

Trinity University

Digital Commons @ Trinity

Philosophy Faculty Research

Philosophy Department

2003

Indiscernability Skepticism

Steven Luper

Trinity University, sluper@trinity.edu

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



Part of the [Philosophy Commons](#)

Repository Citation

Luper, S. (2003). Indiscernability skepticism. In Luper, S. (Eds.), *The skeptics: Contemporary essays* pp. 183-202. Burlington, VT: Ashgate Publishing.

This Post-Print 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.

INDISCERNABILITY SKEPTICISM¹

Steven Luper, Trinity University

Ideally, our account of knowledge would help us to understand the appeal of (and flaws in) skepticism,² while remaining consistent with our ‘intuitions,’ and supporting epistemic principles that seem eminently plausible. Of course, we don’t always get what we want; we may not be able to move from intuitions and principles to an account that fully squares with them. As a last resort, we may have to move in the other direction, and give up intuitions or principles that are undermined by an otherwise compelling account of knowledge, so as to achieve ‘reflective equilibrium.’³ But last resorts come last.

As is well known, Fred Dretske and Robert Nozick devised accounts of knowledge that perform well as measured by two of these three expectations. Their ‘tracking’ theory (to apply Nozick’s term to both Dretske’s and Nozick’s analyses) seems intuitive and it allows us to offer an intriguing explanation of skepticism.⁴ However, their view forces us to reject the principle of closure, even qualified extensively. As far as I know, no one accepts the following stark version of the closure principle:

If S knows p , and p entails q , then S knows q .

This principle implies—falsely—that S knows q even if S does not believe q or realize that q follows from p . Nevertheless, the closure principle, suitably restricted, seems too

¹I thank Curtis Brown and Peter Klein for helpful suggestions about an earlier draft.

²Several people have emphasized this point or something like it. E.g., Robert Nozick (1981) writes, “an account of knowledge should illuminate skeptical arguments and show wherein lies their force” (p. 197).

³Goodman 1955; Rawls 1971, p. 20.

⁴Dretske 1970, Nozick 1981; cf. Goldman (1976), who develops a tracking account of perceptual knowledge.

obvious to give up; because tracking theorists reject it even when thoroughly qualified, several writers have attempted to improve upon the tracking theorists' approach.

In particular, contextualists have offered a way of coming to terms with skepticism without giving up the principle of closure. However, in my view we should not accept their approach, for its disadvantages outweigh its advantages. First, according to contextualists, thinking about skepticism destroys our knowledge. Only people who ignore skepticism know anything, and even *they* lose their knowledge if skeptical possibilities are brought up. Second, while contextualism is consistent with some plausible epistemic principles, it forces us to reject others, such as this metalinguistic version of the closure principle:

If person S correctly attributes to herself knowledge that *p*, then if S *had* believed *q* by deducing *q* from *p*, S *would have been* correct to attribute to herself knowledge that *q*.

Fortunately, there is no need to say that thinking about its implications destroys our knowledge, and no need to give up closure, even on the metalinguistic level. There is an intuitively compelling account of knowledge that helps us to understand the problem with—and appeal of—leading forms of skepticism, *and* that squares with plausible epistemic principles: the *indicator* analysis.⁵ We can meet our three expectations fully.

To defend these claims, I'll review the tracking and indicator analyses, and show that the latter handles skepticism at least as well as the former yet sustains closure. Then I'll consider the contextualist account, and suggest that it is inconsistent with plausible epistemic principles, and cannot handle skepticism as well as the indicator approach.

I. TRACKING V. THE INDICATOR ACCOUNT

Consider rough and somewhat oversimplified versions of the theories offered by Dretske and Nozick:

Dretske: A person S knows p if and only if there is a reason R such that:
S's belief p is based on the fact that R holds, and
 $\text{not-}p \rightarrow \text{not-}(R \text{ holds}).$ ⁶

Nozick: A person S knows p if and only if there is a method M such that:
 S believes p via M , and
 $\text{not-}p \rightarrow \text{not-}(S \text{ believes } p \text{ via } M).$ ⁷

Simplifying even more, I am going to ignore the differences between these two approaches, and lay out a bare-bones version of the tracking account of knowledge. S knows p if and only if there is a reason R such that (using 'T' for 'tracking theory'):

(T1) S's belief p is based on the fact that R holds

(T2) $\text{not-}p \rightarrow \text{not-}(R \text{ holds}).$

What I have set forth is the heart of the tracking view. But now consider a simple modification, in which we replace (T2). S knows p if and only if there is a reason R such that:

(1) S's belief p is based on the fact that R holds

(2) $R \text{ holds} \rightarrow p.$ ⁸

⁵Luper 1984, 1987; Sosa 1999.

⁶Dretske 1971 puts his condition as follows: where R is the basis for S 's belief that P , S knows that P is true if and only if "R would not be the case unless P were the case."

⁷Nozick formulates his account on p. 179 of Nozick 1981; in the text, I simplified his third condition, and ignored the fourth.

⁸Note that the truth condition would be redundant, since

$R \text{ holds}$

(which is implied by (1)), together with

$R \text{ holds} \rightarrow p$

Condition (2) is just the contrapositive of (T2). So (2) and (T2) must amount to the same thing, right?

Tracking v. Indication

No; in fact, the contrapositives of subjunctive conditionals are not equivalent, and while (2) and (T2) work similarly in many cases, they differ in crucial ways. Consider a case in which they converge and a case in which they diverge.

Convergence first. In one of his examples, Dretske asks you to imagine that you are at a perfectly ordinary wildlife show, standing in front of a cage marked 'zebra'. There are no Cartesian demons lurking in the wings, no Gettieresque surprises waiting to spring themselves at you. You are looking right at the zebra inside, and you come to believe z : the animal in the cage is a zebra. The source of your belief is a familiar empirical test by which you examine a scene and see if you get zebra-in-a-cage-type experiences; you take having these experiences to indicate that z is true, and not having them to indicate that z is false. You meet (T2): in the close worlds in which z is false, you do not have zebra-in-a-cage-type experiences. You also meet (2): in the close worlds in which you have these experiences, z is true.

Now divergence. Suppose you use zebra-in-a-cage-type experiences as your basis for believing *not-m*: the animal in the cage in front of me is *not* a mule cleverly disguised to look just like a zebra. That is, you take getting zebra-type experiences to indicate that *not-m* is true. You *fail* to meet (T2). (T2), applied to *not-m*, requires that

$m \rightarrow S$ does not get zebra-type experiences.

entail
 p .

And in the close worlds in which the animal is a zebra-like mule, you still have zebra-type experiences, leading you to think that the animal is *not* a zebra-like mule. Yet you *will* meet (2). Applied to *not-m*, (2) requires that

S gets zebra-type experiences \rightarrow *not-m*.⁹

In the actual world we are imagining, you are in an ordinary wildlife show looking at a caged zebra. In such a world, and in close worlds where you get zebra-like percepts, you are led to think the animal is not a zebra-like mule, and, of course, in these worlds the animal isn't.

So (T2) is not equivalent to (2), even though both are met in many of the same cases.

Tracking, Indication, and Closure

Perhaps (T2) and (2) are distinct, but isn't (T2) *preferable* to (2)? Won't Dretske and Nozick insist that seeming to see a zebra does *not* position us to know that we are not looking at a zebra-like mule, and that it only positions us to know that we are looking at a zebra? But when they say these things we know what comes next: they deny the principle of closure of knowledge under entailment, and say you cannot know that the animal in the cage is a not a made-up mule even if you *deduce* it from something you admittedly know—namely, that it is a zebra. Why do they reject closure? Consider condition (T2). Let us say that when a true belief meets (T2), it *tracks the fact* that *p* via

⁹The point of condition (2) is to require that *p* be true in the *R-holds*-neighborhood of the actual world. But, as Peter Klein reminded me, if we require that this neighborhood be very extensive, the condition is too strong. To use an example Jonathan Vogel (1987) used against Nozick, suppose I left a tray of ice cubes in the sun an hour ago, and believe, correctly, and for obvious reasons, that they have melted. Nonetheless, there is a world in which someone put the cubes back in the freezer without telling me; if we define the relevant neighborhood so that this world is in it, I fail to meet condition (2), for in that world I will still have my reasons for believing the cubes have melted, even though they are frozen. In my view, if there is someone standing by seriously considering putting the tray back into the freezer, or someone who

R. Here's the main problem: a belief can track the fact that p via some reason R without tracking p 's logical consequences. So we can know a claim is true, believe things that follow, yet fail to know that those things are true.¹⁰ Dretske's zebra case illustrates the point.¹¹

I assume that incompatibility with the principle of closure counts heavily against an analysis of knowledge, while compatibility counts in its favor. Hence before we decide on the relative merits of (2) versus (T2), it is important to see that (2) *supports* closure. Let us say that when belief p meets (2), R is an *infallible indicator* that p is true (or, more fully, we might say that, *relative to S's situation*, R is an infallible indicator that p is true). If R is an infallible indicator that p is true, as (2) requires, then R is also an infallible indicator that p 's logically consequences are true. If

R holds $\rightarrow p$

p entails q

hold, then it *must* also be the case that

just might do so, then I do not know the ice has melted. However, those who disagree can weaken (2), even to the point where it requires only that in S's situation as it was when R held, p is true.

¹⁰Consider the following inference:

(a) $not-p \rightarrow not-(R \text{ holds})$

(b) p entails q

(c) So: $not-q \rightarrow not-(R \text{ holds})$.

It is not valid, because the following is not a valid pattern of inference:

1. $not-p \rightarrow not-q$

2. p entails r

3. So: $not-r \rightarrow not-q$

Nor is the following, which commits the fallacy of contraposition twice:

1. $not-p \rightarrow not-q$

2. $q \rightarrow p$ (from 1—fallacy of contraposition)

3. p entails r

4. $q \rightarrow r$ (from 2,3 by strengthening the consequent)

5. $not-r \rightarrow not-q$ (from 4—fallacy of contraposition).

¹¹Actually, it is not obvious that we should reject closure if we accept the tracking account. As I point out in footnote 32, if we track p via some fact, say R, and believe q by deducing it from p , then we track q via p . We track q if we take p itself as our basis for believing q . We will deny closure only if we insist that knowing q entails tracking q via R, the fact via which we track p . But why say that? For other doubts about Nozick's case against closure, see various essays in Luper (1987a).

R holds $\rightarrow q$

is true.¹² So if my belief p is based on a reason R that is an infallible indicator that p is true, and I come to believe q by deducing q from p , then my belief q is also based on a reason that is an infallible indicator that q is true--namely R, or R plus the fact that p entails q (we could even say that the reason q is based on is p itself). For example, my zebra experiences are situationally infallible indicators that the animal in the cage is a zebra: in the close worlds in which I have these experiences, I'm seeing a zebra. But all the worlds in which I'm seeing a zebra are worlds where what I'm seeing is not a disguised mule. The upshot is that while tracking is hostile to closure, infallible indication is not.¹³

Even when strengthened in various ways, the indicator analysis will sustain closure. This is fortunate, since our account does need to be refined and clarified in various ways. We cannot go into all of these, but one adjustment seems especially important, since conditions (1)-(2) can be met in cases in which our belief's source is (generally) unreliable. To make this point clear, let's tinker with one of Goldman's examples. We'll play a prank on Sue, a wacko who takes highly realistic barn appearances to indicate the presence of *papier-mâché copies* of barns. The prank is innocent enough: we bundle her off to a region of the world in which all the real barns

¹²The inference pattern here is strengthening the consequent, which is valid:

1. $p \rightarrow q$
2. q entails r
3. So: $p \rightarrow r$

¹³Our thesis applies directly to Nozick's tracking account, too. Consider the following revision of his view (contraposing his third condition, and leaving out the first, which is redundant): S knows p if and only if there is a method M such that S believes p via M, and S believes p via M $\rightarrow p$. Suppose S knows p via M¹, and believes q by deducing q from p . Then there is a method M² via which S believes q such that

have been replaced with *papier-mâché* duplicates. When she spots a papier-mâché barn in the distance, her visual impressions lead her to think that she's looking at a papier-mâché barn. Here's the point: the fact that she has barn-type percepts is, then and there, an infallible indicator that she sees a papier-mâché barn, so she meets our two conditions for knowledge. Given her peculiar circumstances, if she were to see barn appearances, a papier-mâché barn would be present. Yet she does *not* have knowledge. She takes highly realistic barn appearances to indicate the presence of papier-mâché barns, but barn appearances are not generally reliable indicators of papier-mâché barns (instead, realistic barn appearances are reliable indicators of *barns*). Generally, in the situations in which we have barn appearances, we are not confronted with papier-mâché barns.

These considerations suggest the need to add the following to our list of conditions for knowledge (naturally, those who are not convinced by the example can get by without the addition):

(3) Generally, when an R-type situation holds, a *p*-type situation holds.

Let us say R is a *reliable indicator* that *p* is true if and only if this condition is met.

Requiring a reliable indicator does not stop us from accepting closure. If R is a reliable indicator that *p* is true, and *p* entails *q*, R is also a reliable indicator that *q* is true. (I'm assuming that to identify a *p*-type situation, we start with the situation described by *p* and abstract away the element of time, and that if *p* entails *q*, then when a *p*-type situation holds, a *q*-type situation must also hold.) Now,

- Generally, when an R-type situation holds, a *p*-type situation holds

- *p* entails *q*

S believes *q* via $M^2 \rightarrow q$,
 --namely, *arriving at belief p via M^1 and believing q by deducing q from p.*

together entail

- Generally, when an R-type situation holds, a q -type situation holds.

So if S's belief p is based on a reason R that reliably indicates that p is true, and S believes q by deducing q from p , then S's belief q will be based on a reason that reliably indicates that q is true. By contrast, the tracking counterpart of (3), namely

(T3) Generally, when a p -type situation does not hold, an R-type situation does not hold,

will not support closure.¹⁴

The fact that conditions (1)-(3) endorse closure, while their tracking counterparts do not, is *prima facie* grounds for preferring the former. Nonetheless, Dretske and Nozick would probably argue that the tracking condition (T2) is preferable to the indication condition (2) precisely because the former positions us to acknowledge our ignorance of skeptical hypotheses, which gives us a way to account for skepticism. Skepticism seems correct because skeptical conclusions can be reached using the closure principle. As for where skepticism goes wrong, we can blame the closure principle itself: skeptics assume it, but it is false. Is this response plausible?

Tracking and Skepticism

It is plausible. Some versions of skepticism do rely on closure. An example is the following *skeptical argument from closure*:

(i) If S knows p , and believes q by deducing q from p , then S knows q .

¹⁴Apply it to Dretske's zebra case. z meets (T3): when we are not confronted with zebras, generally we do not get zebra-type experiences that might lead us to suspect the presence of zebras. Furthermore, z entails *not-m*. But *not-m* does not meet (T3): on those (rare) occasions when we are confronted with mules disguised as zebras, generally we get zebra-type experiences, indicating that *not-m* is true. (But recall the sort of reservation expressed in notes 11 and 32.)

- (ii) But S does not know that S is not a brain in a vat on a planet far from earth whose sensory experiences are completely misleading. (S does not know this because S does not track the fact that S is not a brain in a vat.)
- (iii) So S does not know anything that entails that S is not in the vat scenario (including most of S's ordinary beliefs).

Of course, the skeptical argument from closure constitutes a form of indiscernability (Cartesian) skepticism, and has nothing to do with regress (Pyrrhonian) skepticism. By denying closure, tracking theorists leave regress skepticism untouched, and address only one form of indiscernability skepticism. A response to one skeptical argument is not necessarily a response to others. But progress is made if one prominent form of skepticism is defeated. And I think that this is a substantial virtue of tracking accounts.

As powerful as the above response is, we can reject it if we provide an account of knowledge that sustains closure yet explains the appeal of skepticism while pinning down the flaw in the argument from closure. And in fact we have such an account: knowledge as indication. We already know that (1)-(3) support closure. And, as I will explain later, we can locate the flaw in the skeptic's argument at step (ii): by the conditions of the indicator analysis, we know we are *not* brains in vats. We can even explain the appeal of skepticism--in much the same way as the tracking theorists do. Let me elaborate.

Why One Form of Skepticism is Appealing

I begin with a concession. Dretske's and Nozick's story is basically the right explanation of the appeal of (one form of) skepticism. This is true even though the tracking account is wrong and the indication account is right. Skepticism is appealing because people often *think* of knowledge as tracking. They think of knowledge as tracking because

knowledge closely *resembles* tracking.¹⁵ It is easy to conflate (2) with its tracking counterpart (and (3) with *its* tracking counterpart), and if we do, we are likely to believe—falsely—that to know we are not brains in vats we must track the fact that we are not brains in vats. We will then disavow knowing we are not brains in vats. So even if (1)-(3) are the correct conditions for knowledge, and we do know we are not brains in vats, at least part of the explanation of why skepticism is tempting is surely that we do not track the fact that we are not in skeptical scenarios. We also see our way past the temptation: we can see through the skeptical argument from closure by familiarizing ourselves with the differences between tracking and indication, and carefully applying (1)-(3) rather than the tracking counterpart.

At this point Dretske and Nozick might well object. Perhaps (2) closely resembles (T2), and perhaps people run the two conditions together, but aren't these points grounds for concluding that the ordinary notion of knowledge can be captured only if we adopt *both* sets of conditions? If so, we are back to an analysis that undermines closure, and we seem forced to accept the tracking theorists' account of skepticism in its entirety.

This objection has substantial merit, but it can be overcome. True, in conflating indication with tracking, people might end up insisting on both. But it is more likely that people who conflate the two shift back and forth from one to the other, using the first on some occasions and the second on other occasions, oblivious, all the while, to the equivocation. There is plenty of evidence that people equivocate in this way. For example, consider the fact that we vehemently cling to the principle of closure. It is when we interpret knowledge as indication that we adopt this principle. We would not

¹⁵Luper 1984; Sosa 1999.

adopt it if we accepted all of the conditions or just the tracking conditions by themselves. Consider, too, how natural it is for people to claim to know a possibility does not hold on the grounds that the possibility is *remote*, and sometimes they stick to their guns even when we point out that they do not meet the tracking requirement. On the other hand, sometimes they will not stick to their guns: perhaps then they are worried about tracking.

I admit that it is difficult to pin down the ordinary concept of knowledge, and show that I have it right while Dretske and Nozick have it wrong. I do not have a knockdown case against the tracking theorists. We should not expect too much from ‘intuitions’ and appeals to the ordinary concept of knowledge; anyone familiar with post-Gettier attempts to analyze knowledge realizes that ‘intuitions’ and language analysis do not always provide the means to choose among competing accounts. When this happens, however, we can defend a view on the grounds that it *refines* the ordinary conception of knowledge.¹⁶ And in that spirit I recommend the indicator analysis. Indication *improves upon* the commonsense concept of knowledge, for, unlike the ordinary notion, it unambiguously sustains the principle of closure.

Where Our Skeptical Argument goes Wrong

Still another objection is available. On the view I have offered, we know we are not brains in vats. Surely, tracking theorists might insist, we can know no such thing—about *that*, at least, skeptics are correct. And anyone who disagrees needs to explain how such knowledge is possible. This challenge seems fair enough, and I think it can be met. At a minimum, I can make my position as plausible as Dretske’s and Nozick’s. That is, the combination,

¹⁶Compare Quine’s (and Carnap’s) notion of explication in “Two Dogmas of Empiricism,” p. 25 of Quine 1953.

(a) I know I am not a brain in a vat, and

(b) I know ordinary claims that are incompatible with being a brain in a vat,

my view, is at least as plausible as the combination,

(a) is false, yet

(b) is true,

which is their view.

Consider what, specifically, is involved in knowing that a skeptical possibility does not hold. I know that

h: I am sitting in my house typing out an essay

is true. My reason R^1 for believing *h* is my having certain familiar sensory experiences, relating to my house and computer, and these are infallible (and reliable) indicators that *h* is true. Furthermore, *h* entails

not-biv: I am not a brain in a vat on a far-off planet whose experiences are completely misleading.

Seeing this, I believe *not-biv*. And I have an infallible reason for believing *not-biv*, namely, R^1 --my having the experiences on which I based belief *h*. Of course, if asked why I believe *not-biv*, I will likely answer that my belief is based on a compound reason: the fact that *h* is suggested by the experiences involved in R^1 together with the fact that *not-biv* follows from my belief *h*. (Alternatively, I might say that my reason for believing *not-biv* is *h* itself.) No matter: this complex reason (or *h* by itself) is itself an infallible (and reliable) indicator that I am not a brain in a vat on a far-off planet.

What might make us suspicious about knowing *not-biv* is the fact that normally our grounds for it (like our grounds for many other claims, as G. E. Moore (1959) pointed

out) are indirect.¹⁷ Just now, when I explained the typical way in which we know that *not-biv*, I mentioned one source (R¹—involving experiences of sitting at the keyboard and so on) for the commonsense belief *h*, and a compound source for *not-biv*. We base commonsense beliefs directly on experience; why can't we base beliefs about skeptical possibilities directly on experience? Why must the latter be based on the commonsense beliefs instead? Unless we can answer these questions, we will suspect that beliefs about skeptical possibilities are special—and perhaps that skeptics are right about them.

Fortunately, there are answers.

As to why our grounds against skeptical possibilities typically are indirect, here is part of the story: suppose I am asked why I believe I am in my house typing right now. In crafting my answer, I must deal with the fact that I did not go through any explicit reasoning, or apply any explicit test. I rarely do when perception leads me to believe things. Nonetheless, I find that certain sorts of familiar experiences habitually trigger certain sorts of commonsense beliefs. So I can give an answer by laying out the sorts of experiences that typically trigger beliefs about typing, buildings, and so on. But the belief that I am not a brain in a vat is not one I ordinarily entertain, hence it is not possible to link it with experiences that habitually trigger it. It rarely comes up. But when it does, I find myself believing it because I see that it follows from my commonsense beliefs.

Here is another part of the story: suppose I clarify the ways I arrive at perceptual beliefs by explicitly crafting rules of the form, if I have experiences *E*, then *p* is true. For example: If I get zebra-in-a-cage-type experiences, then the statement, 'There is a zebra in the cage' is true. Such rules explicitly address a particular statement. But they need

¹⁷Cf. Richard Feldman (1999, p. 105).

not explicitly address claims that *follow* from this statement. Clearly, my rules can say that a claim p is true without (explicitly) saying that p 's consequences are true.¹⁸ In a fairly straightforward sense, a thermometer (made out on a Fahrenheit scale) that registers 70 degrees *says* it is 70 degrees Fahrenheit, but *not* that the temperature is not 0 degrees Celsius. Nonetheless, we can know the latter indirectly, since the one fact entails the other. The upshot is that my experiences can be infallible (and reliable) indicators that I am at home, and that I am not a brain in a vat on a distant planet, yet I might fail to have belief management rules which *prompt* me to believe I am not in a vat on the basis of my experiences.

Finally, let me address the claim that we *must* base beliefs about skeptical possibilities indirectly on commonsense beliefs, and not directly on experience. *This claim is false.* There is nothing to stop us from converting what is usually indirect knowledge into direct knowledge. I can simply take, as my reason for believing that it is not 0 Celsius, the fact that a Fahrenheit thermometer says that it is above 32. And I can know I am not a brain in a vat directly on the basis of the experiences that led me to believe that I am at home typing, rather than by inferring I am not in a vat from this belief.

One other potential criticism should be addressed. Lately, some philosophers apply the label 'Moorean,' disparagingly, to any attempt to say that we can know skeptical possibilities do not hold. The thought is that Moore merely turned the tables on

¹⁸This does not threaten closure. Suppose our source R for believing p does not address one of p 's consequences q . Suppose also that we know that p is true and we believe that q is true by deducing it from p . Then our source for believing q is: seeing that q follows from p which is itself indicated by R, and this compound source is a reliable and infallible indicator that q is true.

the skeptic, without describing where the skeptic goes wrong. Such critics might try to dismiss my own efforts in the same way. I have two reactions.

First. Moore's demonstration that the tables can be turned on the skeptic is itself important. All of us who are not clinging to skepticism find Moore's antiskeptical argument¹⁹ at *least* as powerful as the skeptic's argument from closure. Second. Unlike Moore, I have explained what is *involved* in knowing that skeptical possibilities do not hold. I have not just *asserted* that the skeptic's argument fails. Skepticism can be resisted even if we grant the closure principle, however qualified. The key is the Mooresque claim that we know skeptical scenarios do not hold.

II. CONTEXTUALISM V. THE INDICATOR ACCOUNT

Let's turn to the contextualists' explanation of skepticism, and begin by confronting a possible source of confusion: those of us who defend subjunctive accounts, whether we view knowledge as tracking or as indication or the like, have always thought that knowledge depends on the knower's context. Yet we are not contextualists—not, in any case, as that term is usually meant these days. In what ways do the two approaches overlap, and how do they differ?

Subjunctive Accounts and Contextual Variables

For the subjunctivist, it is a straightforward matter that our status as knowers depends on features of our situation. In some circumstances a belief source R must be especially versatile to generate knowledge. To be a reliable indicator, R must display a minimal

¹⁹His antiskeptical argument is this:

- (i) I know all sorts of commonsense claims: I am in my house typing an essay, and so on.
- (ii) If I know these things that are incompatible with being a brain in a vat, then I know I am not a brain in a vat.

level of reliability, but to be a situationally infallible indicator, R often must do still better, depending on the specific nature of the circumstances at hand. In optimal conditions I can know I see a barn by looking at it. If there are papier-mâché barns scattered through the area, however, I can look right at a barn, recognize it as a barn, and still not know that it is a barn. I will need to become (or employ) a better barn detector. I will need to walk inside the barn, or touch it, or the like. I will need a belief source that, even in the presence of fake barns, would indicate the presence of a barn only if there were one there, and the more undetectable the fakes in my vicinity, the abler my source must be at detecting the fakes.

The point holds for noninferential as well as inferential knowledge. If our belief is noninferential—if, for example, it is perceptual—then the more deceptive our circumstances, the more acute our perceptual process must be.²⁰ If inferential, then the more deceptive our circumstances, the better our evidence must be. It is easy to imagine circumstances in which I can know the butler did it on the basis of evidence involving an honest eyewitness. But this evidence will not suffice if it turns out that the butler has a twin brother with a grudge against the victim. The required discriminatory powers of our knowledge-producing belief sources varies on a *sliding scale*, starting with a minimal level in ordinary circumstances, and rising with the level of deceptiveness of our situation.²¹

The element of contextualism evident in subjunctivist accounts of knowledge allows us to explain and put aside a skeptical argument that is quite different from the skeptical argument from closure. The argument—from the situation principle—is this:

(iii) So I know I am not a brain in a vat.

²⁰Assuming for the sake of argument that perception is a noninferential process.

- (i) If, in some situation, E is the evidence we have about p , and we know p on the basis of E, then we know p in any situation in which E is our evidence and we believe p on the basis of E. (Call this the *situation principle*.²²)
- (ii) For virtually any evidence E, there are circumstances in which we will not know p on the basis of E. (There are circumstances in which E is the evidence we have about p , and we believe p on the basis of E, yet fail to know p . That is, nearly all of our belief sources fail to produce knowledge in some circumstances.)
- (iii) So our evidence rarely positions us to know anything.

Reflection about Gettier cases reveals what is wrong with the situation principle.²³ When I look at a caged zebra at an ordinary wildlife show, I know the animal is a zebra. Yet in the Gettierized version of this situation, where I stumble into the Hey Presto wildlife show, featuring zebra doppelgangers, I fail to know, even though my belief source is the same in both cases. But there is no mystery here. What is going on is that a belief source that is discriminatory enough in ordinary circumstances is not discriminatory enough in others. The situation principle has to go.

Agent-Centered Contextualism v. Speaker-Centered Contextualism

According to the subjunctivist as well as the contextualist approach, it is harder to know things in some circumstances than in others--our epistemic apparatus must be more

²¹Luper 1986.

²²A weaker principle is also false: if, in some situation, the evidence we have about p is the *fact* that E holds, and we know p on the basis of E, then we know p in any situation in which E is our evidence and we base belief p on the fact that E holds. Goldman's papier-mâché barn case tells us how to construct a counterexample.

²³Luper 1987b. For more on aligning Gettier cases with skeptical scenarios, see Luper 1984.

versatile in some circumstances in order to produce knowledge.²⁴ Both accept a view we might call *agent-centered* contextualism, which claims that whether it is correct to attribute knowledge to a given agent S depends on features of S's context. But subjunctivists and contextualists diverge when it comes to a further issue. The latter usually defend a position we might call *speaker-centered* contextualism,²⁵ which says that whether it is correct for a speaker to attribute knowledge to a given agent depends on features of the *speaker's* context. Whether I, the speaker, am right when I say that Mary, the agent I am discussing, knows *p*, depends on my context, as opposed to Mary's (it could also depend on Mary's context, too—this is not ruled out by definition). The truth conditions for knowledge attributions vary with speakers' contexts. Thus when I say Mary knows *p*, and you, referring to the same person, say Mary does not know *p*, we can both be correct—we are not contradicting each other. Subjunctivists acknowledge that all sorts of features of a speaker's context might affect whether she will *say* (or be *warranted* in saying) that an individual knows something. But subjunctivists typically deny that the *truth conditions* of knowledge attributions vary with speakers' contexts. Hereinafter, I'll use the term 'contextualism' to refer to speaker-centered contextualism.

The Contextualist Explanation of Skepticism

So how do contextualists deal with skepticism? Their approach has two parts: a diagnosis (accounting for the appeal of skepticism), and a cure.

Here's the diagnosis. When skeptics discuss their hypotheses with us, they put us into a special context; there, we may correctly attribute knowledge to an agent S only if S meets very demanding epistemic standards. These are so rigorous that it is almost always

²⁴In this sense subjunctivists have always accepted a claim contextualists have been emphasizing lately: the standards knowers must meet vary with context.

a mistake to attribute knowledge to anyone (whether ourselves or others). So there is an element of truth in skepticism, according to contextualists: when we discuss skeptical possibilities, we enter a context where heightened standards apply, forcing us to conclude that no one knows that skeptical possibilities do not hold. Because skepticism prevails in this special context, skeptics think their view must be accepted: Surely if we have to deny anyone knowledge when pressed by skeptics, then we have to deny anyone knowledge period—we have to deny it when we are not thinking about skepticism. Still, there is a suppressed assumption here: namely, that proper attributions of knowledge do not depend on the context of those who judge. The assumption is analogous to the situation principle:

The context principle: if it is correct for judges in one context to attribute knowledge to an agent S, then it is correct for judges in any context to attribute knowledge to S. (If it is improper for judges in one context to attribute knowledge to S, then it is improper for judges in any context.)²⁶

Fully spelled out, then, the skeptic's argument (from the context principle) is this:

- (a) The context principle is true.
- (b) People who raise skeptical doubts create a context in which they must say that no one knows very much.
- (c) So everyone, regardless of her context, must say that no one knows very much.

²⁵For example, Unger 1984 and 1986, Cohen 1988, 1991 and 1999, DeRose 1995, and Lewis 1996.

²⁶If challenged about this principle, a skeptic might respond that those who deny it are probably confusing it with a principle other skeptics adopt: the situation principle, used in the skeptical argument from the situation principle. One might defend the argument from the context principle without defending the argument from the situation principle. Alternatively, skeptics might claim that their critics are confusing the context principle with the following warranted assertability principle, which is also false: if a judge in

Now the contextualists' cure: the skeptic's argument is no good because the context principle on which it relies is false. What skeptics fail to see (what they deny?) is that the epistemic standards that are appropriate for knowledge attributions *vary* from context to context. In ordinary contexts, we may attribute knowledge to a given agent S when S meets quite low epistemic standards. So in ordinary contexts it is correct to say that people know ordinary knowledge claims. Skeptics think that by getting us to admit ignorance in the special context they create when they raise their doubts, they can show that everyone is ignorant regardless of context. But people who ignore skepticism and skeptical possibilities are in a context where lower standards are in place, and it is correct for them to attribute knowledge to themselves and to others who ignore skepticism.

The upshot is that contextualists can say that skeptics are more right than they are wrong: anyone who discusses skepticism must end up embracing it, while those who ignore skepticism can escape it. An added bonus is that contextualists do not have to reject closure, so long as its application is confined to one context at a time. In all contexts, the principle of closure is correct, but if we shift from one context to another in the course of applying closure, "all bets are off," as Lewis (1996) says. This is the error involved in the skeptical argument from closure; it is also involved in the Moore-style inversion of the skeptic's argument.

To make the contextualist approach more concrete, let us outline a version offered by Keith DeRose (1995). Suppose we ask whether some agent S knows p . And suppose that sk is the most remote alternative to p that we are considering. Let us say that p 's *DeRose zone* is the sphere of possible worlds centered on the actual world that includes

one context is *warranted* in attributing knowledge to an agent S, then a judge in any context is warranted in attributing knowledge to S.

the closest worlds in which p is false. For us to correctly say S knows p , two requirements must be met. First, S's belief as to whether p is true must match the fact of the matter throughout p 's DeRose zone: that is, S must believe p throughout the p portion of p 's DeRose zone, and $not-p$ throughout the $not-p$ portion of p 's DeRose zone (DeRose 1995, reprinted in DeRose 1999, p. 206). Second, S's belief as to whether p is true must match the fact of the matter throughout $not-sk$'s DeRose zone.

Using DeRose's account, let us see how the contextualist will deal with skepticism. DeRose's second condition will be hard to meet if sk is a very remote alternative to p : in a context where skeptical possibilities are raised, knowledge is difficult to attain. For example, we will not qualify as knowing there is a zebra before us in a context where the possibility of cleverly disguised mules arises, for there are $not-z$ worlds in which we believe z within the $not-m$ DeRose zone. So if the context principle were true, we would be judged to know little in any context. This is the diagnosis of skepticism. Now the cure: both of DeRose's conditions easy can be met when skeptical possibilities do not arise. In an ordinary context, we will count people as knowing there is a zebra in front of them. So the context principle is false; people can be credited with knowledge in ordinary contexts.

There you have the contextualists' story. I will argue that it should be rejected, since we have already offered a better diagnosis of skepticism as well as a better cure.

Problems with the Contextualists' Approach

This essay began with three requirements an account of knowledge ought to meet:

- (a) It should square with our intuitions about clear examples of knowledge and ignorance.

(b) It should help us to diagnose and cure skepticism.

(c) It should not be in tension with plausible epistemic principles.

Like subjunctivists, contextualists can defend their view on the basis of these requirements. Their approach to skepticism is superior (they might say) because tracking theorists reject the principle of closure, thus violating (c), while contextualists need not, and otherwise the two approaches are fairly evenly matched: both explain the appeal of skepticism by saying that people do not know that skeptical possibilities fail to hold, and both say we know about ordinary possibilities, such as that we have arms and legs, so both look good from the standpoint of (a) and (b).

As for the indicator theorist's diagnosis of and cure for skepticism, contextualists would reject it, too, even though it supports closure, as suggested by (c). They would argue that on our account we know that skeptical possibilities do not hold, and this stops us from explaining the appeal of skepticism, in violation of (b).

However, this argument is unconvincing. We need not accept skepticism to account for its appeal. In one form—the argument from the situation principle—skepticism's appeal is due to the (limited) charms of the situation principle. This form of skepticism will be especially irresistible if we confuse the situation principle (which is false) with the context principle (which is true even while it is denied by contextualists). In another, skepticism derives from the persistent impression that we must jump-start our justificatory efforts in a way that seems ruled out by the Pyrrhonian regress argument. And the skeptical argument from closure has its own source of appeal. As we have already explained, its attractiveness derives from an *error* that is very hard to detect and

avoid--an error that results when we conflate tracking conditions and their non-tracking counterparts, and conclude that we cannot know that skeptical scenarios don't hold.²⁷

Our explanation can be adapted to the metalinguistic level, too, so there is no need whatever to accept attribution conditions that support skepticism. The view we need is this: it is correct for a speaker to attribute knowledge that p to agent S if and only if S meets the indicator conditions for knowing p . On this view, it is easy to explain people's hesitation to attribute knowledge when a skeptical scenario sk arises, either for the speaker or the agent, since they confuse indication with tracking and notice that the tracking conditions are not met. They then reason: S would not know *not-sk*, so (by closure) S does not know ordinary things. Later, when the skeptical possibilities are ignored, people switch back to the indicator conditions, attribute knowledge accordingly, and refresh their confidence in the closure principle.

This reasoning shows that our way of dealing with skepticism is at least as plausible as the contextualists'. Now let's argue that it is superior.

The main concern is that the contextualists' approach forces us to reject very plausible epistemic principles, thus violating (c). All things being equal, we are entitled to resist a theory that forces us to abandon such principles, as contextualists themselves say in defense of their view against Nozick and Dretske. In fairness, we must note that contextualists (or at least David Lewis) have seen this sort of objection coming, but, as some of Nozick's critics said, a theorist cannot preclude criticism by being first to point out the troublesome consequences of his own view. Applying *modus ponens* does not

²⁷As we have said, the skeptical argument from closure is by no means the only defense of skepticism. Different errors are involved in some of the other defenses. For example, the argument from the situation principle is flawed because the situation principle is false.

stop others from applying *modus tollens*. Let me mention a couple of principles contextualists must reject.

The first we might call the *principle of stability*:

If at time 1 S is correct in saying, 'I know *p* at time 1,' and at time 2 S's situation remains the same (as far as possible) except that S thinks of one or more implications of *p*, then, at time 2, S will (still) be correct in saying, 'I knew *p* at time 1.'

To see why contextualists reject this principle, let's use Dretske's example again. You are at an ordinary wildlife show, looking at a caged zebra, believing that a zebra is there. Speaking at time 1, you are correct when you say, 'I know *z*', for you meet the low epistemic standards in place in this context. A moment later, at time 2, you continue to believe *z* on the same basis, but you start thinking through the implications of *z*; realizing that *z* entails *not-m*, you come to believe *not-m* thereby. Speaking at time 2, you have to say that you do not know *not-m*, and further that you do not know *z*, for you are now considering a skeptical possibility, which triggers high standards that are not met by either belief. Worse: judging at time 2, you have to say that you did not know *z* at time 1! For you-at-time-1 cannot meet the standards which you-at-time-2 are relying on to issue correct attributions of knowledge. Thinking through what is implied by your knowledge destroys your knowledge. More carefully, it destroys your knowledge as assessed while considering skeptical possibilities: judging at time 2, you will admit that *what you said* at time 1 when you uttered the words 'I know *z*' was *true* (since what you *said* was, in effect, that your belief *z* met the lower epistemic standards in place at time 1)

but you will be forced to add that you did not know z at time 1 (since you did not meet the higher standards *now* in place)!²⁸

Another plausible principle contextualists must reject is a metalinguistic version of the closure principle:

If S correctly attributes to herself knowledge that p , then if S *had* believed q by deducing q from p , S *would have been* correct to attribute to herself knowledge that q .²⁹

This principle is extremely plausible, since p *must* be true if S correctly attributes to herself knowledge that p , so that if S were to believe q by deducing q from p , q must be true, too. Yet the contextualist will have to deny this principle. Why? –Back to Dretske’s zebra case. I know z (the possibility that m not having occurred to me). But the close worlds in which I deduce *not- m* from z are worlds in which heightened standards apply; there, I cannot be said to know *not- m* . (I cannot be said to know z in such worlds either, but this does not stop me from knowing z in the actual world.) So the antecedent of the metalinguistic closure principle is true, while the consequent is false.

Does the contextualist give plausible grounds for denying the metalinguistic closure principle? I don’t think so. According to the contextualist, *S-in-the-actual-world* does not meet the higher epistemic standards applied to *S-in-the-hypothetical-world-where-S-believes not- m* . So what? The fact that actual- S meets the *lower* epistemic

²⁸As Curtis Brown pointed out to me, a bizarre consequence of contextualism is that the following reasoning is invalid:

At time 1 S said ‘I know p .’
What S said was true.
So S knew p at time 1.

²⁹Here is another metalinguistic closure principle contextualists must deny:

If at time 1 S is correct in saying, ‘I know that p ,’ and at time 2 S ’s situation remains the same (as far as possible) except that S comes to believe q by deducing q from p , then, at time 2, S will be correct in saying, ‘I know that q .’

standards *entails p*! Even by the lowest plausible standards, knowing *p* entails *p*. So if, as the antecedent of the metalinguistic closure principle says, S's knowing *z* is certifiable (by weak standards), *z* must be true, and if S believed *not-m* by deducing it from *z*, S could not be wrong about *not-m*, even if neither the belief *z* nor the belief *not-m* meets the contextualist's higher epistemic standards. Truth is preserved from the one context to the other.

Once we notice that contextualists reject the metalinguistic closure principle, their solution to skepticism resembles that of the tracking theorists quite a bit. For the skeptic could draw on the metalinguistic closure principle to offer an argument against all commonsense knowledge attributions. As follows:

1. No matter what S's circumstances are, it is *never* proper for S to say she knows that she is not a brain in a vat (*not-biv*).
2. So if S were to believe *not-biv* by deducing it from a claim *h* that entails *not-biv*, she still could not properly say that she knows *not-biv*, no matter how well grounded her belief *h* happens to be.
3. The metalinguistic closure principle is true.
4. So S will always be incorrect in attributing to herself knowledge of anything that entails *not-biv*.

The contextualists' response will be that the principle at step 3 is false. But given this response, do we really want to say their solution to skepticism is better than the solution offered by the tracking theorists? Tracking theorists confront a skeptic who argues at the object level; they reject the object level version of the closure principle. Contextualists simply move up a level, and confront a skeptic who appeals to the metalinguistic closure

principle. This is progress? Perhaps the metalinguistic closure principle is not quite as plausible as the principle Dretske and Nozick deny, but isn't it much too close for comfort? Both are plausible for similar reasons. Can't we at least say that we should avoid denying either if we can? And we *can*—by using the indicator account as a theory of both the truth conditions and attribution conditions for knowledge.

A final point: suppose the indicator account of attribution conditions did *not* square with our intuitions (if we have any!) about attribution conditions quite as well as the contextualist account. Suppose that the contest were very close, but lost by the former. There is a further important consideration against contextualist theories. In the spirit of revision, it would be much better to adopt the indicator account, for thereby we avoid the puzzling complications introduced when we say that truth conditions for knowledge claims vary with the speaker's context. In particular, our knowledge will not collapse in epistemic contexts. For the contextualist, knowledge can be correctly attributed only to people who have overlooked possibilities, and once those possibilities are looked at squarely, the knowledge attribution must be retracted. Not so on our account. We can boldly claim to know that skeptical possibilities do not hold, while still seeing the appeal of many sorts of skepticism. I endorse a remark Lewis³⁰ makes:

It is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary. One of those flawed premises is the claim that I do not know I am not a brain in a vat.

III. KNOWLEDGE AND RATIONAL BELIEF

³⁰Lewis 1996, p. 220.

So far we have ignored the issue of how knowledge is related to justified, or rational, belief. There is not enough space for a full discussion of this matter, but a few comments are in order.

Knowledge and Nonjustified Belief

First, we ought not to rule out noninferential (hence non-justified yet rational, as opposed to unjustified and irrational) knowledge; nor does the indicator account do so. If perception produces noninferential beliefs, it sometimes generates noninferential knowledge about such things as tables and chairs. Here's why: In optimal circumstances, and circumstances that are close to optimal, perception causes me to believe there is a chair in front of me (or what have you) only if there *is* a chair.

Moreover, generally, perception is accurate when it leads me to believe in the presence of chairs (and such).

Knowledge and Rational Belief

Second, we will want to specify that an inferentially known belief p be rational. From an internalist perspective, this entails that p be sufficiently justified, so that the preponderance of one's evidence counts in favor of believing p . Gettieresque concerns will be handled by (1)-(3), which are externalist conditions. In other words, we can say that S knows an inferential belief p if and only if S believes p on the basis of some fact R that provides sufficient justification for p , and R is both a reliable and infallible indicator that p .

So understood, inferential knowledge is closed under entailment, since we are justified in believing things that are entailed by individual things we justifiably believe.

That is, the following principle of closure of justification is true:

If S has sufficient justification for S's belief p , and S believes q because it follows individually from p , then S has sufficient justification for S's belief q .

Our principle is restricted to things that are entailed by *individual* things we justifiably believe in order to avoid the lottery paradox, which is generated if we say that we are justified in believing the *conjunctions* of things we are justified in believing.³¹

Many objections to the principle of closure of justification can be met once we notice two points. First, as Peter Klein (1981) pointed out, our evidence for p need not be the evidence for p 's consequences. The evidence for the latter might be p itself, which suffices since p entails its consequences.³² The second point is that evidence against any of the consequences of p counts against p . So if we are justified in believing p , we will not have strong evidence against any of p 's consequences, and we will have powerful grounds *for* them—namely p .

Needless to say, many theorists will not be satisfied with the above analysis of knowledge, combining, as it does, internalist with externalist elements. For them, I can only sketch a few suggestions. The most promising way to clarify the idea of a rational belief involves developing the notion of belief management practices, which are the predictable patterns by which each of us adopt, maintain and revise our beliefs. These belief management practices can be described, at least roughly, in the form of rules. An example of such a rule might be: accept spontaneously occurring experiential beliefs, unless they contradict other, firmly accepted beliefs. Such rules need not be consciously

³¹A particular lottery ticket has little probability of winning given that the game has many entries. So you are justified in believing that ticket 1 will lose. Similarly, you are justified in believing that ticket 2 will lose, and that 3 will lose, and so on. But you are not justified in believing the *conjunction* of these claims, which is tantamount to the claim that every ticket will lose.

³²Klein's point can also be pressed against the tracking theorists' denial of the closure of knowledge: if we track p , and believe q by deducing it from p , then we track q if we take p as our basis for believing q .

applied, and it might be difficult indeed for anyone to fully articulate rules that accurately capture their practices. Furthermore, some rules may not function to steer us toward an accurate picture of the world: wishful thinking, for example, preserves assumptions that make for happiness. We must set these aside in assessing the *epistemic* status of our views, and focus on rules used to get us to an accurate picture. (Hereinafter, the term ‘belief management rules’ will refer to the latter.) With these points in mind, and drawing on the work of Alvin Goldman (1979, 1988), we can offer a characterization of rationality. Roughly speaking, a person S’s belief *p* is *objectively* rational if and only if:

- (a) S believes *p* in conformity with S’s belief management rules
- (b) Generally, when S holds a *p*-type belief in conformity with S’s rules, it is true.³³

The reason: again roughly speaking, a belief is objectively rational if and only if it is held in conformity with belief management practices that are reliable when they endorse *p*-type beliefs. The contrast is with beliefs that are *subjectively* rational—i.e., held in conformity with management practices that one accepts for the purpose of arriving at an accurate picture of the world. (That my belief *p* is subjectively rational is ensured by (a).) Thus, a belief is objectively rational when and only when believing it in conformity with one’s belief management practices is a *reliable indicator* that it is true.

Steven Luper

Trinity University

³³Certain sorts of objection can be forestalled by qualifying this condition (and the same goes for the tentative externalist account of inferential knowledge to follow):

Generally, when S holds a *p*-type belief in conformity with (that subset of) S’s rules, in the circumstances for which those rules were developed, *p* is true.

We might also want to refer to S’s community:

Generally, when a member of S’s community holds a *p*-type belief in conformity with (that subset of) S’s rules, in the circumstances for which those rules were developed, *p* is true.

REFERENCES

- David Annis (1978). "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly* 15.
- D. M. Armstrong (1973). *Belief Truth and Knowledge*. Cambridge: Cambridge University Press.
- Laurence Bonjour (1980). "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* V, 53-71.
- Stewart Cohen (1988). "How to be a Fallibilist," *Philosophical Perspectives* 2, 91-123.
- (1991). "Skepticism, Relevance, and Relativity," in McLaughlin 1991.
- (1999). "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13, 57-89.
- Keith DeRose (1995). "Solving the Skeptical Problem," *Philosophical Review* 104.1, 1-15.
- Fred Dretske (1970), "Epistemic Operators," *Journal of Philosophy* 67, 1022.
- (1971). "Conclusive Reasons," *Australasian Journal of Philosophy* 49, 1-22.
- Richard Feldman (1999). "Contextualism and Skepticism," *Philosophical Perspectives* 13, 91-114.
- Alvin Goldman (1976). "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73, 771-791.
- (1979). "What Is Justified Belief?" in George Pappas, editor, *Justification and Knowledge* (Dordrecht: Kluwer Academic Publishers) 1-23.
- (1988). "Strong and Weak Justification," in *Philosophical Perspectives* 2, 51-69.

- Nelson Goodman (1955). "New Riddle of Induction," in *Fact, Fiction and Forecast*.
Cambridge, MA: Harvard University Press.
- Gilbert Harman (1986). *Change in View*. Cambridge, MA: MIT Press.
- Peter Klein (1981). *Certainty: A Refutation of Scepticism*. Minneapolis: University of
Minnesota Press.
- David Lewis (1996). "Elusive Knowledge," *Australasian Journal of Philosophy* 74.4,
549-567.
- Steven Luper(-Foy) (1984). "Knowledge, Nozickian Tracking, and Scepticism,"
Australasian Journal of Philosophy 62, 26-50.
- (1987a). *The Possibility of Knowledge*. Totowa, NJ: Rowman and Littlefield.
- (1987b). "The Possibility of Skepticism," in Luper (1987a).
- (1987c). "The Causal Indicator Indicator Analysis of Knowledge," *Philosophy and
Phenomenological Research* 47, 563-587.
- B. McLaughlin (1991). *Dretske and his Critics*. Oxford: Basil Blackwell.
- G. E. Moore (1959). "A Defense of Common Sense," in *Philosophical Papers*. George
Allen & Unwin, pp. 32-59.
- Robert Nozick (1981). *Philosophical Explanations*. Cambridge: Harvard University
Press).
- W. V. O. Quine (1953). *From a Logical Point of View*. New York: Harper & Row,
publishers.
- John Rawls (1971). *A Theory of Justice*. Cambridge: Harvard University Press.
- Ernest Sosa (1999). "How to Defeat Opposition to Moore," *Philosophical Perspectives*
13, 141-152.

Gail Stine (1976). "Skepticism, Relevant Alternatives, and Deductive Closure,"

Philosophical Studies 29, 249-261.

Peter Unger (1986). "The Cone Model of Knowledge," *Philosophical Topics* 14, 125-

178.

----- (1984). *Philosophical Relativity* (Minneapolis: University of Minnesota Press).

Jonathan Vogel (1987). "Tracking, Closure, and Inductive Knowledge," in Luper(-Foy)

1987a.