

Knowledge Externalism

Marc Alspector-Kelly

Knowledge Externalism

Abstract

A popular counterexample directed against externalist epistemological views is that of an agent (Leher's 'truetemp' for example) whose beliefs are clearly neither justified nor known but that were generated in the manner that the externalist requires, thereby demonstrating externalism to be insufficient. In this essay I develop and defend an externalist account of knowledge—essentially an elaboration of Fred Dretske's information-theoretic account—that is not susceptible to those criticisms. I then briefly discuss the relationship between knowledge and justification.

Knowledge Externalism

A popular counterexample directed against externalist epistemological views is that of an agent (Leher's 'truetemp' for example) whose beliefs are clearly neither justified nor known but that were generated in the manner that the externalist requires, thereby demonstrating externalism to be insufficient. In this essay I develop and defend an externalist account of knowledge—essentially an elaboration of Fred Dretske's information-theoretic account¹—that avoids the counterexample. I then briefly discuss the relationship between knowledge and justification.

1. Defining Terms

Many proposals for explicating the internal/external distinction are on the books. Without arguing for it here in any detail, I will understand the internal characteristics of the epistemic agent to be those that would survive transportation to an evil-demon world, that is, a world in which one becomes the victim of a Cartesian evil demon. (Call such characteristics 'demon-invariant.')

This lines up nicely with 'access' internalism, according to which the internal is that to which the believer has access, since it is plausible that we enjoy such access only to whatever it is that the demon cannot hide from us. It also lines up nicely with 'perspectival' internalism, according to which the internal is whatever lies within the believer's first-person perspective, since it is plausible that the believer's first-person perspective is that which travels with him to the demon-world.²

If the internalism is with respect to justification, then the claim is that ‘justifiers’—that is, whatever it is that turns the merely believed into the justifiably believed—are demon-invariant. If the internalism is with respect to knowledge, then the claim is that whatever it is that turns true belief into knowledge (which may or may not be justification) is demon-invariant.³ Externalism is easier to describe: the externalist just denies that the status of a belief as justified or as known depends on what aspects of the epistemic agent are demon-invariant.⁴

3. The Problem with Bob

A standard counterexample against externalism begins by describing someone whose beliefs (or whose beliefs of a certain sort or about a certain subject-matter) do satisfy the externalist’s condition—they are reliably produced, or produced by information, etc.—but do so in an unusual and unexpected manner concerning which the believer herself is entirely unaware.⁵ Since she is clueless as to her fortuitous circumstances, she has absolutely no reason to believe that her beliefs satisfy the externalist’s condition, and so no reason to think that they are in fact true. Indeed the cases are often elaborated in such a way that she has considerable evidence against them.⁶ Her beliefs are therefore, the internalist claims, surely not justified. For you cannot be justified in a belief while lacking any evidence or reasons for it; at least, you surely cannot be so justified if you are in possession of overwhelming reasons or evidence against it. The internalist typically concludes further that her beliefs are also not

knowledge, notwithstanding satisfaction of the externalist's condition, because they are not justified.⁷

Keith Leher's 'Truetemp'⁸ and Laurence Bonjour's 'Norman'⁹ are perhaps the better-known such cases, but I am going to present a series of cases of my own in order to avoid certain irrelevancies.¹⁰

Bob has joined a religious cult. The cult employs those methods for the assurance of its followers' loyalty for which such cults are reviled: Bob is kept only barely above starvation, communication with his family is not permitted, he is not allowed access to television or newspaper, and so on. The indoctrination having been completed, Bob is now incapable of critical evaluation of his leader's veracity; he cannot help but believe the leader's every word.

The leader hands Bob a (working, accurate) thermometer, and tells him that whatever number appears on its face is in fact the current temperature. And Bob believes him; not because he knows that it is a thermometer, but only because the leader has so instructed him. Suppose, for good measure, that the indoctrination process has eradicated his memory of thermometers, so that he entertains no suspicion that device in his hand is one. Bob's beliefs concerning the temperature are reliable, information-produced, and so on. But he certainly does not know what the temperature is, nor are his beliefs on the matter justified.

4. What the Problem with Bob Is Not

At this point the internalist will suggest that satisfaction of her favorite internal constraints will change matters. Suppose that one day Bob, in response to a vague dissatisfaction that he cannot bring into sharper focus as a result of his weakened condition, sneaks into the leader's office and discovers therein a refrigerator whose contents he proceeds to devour. In a sharper and more skeptical frame of mind as a result, Bob becomes suspicious of his leader's veracity. In particular, he is no longer prepared to assume that the thermometer gives the temperature solely on the basis of having been told as much by his leader. But upon further inspection of the office Bob discovers a pamphlet from the thermometer's manufacturer that explains what it is and how it works. As a result he trusts the thermometer once again; but he does so now on the basis of evidence.

His temperature beliefs are now justified in the internalist's sense. And it seems reasonable to consider them known. After all, he actually read (and understood) the description of the thermometer's mechanism, something most of us, who surely can and often do know what the temperature is by reading a thermometer, have not done.

That Bob's beliefs concerning the temperature are now justified where they were not before is a result of those beliefs' being now based on evidence. But the internalist is wrong to think that their being so based is, in itself, responsible for his new knowledge.

Suppose that well before Bob entered his office the leader, having noted a certain wavering in Bob's convictions, anticipated his disloyal behavior. The thermometer's manufacturer distributes no such pamphlet as Bob discovered in the office. The device operates, in fact, according to entirely different physical principles than those described therein. The leader fabricated the pamphlet and planted it for Bob to discover, so as to restore Bob's faith and bring him back to the fold. Bob's temperature beliefs are still true,

reliably produced, and so on, as before; and he is now justified in the internalist's sense. But he still does not know what the temperature is.

5. What the Problem with Bob Is

Consider, first, Bob's circumstances before his discovery of the (putative) manufacturer's pamphlet. The problem with Bob, I suggest, is that while his temperature beliefs are shaped by an information-generating process—they are, as the terminology goes, the product of an information channel—the fact that it *is* an information channel has utterly no impact on Bob, whose conformity with the thermometer is solely the product of indoctrination and his leader's instruction. His beliefs are the product of a reliable source of information and he does rely on that source, but he does not rely on it *because* it is reliable.

The internalist will read the 'because' of the last sentence as concerning Bob's reasons or evidence: the problem with Bob, the internalist suggests, is that he does not rely on the thermometer in response to evidence or reasons to believe that it is reliable. This looks plausible because in the typical course of events one's beliefs' being responsive to evidence will ensure in addition that they are generated by an information channel. The fact that appeal to evidence typically has this effect, I suggest, is precisely why we consider it to be epistemically valuable.

But, as the elaboration of Bob's tale above indicates, what is typical is not inevitable. Had the putative manufacturer's pamphlet been genuine, it would have been informative precisely because the pamphlet's author would have been familiar with that

mechanism (or at least would have learned it from someone who is). Its contents would therefore contain the information that the mechanism is an information channel. Bob's conviction upon reading the pamphlet that the mechanism is such a channel and, as a result, the correspondence between his temperature-beliefs and the thermometer, would then itself be a product of the mechanism's informativeness. Many links intervene; but the information that the mechanism is an information channel is preserved throughout, and responsible for Bob's temperature beliefs' correspondence with the thermometer. There is, in other words, a *second* information channel, carrying the information that the thermometer is an information channel, and terminating in the conformity of Bob's beliefs to the thermometer, and thereby to the temperature.

But since the pamphlet was fabricated, the fact that the thermometer is an information channel played no role in the production of the pamphlet's contents, and so no role in the restoration of Bob's trust in the thermometer. The link between the informativeness of the thermometer and Bob's reliance on it is severed, notwithstanding the fact that his belief in its informativeness is based upon evidence. The second information channel described in the previous paragraph no longer exists.

I propose, then, that S knows that P iff S's belief that P is produced in such a way as to contain the information that P, and S's belief is produced that way precisely because its being produced that way contains the information that P. More precisely:

- 1) S's belief that P is produced in circumstances C as a result of S's being in state T, in which state S would, in C, believe that P if and only if P. (T is an information channel for P in C.)

- 2) The process that produced S's being in T is such that it would only produce S's being in T if T is an information channel for P (in C). (T is produced by the information that T is an information channel in C.)
- 3) The probability that circumstances C are realized given that S is in T is high.

Condition 1 ensures that S's believing that P contains the information that P. This is, roughly, Dretske's¹¹ condition on knowledge.¹² Condition 2 ensures that S would not believe that P unless the process that produced it is an information channel in circumstances C that in fact obtain.

Circumstances C are the so-called 'channel conditions' that are required in order for T to constitute an information channel, and so include whatever conditions, both within and outside S, must be met in order for S's belief to be a function of P. They include, in Bob's case, his determination to find out what the temperature is by reading the thermometer, his enjoying a clear, unobstructed view of it, his looking attentively at the thermometer, his not hallucinating, the thermometer's being accurately calibrated, the temperature's not being so low or high as to be beyond the thermometer's limits, and so on.

Condition 3 requires that these circumstances are in fact likely to hold—that the information channel will be open—if Bob does shape his temperature beliefs by the thermometer. This might be either because they hold generally whether or not Bob so shapes his beliefs or because Bob's so shaping his beliefs is somehow sensitive to their

holding, so that in either case he would typically believe what the thermometer says only if his resulting belief would be correct.

6. Understanding Bob's Problem

Condition 2 is not satisfied when the conformity of Bob's beliefs to the thermometer—that is, his being in state T—is due only to his leader's having told him to believe that the thermometer gives the temperature. Since Bob would have believed whatever the leader said, he would have conformed his beliefs to the thermometer whether or not the thermometer accurately reported the temperature.

Condition 2 is also not satisfied when Bob trusts the thermometer in response to having read the fabricated pamphlet. Since the pamphlet is fabricated, the contents of the pamphlet are not a function of the accuracy of the description therein, and so Bob's being in T is still not a function of the informativeness of T in C.

Condition 2 could have been satisfied in a wide variety of ways. Bob could have taken the thermometer apart and figured out what it does and how it works, and relied on it as a result. Or he could have read a (genuine) manufacturer's pamphlet. Or Bob might not himself have learned of the thermometer's informativeness in any way at all, but was instead somehow 'socialized' to rely on it, which socialization he would not have received if the thermometer were not informative.¹³ Or, Bob's reliance on the thermometer could even be the result of his genetic inheritance, an inheritance he would again only have received if the thermometer, or thermometers like it in evolutionary prehistory, were informative. This latter possibility is of course not the right one given

the absence of thermometers at the time our genetic inheritance was fixed, although it is realized in other cases, as we will soon see.

Bob himself is not, however, responsive to the thermometer's informativeness in any of these ways. And that is why he does not know what the temperature is.

7. Perception

This is a thoroughly externalist account. None of the conditions require that S be aware of, have access to, have evidence of, etc. the fact that the condition holds. Notice that this is so even if these conditions hold because S is responding to evidence, as would have been the case had the putative manufacturer's pamphlet been genuine. The externalist does not suggest that an agent's responding to evidence for P prevents her knowing that P, but only that her being so responsive is not a conceptual component of her knowing that P. Responsiveness to evidence may be a wonderful way to ensure that the conditions hold; but it is possible for them to hold in other ways.

It is therefore possible for them to hold in circumstances in which talk of evidence, reasons, and even justification is strained at best. If you are asked why you believe P and you say that you do so because you saw that P, that could reasonably be seen as the citation of evidence. But if asked why you believe what you (think you) saw, you are likely to have little to say; there is the feeling that the question is misplaced. Perception is, one might say, the source of evidence; how could there be evidence for it? And even if one thought the request legitimate, very few of us are in a position to explain

how perception *does* provide us with information; the neurophysiology involved, for a start, is beyond most of us, and even the neurophysiologists still have a lot to learn.

But on the present analysis none of this prevents your knowing what you know as a result of the operation of your perceptual system. That system is, when used in circumstances for which it is suited, a remarkably powerful source of information, and the product of evolution. Another product of evolution is your innate trust in your perceptual system as a source of information. Yet another product of evolution is the fact that you would not trust your perceptual system if it were not such a source. No increase in fitness accrues to an organism in virtue of having a powerful source of information unless the organism relies on that source in generating representations upon which it acts. So the fact that your perceptual system is informative is responsible for your relying on it. Moreover, the environments in which humans now find themselves are not so different from that of our evolutionary predecessors as to prevent the perceptual system that was informative then from being informative now. So true beliefs generated by that system will often satisfy conditions 1-3, and constitute knowledge when they do so.

8. Unjustified Knowledge

Because this account of knowledge is thoroughly externalist, it is possible to construct scenarios within which the believer has no evidence in favor of the claim putatively known, and even ample evidence against, which satisfy conditions 1-3. The threat of counterexample at a higher level, so to speak, therefore reappears.

Here is an example. The drug that Jason's doctor has prescribed for him has, she informs him, the unique side-effect of producing the vivid hallucination that snow is falling. Jason begins to take it the next day, walks outside, and observes what looks to be snowfall, notwithstanding its being July. He stands wondering for a few moments at the remarkable phenomenon until he recalls his doctor's explanation of the side-effect. But his experience of the snowfall is so vivid—he can see the individual flakes, catch them on his hands and tongue, and so on—that he can't help believing it all to be real. He decides that it is just a remarkable coincidence that it should snow just now, after his taking a drug that is known to produce hallucinations of snowfall.

Jason's belief that it is snowing is, presumably, unjustified. In light of the fact that the snow he (seems to) see and feel is falling in July, in conjunction with the doctor's explanation of the side effect, the hypothesis that he is hallucinating is far and away the more reasonable explanation of his experience.

As it happens, however, the doctor has confused her drugs; the drug she prescribed actually produces no hallucinations. The snowfall that Jason sees, and believes to be falling as a result of seeing it, is a real meteorological aberration.

Jason satisfies conditions 1 and 2. His belief is produced by the usual perceptual process operating successfully, so that his seeming to see snow is a function of the existence of real snow; and he trusts his perceptions in response to the same evolutionary history that produced the trust that you and I put in our perceptions (although he trusts them somewhat too implicitly).

Condition 3 is satisfied as well. Notice that not only is Jason not hallucinating, he would also not be likely to be doing so. He would only be likely to be hallucinating

(whether he was or not in this particular case) if the doctor was right about the drug's side-effects; but the doctor was wrong.¹⁴

Does Jason know that it is snowing? Intuitions at this point, I suspect, are far less reliably one-sided, so that it is at the very least not obvious that the internalist now has a counterexample in her favor. There are, moreover, reasons for favoring the claim that Jason does know that it is snowing in this case, even though his belief that it is so is unjustified.

Suppose that, instead of recalling his doctor's warning of the side-effect but discounting that warning, Jason initially forgot the warning entirely, along with the entire visit to his doctor the previous day. This memory lapse is, say, a real side-effect of the drug, as a result of which Jason's memory of the previous day has been entirely erased.

In such circumstances Jason surely does know that it is snowing. What was erased from his memory was, after all, misleading information. Aside from that unfortunate past history now forgotten, he is in precisely the situation that you or I would be in if we were to see the snowfall.

The internalist has, moreover, good reason to count this as a case of knowledge. For the memory erasure has the effect that Jason is now justified in the internalist's sense, since his memory of the previous day is now, thanks to the drug's real side-effect, not demon-invariant, and so cannot undermine his justification now for his belief that it is snowing. Since the internalist's point is supposed to be that the evidence against his perceptual veracity—what the doctor told him—prevented Jason from knowing when he remembered but discounted that evidence, she can hardly deny that Jason is justified, and

so knows, when no such misleading evidence against his perceptual veracity is available to him.

So, then, if we say that in the one case—when Jason remembered but discounted his doctor’s warning—he did not know that it was snowing, but that in the other case—when he did not remember the doctor’s warning—he did know that it was snowing, then we would have to say that a person who does not know that P can come to know that P by forgetting evidence against P, and conversely, that a person who does know that P can cease to know that P by remembering evidence against P. And if Jason, when he did remember his doctor’s warning, thereby knew what his doctor said—which he surely did if memory can provide us with knowledge at all—we would in addition have to say that a person who does not know that P but does know that Q can come to know that P simply by ceasing to know Q, and also that a person who does know that P but does not know that Q can cease to know that P simply by coming to know Q.

That, surely, is not right. One cannot come to know something new simply by the elimination of some other piece of knowledge one already has. Nor can coming to know something new in itself eliminate knowledge currently in one’s possession. But these consequences can be avoided only if we concede to Jason the knowledge that it is snowing when he remembered but discounted his doctor’s warning, notwithstanding the fact that his belief that it is snowing is, in that case, unjustified.¹⁵

9. Science Fiction and Accidental Informativeness

Reflection on this case allows us to explain an odd feature of the counterexamples that one does find in the literature. Lehrer's Truetemp, Bonjour's Norman, and the rest are all cases in which the externalist condition is met through an *unusual, bizarre, unlikely* process. But, given the purpose of the examples, there is no reason why this ought to be the case. The examples are only supposed to be ones in which, first, the belief is caused by a reliable process (or information-generated, or whatever externalist condition is under attack), and second, the believer has no evidence for believing, or ample evidence against believing, that this is so. But a run-of-the-mill informative process like perception, along with evidence that it is not informative, as in the example above, ought to do nicely. So why the science fiction?

The answer, I suggest, is that the science fiction smuggles in failure of condition 2. Truetemp's believing the beliefs he does has nothing to do with the fact that being-a-Truetemp realizes an information channel. Norman's relying on his clairvoyant beliefs is not itself a result of that process's being informative. And so on. That they do not satisfy condition 2 is what makes the examples feel, somewhat paradoxically, like cases of informativeness-by-accident. The informativeness *is* accidental, in that the agent's cognitive system is in no way shaped by what belief-forming processes are information channels, and so when that system lands on one that is informative it really does do so entirely by accident.

But that is precisely what is ruled out by condition 2. Its satisfaction ensures that it is only in response to the information that the process is an information channel that the agent relies on it in forming her representations of the world. Jason's reliance on his perception does meet this condition; and that, I suggest, is both why no strong intuition to

the effect that Jason lacks knowledge results, and why internalists avoid constructing such examples as cases against externalism.

That condition 2 rules out such ‘accidental informativeness’ is, I suggest, the reason why it is required before a belief counts as known. It is not enough to characterize the externalist’s condition as one which only requires that the cognitive system act on the basis of information, since that requirement provides no assurance that the system is responsive to what processes *are* information channels from which representations should be produced. The problem with Bob is ultimately that even if he is lucky enough to instantiate an information channel such luck is rare: beliefs that ultimately result from indoctrination are very unlikely to be information-produced. Cognitive efficiency requires a tighter relationship with information than that.

10. Justification and Cognitive Flexibility

With many organisms the tighter relationship that is ensured by satisfaction of condition 2 is a product of evolution and fixed within that organism. The frog’s response to moving spots of light—lashing out its tongue—is the product of evolution, a reaction instantiated because of the information that such a moving spot might convey concerning the location of a tasty insect. But a particular frog is a prisoner of this information channel; his reliance on it is inevitable, even if a change in circumstance should render it no longer informative.

But we are more flexible than frogs, and capable of responding to the degradation, or the limits, of old information channels, and the availability of new ones. This

flexibility has the obvious advantage that we can function effectively in a much wider variety of environments. But it comes at a price. The ability to disengage from one channel that has degraded or whose limits have been reached in order to engage with a more reliable one raises the possibility that other forces might lead us to engage instead with a less reliable channel, or with no information channel at all.

There is thus considerable value in external pressure being applied to the agent to seek and employ the best information channels available, and even more value in the agent's applying such pressure on herself. The deontological quality of the demand that your beliefs be justified—that many have identified as underlying justification's internalist character¹⁶—amounts, I suggest, to that sort of pressure. It is exerted on you by your fellow epistemic agents who, after all, often rely on you as one of their sources of information. And, given a suitable epistemic upbringing, it is a pressure you have come to exert on yourself.

The demand for justification, so understood, is thus in the service of ensuring that flexible cognitive agents such as ourselves do not stray from the doxastic paths traced out by the channels of information open to them. Justification is, so conceived, not a conceptual constituent of knowledge, but is instead a means to it, one that is indispensable for such creatures as we.

Department of Philosophy

Western Michigan University

References

- Alston, William. (1986) "Internalism and Externalism in Epistemology," *Philosophical Topics* 14.
- Beebe, J. R. (2004) "Reliabilism, Truetemp and New Perceptual Faculties", *Synthese* 140(3), 307-329.
- Bonjour, Laurence. (1980) "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* 5.
- Conee, Earl and Feldman, Richard. (1998) "The Generality Problem for Reliabilism," *Philosophical Studies* 89, 1-29.
- Dretske, Fred. *Knowledge and the Flow of Information* (Cambridge, Mass.: MIT Press, 1981).
- Foley, Richard. (1985) "What's Wrong with Reliabilism?" *The Monist* 68, No. 2, 188-200.
- Goldman, Alvin. (1976). "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73, 771-791.
- Lehrer, Keith. *Theory of Knowledge* (Boulder: Westview, 1990).
- Plantinga, Alvin. *Warrant: The Current Debate and Warrant and Proper Function* (Oxford: Oxford University Press, 1993).

¹ See esp. Dretske 1981.

² The distinction between access and perspectival internalisms originates in Alston 1986.

³ Alvin Plantinga calls whatever transforms true belief into knowledge (which is not, he thinks, justification) ‘warrant’ in the interests of avoiding the deontological and internalist connotations of ‘justification.’ I am not sure why ‘warrant’ doesn’t share those connotations, however, and would rather avoid both terms. See Plantinga 1993.

⁴ Since externalists reject the internalist’s constraint on justification or knowledge, they need to put something in its place. What they each do put in its place, while admitting of substantial variation, display a family resemblance centered on a relation between the agent and her environment (in particular, a ‘natural’ relation, amenable to scientific investigation), one which somehow contributes to her beliefs’ being correct. So there is this much in positive character in common among externalists, as well as the negative ‘not-internalism’ character above. But this strikes me as a ‘contingent’ commonality that just happens to be shared by most (but not all) externalists, and that denial of internalism is the core. Nothing, however, turns on this here.

⁵ Another species of counterexample—an unwitting victim of a Cartesian evil genius, for example—purports to show that beliefs can be justified notwithstanding failure of the externalist’s requirements and therefore that externalism is unnecessary (See Foley 1985, for example). But the intuition that these beliefs are justified is matched by an equally strong one that they are not known. So the counterexamples do not count against externalist accounts of knowledge. As will be apparent later in this essay, I endorse an internalist view of justification. These counterexamples are therefore no threat to the view developed here.

⁶ See, for example, Bonjour 1980 and Lehrer 1990.

⁷ See, for example, Bonjour 1980.

⁸ Lehrer 1990.

⁹ Bonjour 1980.

¹⁰ Concerns have been raised concerning the realizability of the Truetemp story, for example (Beebe, 2004). More generally, the beliefs generated are often described as ones the agent ‘finds within himself,’ or that ‘well up’ within him, as though discovering that one has a belief that one can’t help but believe can be like

discovering that one has, say, a third foot. One might wonder whether beliefs can really be like this. The issue is avoided by the examples I present here.

¹¹ See Dretske 1981.

¹² Roughly, but not precisely. Dretske requires that S's belief that P be caused by the information that P. This does not, however, ensure that S's believing that P itself represents the information that P, although this is clearly Dretske's intent. It is compatible with T's being caused by R, where R represents the information that P, that if R had not occurred, T would have been produced anyway by another event W that is incompatible with P. If that is so R would still have been caused by the information that P but would not itself represent that information.

¹³ Notice that this is the case for most of us when it comes to trusting thermometers.

¹⁴ Jason might remind the reader of Goldman's barn-facades example, in which an agent correctly takes herself to perceive a barn, but where the unusually high proportion of false barn-facades to real barns in the region of the country in which the agent finds herself undermines her claim to knowledge (Goldman 1976). But the similarity is superficial, because the belief of Goldman's perceiver does not (and Jason does) satisfy condition 3. Since no more of the object is seen than its presenting surface, where the belief induced concerns the kind of object whose surface it is, condition 3 requires that the agent's environment be one in which such surfaces are, at least typically, of objects of the sort the agent perceptually judges it to be (which is false in Goldman's example). The barn-facades case would parallel Jason's situation only if the drug Jason took was in fact likely to produce hallucinations. But this is not true of the drug Jason actually took.

¹⁵ Jason does not, however, know that he knows that it is snowing. For it would then have to be the case that he would not believe that he knows this if he did not in fact know this. But Jason has demonstrated an insensitivity to evidence against his perceptual veracity which, although misleading in this case, indicates that he would still have believed that his perception of snow is veridical when it wasn't. Jason's belief that he knows would therefore violate the first condition: his belief that his perception is accurate is not a function of its accuracy.

¹⁶ See Alston 1986 and Plantinga 1993.