiskeptical*. But skeptics make a second assumption that is objectionable to Nozick. Namely, they assume the Principle of the Closure of Knowledge Under Entailment (CK):

(c') 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*.

The rest of the skeptic's argument is a *reductio* against the claim that we know anything by processing sensory information. Take a claim that is incompatible with *sk* but which you think you know by processing sensory information; for example, take *h*: *S* is at the Harding in Cambridge, Massachusetts. You cannot be in Cambridge if you are on a planet near Alpha Centauri, so *h* entails not-*sk*. 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 skeptical argument is that it assumes CK. If, as Nozick argues, *N* is necessary for knowledge, then CK is *false*. Hence, the skeptic's argument is both inconsistent and based on a false premise.

What is to be made of Nozick's explanation of skepticism? 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 the entailment virtually assures the closure of condition N . That is, something like the following inference form must be valid:

(A) For any e and h ,

1. $\text{not-}e \rightarrow \text{not-}(S \text{ believes that } e \text{ via } M)$
2. $S \text{ believes that } h \text{ by deducing } h \text{ from } e \text{ (so that } e \text{ entails } h)$

3. Therefore, $\text{not-}h \rightarrow \text{not-}(S \text{ believes that } h \text{ by deducing } h \text{ from } e)$.

But (A) is not valid.

Here is a counterexample to (A). Suppose that S , looking at a zebra in a cage marked 'zebra', believes that z : There is one and only one large animal in the cage in front of S , and it is a zebra.⁸ The method M used by S is

If: S (believes that S) has percepts of a single, large, striped animal in a cage marked 'zebra',

Infer: z .

S 's belief that z meets N . If there were no caged zebra, the cage would be empty or it would contain some other animal, and S would not have the percepts necessary for M to indicate that z . Now suppose that S deduces from z the belief that $\text{not-}m$: There is not a mule in the cage that is cleverly disguised to look like a zebra. S 's belief that $\text{not-}m$ does not meet N . If the animal were a disguised mule, M would still indicate it is a zebra and S would still infer that it is not a disguised mule.

Although difficult it is not impossible to defend CK if N is adopted. Even if N is included as a necessary condition for knowledge, by *some* (implausible but) possible analyses of knowledge CK is true. The most straightforward way for knowledge to be closed is for each of its necessary conditions to be closed on its own. So when one of these is *not* closed, knowledge cannot be either, except in a roundabout way. But in theory⁹ it is possible for knowledge to be closed even though some necessary conditions are not closed on their own. Closure might be the result of the force of the other necessary conditions, perhaps acting jointly. Consider a trivial example. When S meets Nozick's conditions for knowing that p , Nozick says that S tracks p . (For simplicity, let us focus on N alone, and suppose that S tracks p if and only if S 's belief that p meets N .) Let us say that S *contract*s p if and only if:

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

The analysis of knowledge as contracting is closed even though N is not. However, it should be apparent that virtually any analysis that includes N as a necessary condition will not sustain CK. Short of something like ' S tracks every logical consequence of p ' (the addition of which is unmotivated and leads directly to skepticism¹⁰), no additional condition will close knowledge once N is introduced.

Nozick is, then, correct for all practical purposes in saying that CK would be false if N were a condition for knowledge, but he has not succeeded in providing a satisfactory explanation of skepticism and its flaw(s). The problem with Nozick's explanation is that it forces us to deny CK when CK is true. Although we must reject the argument (a') – (f'), that is not because it assumes CK. We must reject it because it assumes that N is a condition of knowledge.

3. Nonskepticism

I want to argue that knowledge *is* closed under entailment, so Nozick's semi-skeptical reply to the skeptic must be rejected. Due to space limitations, I can only provide an oversimplified analysis that sustains closure in a manner that is not *ad hoc*.¹¹ The Principle of Closure 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 either conjoining or disjoining closure to it. Consider Nozick's analysis, for example. The analysis of knowledge as contracting conjoins closure to the tracking analysis. But closure can be disjoined as well. 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 analysis of knowledge such that *S* knows that *h* if and only if

L: $F \text{ occurs} \rightarrow h$,¹²

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) $F \text{ occurs} \rightarrow e$.

(2) And *e* entails *h* & *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) $F' \text{ occurs} \rightarrow h$.

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 that 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 [$r \rightarrow q$]).

An intuitive way of explaining why (B) is valid (and therefore why L closes knowledge) is this: L is a restriction not just on *h*, but also on the consequences *ch* of *h* (that is, everything entailed by *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.

Even if it is granted that analysis L sustains closure, however, we still need an explanation of how it is *possible* to know such things as that we are not disembodied brains.

Suppose that I came by my belief that not-*sk* (I am not floating in a vat) by deducing it from my belief that *h* (I am in the Harding). And suppose I believe

that h through a causal sequence F that includes the stimulation of my retinas by light from the engraved sign 'Harding' on the face of the building. I know that h . Consider what L requires of the consequences of h . The consequences of h must be true in every possible F occurs-world except those that are remote from the actual world. Knowledge is closed because given the origin of our beliefs, that what we deduce from them is false is only a remote possibility. To deduce one of these consequences when we know that h is to believe it because of an event F such that for each possible F occurs-world, either that consequent is true in it, or else it is a world that is remote. Take *not-sk*, one consequence of h . When I deduce *not-sk* from h , I know it in the sense that given the way I came by my belief that h , worlds in which *not-sk* is false are quite remote.

Another way to express the sense in which we know that *not-sk* upon deducing it from h is this: my account says that given the origin of our belief that *not-sk* (i.e., a sequence including the cause of the belief that h), the possibility that *sk* is an *irrelevant alternative* to *not-sk*. Assuming that we know statements whose alternatives are either irrelevant or ruled out by us, then we know *not-sk*. To see that *sk* is an irrelevant alternative to *not-sk*, notice that our account would yield the following criteria: first, A is an *alternative* to h if and only if A entails *not-h*. And secondly, A will be *relevant* if and only if $\text{not}-(F \text{ occurs} \rightarrow \text{not-}A)$, i.e., A might hold even if F occurred. According to these criteria, *sk* is an irrelevant alternative to *not-sk* since $F \text{ occurs} \rightarrow \text{not-}sk$.¹³

Contrast analyses of knowledge that leave it open; for example, knowledge as tracking. We can spell out a criterion for when an alternative A to h is relevant in terms of Condition N (as Nozick does¹⁴). A is relevant if and only if $\text{not}-(\text{not-}h \rightarrow \text{not-}A)$. By this criterion, however, the negation of *every* statement p is a relevant alternative to that statement p . Thus although I am not much inclined to think I am hovering over your head in an invisible spaceship watching you read, the possibility that I *am* is one that you must consider relevant. Yet that possibility is an irrelevant alternative to the fact that I am nowhere near you.

The plausibility of my nonskeptical explanation of skepticism is enhanced, I think, by the fact that if knowledge is analyzed in a way strong enough for the skeptic's purposes, then CH is compatible with CK. Let us say that S believes that p through a *CAB process* if and only if S 's belief that p via 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.¹⁵ Then the skeptic's analysis is this: 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 skeptics would not be guilty of inconsistency if they assumed CH as well as CK, both of which follow given this analysis. Moreover, on such an analysis we would know exactly what skeptics 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 nonaccidentally *while being in a skeptical scenario*.

In view of the strength of the analysis of knowledge as belief through a CAB process, it is unappealing as an analysis of knowledge. However, belief through a CAB process approximates a good analysis of *certainty*, I would argue. The root idea I would defend is that S is certain that p if and only if S believes that

p via a process P such that S knows that p via P in the actual world *as well as* in every possible world in which S believes that p via P .

Let us see what this conception of certainty requires when applied to a particular analysis of knowledge. Now in order for S to know that p , we said it is not quite sufficient that S 's belief be nonaccidentally correct. Knowledge has both a metaphysical as well as a doxastic component, so that a justification condition or rational belief condition is also necessary for inferential, though not for non-inferential, knowledge. But let us simplify and reduce knowledge to its metaphysical component, which requires nonaccidentally correct belief. Then the analysis of certainty that falls out is this: S is *certain** that p if and only if there is a process P such that S 's belief that p via P is correct but not accidentally and there are no circumstances in which S believes that p via P yet S 's belief is either false or accidentally correct.¹⁶

It is, I think, fruitful to interpret the notion of certainty involved in Descartes' *Meditations* along these lines. Thus something p is indubitable if and only if it is *certain** for S . So if p is indubitable for S , and if being subjected to the demon is compatible with using the CAB process at hand, then even if S were being subjected to the demon, S 's belief would still be true.¹⁷

I have refuted one version of the skeptic's argument by showing that it contains a false premise, namely, CH or CH'. My refutation focused on a confusion concerning the metaphysical component of knowledge. It does not address possible confusions that might arise concerning the doxastic component.

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

That neither refutation is possible, however, is beside the point. The point is that there is not a sound *argument* from the existence of skeptical scenarios to the conclusion that we do not know what we believe by processing sensory information. The possibility that we are in some skeptical scenario gives us no reason whatever to doubt our knowledge.

Trinity University
San Antonio, Texas

NOTES

¹I wish to thank Robert Nozick for the many things I have learned from him in conversations about his analysis. I also wish to express my gratitude for the comments of Curtis Brown and Steven Boër, a reviewer for *Pacific Philosophical Quarterly*. An early draft of this paper was read at the 1983 meeting of the Western Division of the APA. The present version is rather similar to part 3 of my paper "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism," forthcoming

in the *Australasian Journal of Philosophy*, Vol. 62, No. 1; March 1984. However, various claims made in the longer paper are expanded upon here.

²Cambridge: Harvard University Press, 1981, Section 3.

³If you think that *S* does know in the rigged circumstances I described, replace it with a version just like it except that *S* actually sees the hologram. If we construe the visual process as beginning with light of a certain configuration striking the retina, then the same process is used in both situations, even though in one *S* knows while in the other *S* does not.

⁴P. 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 "The Analysis of Factual Knowledge," *The Journal of Philosophy* 65 (1968), pp. 147–170.

⁵See "The Epistemic Predicament," op. cit.

⁶For an analysis of accidental correctness, see "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism." op. cit. The analysis of noninferential knowledge is my analysis of nonaccidental correctness. (Also see my "The Reliabilist Theory of Rational Belief," forthcoming in *The Monist*, where I improve my analysis of noninferential knowledge.)

⁷*N* is actually a (minor) revision of Nozick's third condition (see his p. 179). I argue that the revision is necessary in my paper "The Epistemic Predicament," op. cit.

⁸Borrowed from F. Dretske, "Epistemic Operators," *The Journal of Philosophy* 67 (1970), pp. 1007–1023.

⁹That the conditions for knowledge could act jointly to sustain CK even though one of them is not closed is a point I owe to Peter Klein, who commented on an earlier draft of this paper read at the 1983 meeting of the Western Division of the A.P.A. Responsibility for the example is mine.

¹⁰On pp. 240–241 of *Philosophical Explanations*, Nozick provides some general conditions that describe the analyses that skeptics and antiskeptics would adopt if they extended his tracking analysis. The conditions for skeptical analyses are met by the analysis of knowledge as contracting. Those for antiskeptical analyses are met by the analysis of knowledge as distracking (see below).

¹¹A complete defense of closure would require that I provide an analysis of knowledge which meets two requirements: first, it sustains closure in a way that is not *ad hoc*. Second, it must handle problem cases at least as well as an analysis such as Nozick's which leaves an opening in knowledge.

I explain the necessity of the first requirement in the text. The reason for the second requirement is this: if my analysis handles problem cases just as well as Nozick's but no better, and it also sustains closure, it is preferable on grounds of simplicity and intuitive appeal. An epistemology which is not complicated by rules governing when we can know via deduction is a simpler and more intuitive epistemology. Showing that an elaboration of *L* handles problem cases better than Nozick's analysis is a task I must leave aside for another occasion. See my "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism," op. cit. There I note that the case against *N* and hence against Nozick's explanation of skepticism is not strong.

¹²As I will understand *L*, it implies that *p* is true throughout the (*F* occurs)-worlds near to the actual world, throughout the (*F* occurs)-neighborhood of the actual world. (On the notion of a neighborhood of possible worlds, see Nozick's suggested emendation of David Lewis' theory in *Philosophical Explanations* footnote 8 on p. 174 (pp. 680–681).)

¹³See also G. Stine, "Skepticism, Relevant Alternatives and Deductive Closure," *Philosophical Studies* 29 (1976).

¹⁴See *Philosophical Explanations*, p. 175.

¹⁵Obviously, it is possible to define a type of certainty using principle CAB' as well. This notion is somewhat weaker, and hence more desirable, but still inadequate as an analysis of knowledge.

Incidentally, that *S* believes that *p* through a CAB process does not require that *S* have evidence *e* such that *e* entails that *p*.

¹⁶Note that certainty* appears to be a somewhat stronger notion than certainty defined as belief through a CAB process. For it appears to be possible for *S* to be certain that *p* (via *P*) even though in some circumstances *S*'s belief that *p* via *P* is false, while this clearly is *not* possible if *S* is certain* that *p* (via *P*). But given the following assumption, the two notions come to the same thing. Let us assume that any time there are circumstances in which *p* is true and others in which *S* believes that *p* via *P* yet *p* is false, then there are also circumstances in which *S*'s belief that *p* via *P* is accidentally correct. Then to eliminate circumstances in which the belief via *P* is accidentally correct, as both notions do, is also to eliminate circumstances in which it is false. (I wish to thank Curtis Brown for useful discussion of this assumption.)

But is the assumption true? So it would seem if the following claim is correct: if there are (not-*p*)-worlds in which *S* believes that *p* via *P*, and if in the actual world *S*'s belief that *p* is correct, then

some of those (not- p)-worlds are going to be close enough to worlds in which S correctly believes that p via P to render S 's belief in these worlds merely accidentally correct.

It may seem that the appeal made by both of these notions of certainty to the notion of accidental correctness is unnecessary. After all, it might be argued, it is not possible for there to be any circumstances in which S 's belief that p via P is accidentally correct if there are no circumstances in which S 's belief that p via P is false. So we may as well say that S is certain⁺ that p if and only if S believes that p via P and there are no circumstances in which S arrives at the belief that p via P yet p is false.

Surprisingly, the notion of certainty⁺ is weaker than the notions of certainty* and certainty. For suppose that P is a causal chain. In order for S to know that p via P , not only must P meet condition L , but so must each *link* in P (this is argued in "The Epistemic Predicament," *op. cit.*). But in saying that S is certain⁺ that p , we do not imply that these links meet L in all possible circumstances in which P is used.

¹⁷Incidentally, the hypothesis that there is a benevolent God is, I would argue, *itself* a skeptical scenario! Notice that this puts Descartes in the peculiar position of (unintentionally) trying to prove that one skeptical scenario (the benevolent God hypothesis) actually holds in order to prove that another (the malevolent demon hypothesis) does not.