

DEDUCTIVE CLOSURE AND RELEVANT ALTERNATIVES

Ted A. Warfield

Responding to skepticism is one of the most important, if not the most important, tasks of epistemologists. In section one of this paper, the version of skepticism arising from the problem known as the epistemological problem of closure will be presented.¹ In section two, I introduce what I take to be the standard response to this problem. Finally, in section three, the consequences of accepting this response will be explored, with special attention given to second-order knowledge.

Our subject for this paper is Steve. Steve is a normal human being with excellent vision. He is sitting at his desk with his eyes closed thinking about his schedule for the day. He opens his eyes and has a visual experience that would normally be described as "seeing a book." It is also true, let us say, that lighting conditions are normal, that Steve is in no way impaired by hallucinogens, and that there is a fairly large book in front of him, only four feet away.

If there is such a thing as perceptual knowledge, Steve seems to be an ideal candidate for having it. Given the specified conditions, it seems that Steve knows the following proposition:

A. There is a book in front of me.

The following, however, is something that Steve most certainly does not know:

B. I am not being deceived by a Cartesian demon into believing there is a book in front of me when in fact there is not.

The epistemological problem of closure now comes into

the picture, for "A" surely entails that "B." This principle can be stated as follows:

C. If S knows that P and knows that P entails Q, then S either knows that Q or could come to know Q simply by reflecting on P and the entailment.

To make the case of Steve fit the closure principle, we need only add that Steve knows that:

D. There being a book in front of me entails that I am not being deceived by a Cartesian demon into believing there is a book in front of me when in fact there is not.

Given "D," it is clear that, if we accept the closure principle, we will be forced to admit that Steve does not know "A." The following argument illustrates this:

P1	$[Ksa \ \& \ Ks(a \rightarrow b)] \rightarrow K^*sb$	The Closure Principle
P2	$\sim K^*sb$	The Demon
C1	$\sim [Ksa \ \& \ Ks(a \rightarrow b)]$	Modus Tollens P1, P2
C1'	$\sim Ksa \vee \sim Ks(a \rightarrow b)$	DeMorgan's Law C1
P3	$Ks(a \rightarrow b)$	Stipulated
P3'	$\sim \sim Ks(a \rightarrow b)$	Double Negation P3
C2	$\sim Ksa$	Disj. Syllogism C1', P3'

{NOTE: Read K^* as "knows or could come to know by reflecting."}

If we accept the Closure Principle (a highly intuitive principle), we will, it seems, be forced to admit that Steve does not know that there is a book in front of him.² In other words, we will have to accept skepticism about perceptual knowledge.

The appeal of the Closure Principle is considerable. The skeptic, it seems, finds this principle even more plausible than the assertion that Steve knows that there is a book in front of him. Nonskeptics would accept, at least at first glance, that one knows what one can deduce from things that one knows, that is,

they would intuitively accept the Closure Principle. For example, if Steve knows that there are two books on his desk, then he also knows (actually K*) that there is (at least) one book on his desk.

One who wants to maintain that Steve knows there is a book in front of him would ideally be able to give an account of knowledge on which he does. And, ideally, this account will allow for knowing things like the proposition deduced in the preceding paragraph. Fred Dretske gives us something like this in his paper "Epistemic Operators." In this paper, Dretske restricts the Closure Principle by introducing the notion of "relevant alternatives."³

II

Palle Yourgrau recognizes Dretske's important contribution to epistemology in the following quotation from Yourgrau's "Knowledge and Relevant Alternatives":

Traditionally, skeptics as well as their opponents have agreed that in order to know that 'p' one must be able, by some preferred means, to rule out all the alternatives to 'p.' Recently, however, some philosophers [following Dretske] have attempted to avert skepticism not (merely) by weakening the preferred means but rather by articulating a subset of the alternatives to 'p'—the so-called relevant alternatives—and insisting that knowledge that 'p' requires only that we be able (by the preferred means) to rule out members of the set. (175)

The details of Dretske's paper need not concern us, for in introducing the notion of relevant alternatives, Dretske makes an unconvincing attempt to give an account of why the closure principle fails. His arguments are not persuasive and have been discussed in much of the literature, so I will not concern myself further with them.⁴ Though Dretske's account of the

failure of the Closure Principle has been laid to rest, the idea of "relevance" is still with us today as a popular response to the skeptic.⁵

Unfortunately, "relevance" is perhaps the most ambiguous concept in epistemology today.⁶ Alvin Goldman has this to say: "The brain-in-a-vat alternative just is not a *relevant* alternative. (I do not, however, have a detailed theory of relevance)" (55). Dretske says:

A relevant alternative is an alternative that might have been realized in the existing circumstances if the actual state of affairs had not materialized. (1021)

He admits that this "might" is vague, and he is able to add only that it will have to be explained in terms of counterfactuals.

The account of knowledge that comes out of this idea of relevance is, despite the ambiguities, plausible. S knows that A only if:

- F. S believes that A.
- G. It is the case that A; and
- H. S is able reliably to discriminate A from its relevant alternatives.⁷

This account follows the spirit of considerations made by Goldman in noting that different propositions require differing levels of discriminatory ability:

To know there is a keyboard before me, I needn't discriminate this state of affairs from envatment. But to know I am not envatted, I do need to discriminate this state of affairs from envatment. So I may not know I am not envatted. Yet this possibility does not preclude my knowing there is a keyboard before me. (55)

This account also seems to fit with the project I claimed

was needed for one who wanted to hold that, in our original case, that Steve knows that "A." For in that case, which is the case in which Steve believed that there was a book in front of him, there was in fact a book in front of him, and he could surely discriminate this from the relevant alternatives (there being a telephone, a pencil, or other things he would "normally" find on his desk in front of him). And not only is knowledge possible on this account, but one can also know the deduced consequences of what one knows, so long as one confines one's deductions to the framework of relevant alternatives.

Of course, the vagueness of "relevance" presents problems. For example, if in fact Steve's visual process is being toyed with by a Cartesian demon, then his being deceived by such a demon is a relevant alternative, even if at that particular moment, the demon is not giving him a false belief. Likewise, what if fake books (things that are indistinguishable from actual books at the distance from which Steve is viewing the actual book) have been placed all around the office building?⁸ If this is so, it seems that Steve is only *accidentally* seeing a real book—and, generally speaking, knowledge can't be accidental. To illustrate further the vagueness of "relevance," consider the following cases involving fake books in which it is difficult to decide if the fake book alternative is relevant:

- J. There aren't any fake books in the building today, but there were yesterday (or a year ago).
- K. There aren't any fake books in this building, but there are some next door (or a hundred miles away).

Other borderline cases could be presented, but I think the point about the vagueness of relevance is clear. Now I would like to talk a little more about the status of relevant alternatives views in epistemology today.

There are at least three points on which the relevant alternatives view has been challenged: content, necessary truths, and intuitive considerations. Regarding content, Yourgrau argues that the content of what a person knows is, on

the relevant alternatives view, different than the proposition that 'P.' He argues that the content on the relevant alternatives account is something like: "IF some set of relevant alternatives is the actual set *then* 'P'" (179).⁹

Regarding necessary truths, Brueckner has noted that relevant alternatives accounts may have a difficult time incorporating knowledge of necessary truths. This is because on most accounts of the relevant alternatives of contingent truths, something must be logically possible to count as a relevant alternative. At the very least, Brueckner argues, considerations of necessary truths will "serve to constrain accounts of the notion of a relevant alternative" (n36).

The intuitive consideration that some feel weighs against the relevant alternatives approach is that what it is to have knowledge on this account just is not what it means to have knowledge. Knowledge on the relevant alternatives account is "weaker" than what is really meant by knowledge—"weaker" in some way other than Yourgrau's content worry.

It is this intuitive consideration that I wish to try to flesh out a bit in the remainder of this paper. My plan is to explore the possibility of second-order knowledge on the relevant alternatives account and to argue (provisionally—see note 13) for its impossibility. If this argument is successful, we will have some reason to think that the relevant alternatives approach shifts ordinary skepticism from ordinary knowledge claims up to the level of "knowing that one knows." One must then decide if this result is due to a weakness in the skeptic's argument or to the allegedly "watered down" view of knowledge that the relevant alternatives view promotes.

III

Given the account of knowledge spelled out in conditions F, G, and H, and duly noting the vagueness of relevance, I turn now to the following question: Does or can Steve know that he knows that "A"? The way to get nearer the answer to this question seems to be by applying the requirements for knowledge to "He knows 'A'" instead of to just "A." This gives us

the following:

Steve knows that he knows "A" only if:

- R. He believes that he knows there is a book in front of him.
- S. He in fact knows that there is a book in front of him.
- T. He is in fact able to discriminate his knowing that there is a book in front of him from the relevant alternatives.

"R" and "S" seem harmless enough, but how exactly is one to understand "T"? Given the already demonstrated vagueness of "relevance," it seems as if we are in for trouble here. What are the relevant alternatives for Steve *knowing* there is a book in front of him?

There seem to be two routes to explore in spelling out condition "T." One route requires the candidate for knowledge to stand in a more favorable position regarding the relevant alternatives of the first-order case. The second route requires that the person *know* that the process that he is using is reliable (as opposed to the first-order case where the person merely had to be using such a process).

The intuition behind the first route is linked to the *relevant alternatives* of the first-order case. The intuition behind the second route, as already touched on, is linked to the *process* used in the first-order case. Looking at the first route, we seem to have two choices for spelling out condition "T":

- U. The relevant alternatives for "knowing there's a book in front of him" are the relevant alternatives for "there being a book in front of him" with the prefix "knowing" affixed to them; or,
- V. The relevant alternatives for "knowing there's a book in front of him" are such that he must know that in the first-order case, no alternative that he could not discriminate was relevant.

The problem with "U" is that, if that is all there is to it, then Steve knows that he knows that "A" if he knows that "A" and believes that he knows that "A." The main point here is that in knowing that "A" (the first-order case) one eliminates all the alternatives of "U." This is due to the fact that in the first-order case, one must be able to rule out the relevant alternatives; and, if one can do this, then one can discriminate necessary conditions for the relevant alternatives that "U" suggests. So, if one knows and believes that he knows, then he knows that he knows. But second-order knowledge must require something beyond mere first-order knowledge, or the account in question suffers from level confusion.¹⁰

If we adopt this first route, we must, if we are to avoid having first-order knowledge entail second-order knowledge, adopt "V" as our explication of condition "T." Doing so, however, tells us that Steve must, if he is to know that he knows that "A," know that no alternative that he could not discriminate from "A" (such as B) was relevant. This, I am sure we can agree, is beyond the capacity of normal humans. So, while second-order perceptual knowledge is *logically* possible on this account, it is humanly impossible.¹¹

Does the second route leave us any better off? First, let's look at an example of it. In *Epistemology and Cognition*, Goldman gives the following characterization:

To know that we (sometimes or often) know, we would have to know that we (sometimes or often) use reliable processes of belief formation. But since the analysis makes it (logically) possible for us to know what processes we use, and makes it (logically) possible for us to know all sorts of truths about the world (which is essential for knowing the reliability of our processes), the analysis makes it logically possible for us to have higher order knowledge. (56-57)

Allow me to note that I do not take this to be the route that Goldman himself prefers.¹²

What this account seems to suggest is this: we can know what processes that we use, and we can know many truths about the world; therefore, we can assess the reliability of a given process. As an example, take a generic process "Z." All we need to do to assess the reliability of process "Z" is to see (roughly) what percentage of the beliefs we form using process "Z" are true. Can we do this?

I believe the answer is "No." We are unable to do this because having only first-order knowledge does not give us enough information. We can indeed know many truths about the world, but our doing so does not allow us to assess statistical reliability. For example, let's say that I'm trying to assess the reliability of process "Z" and I know that process "Z" led me to a belief that "Y." In deciding whether this token counts as a success or a failure of process "Z," I need to know one thing: Is "Y" true? To know that my process is reliable, I need to know that "Y" is true (as well as a significant percentage of my other beliefs). But if all I have is first-order knowledge, then while this means that it is the case that "Y" (knowledge entails truth), I don't have the higher-order knowledge that "Y is true." And if I don't know whether *Y is true*, I can't assess statistical reliability.

If this analysis is correct, it seems that the second route leaves us no better off than did the first. We cannot know that we know, and therefore we don't.

In his paper "Level Confusions in Epistemology," William Alston claims that one can reach this conclusion only with the dubious assumption that "one cannot know that 'p' unless one knows, with respect to each of the necessary conditions of 'p,' that it obtains" (146-47). But it seems that the conclusion can be reached by my more direct route. In applying the relevant alternatives account to "knowing that p" one must avoid the conclusion that first-order perceptual knowledge plus a belief that one has this knowledge entails second-order perceptual knowledge. This can be done, but only at the cost of sacrificing the possibility of our having such second-order knowledge.

As I mentioned in introducing the intuitive considerations

against the relevant alternatives approach, it is unclear whether this resulting skepticism about higher-order knowledge demonstrates that skeptics about ordinary knowledge have committed level confusion or that the relevant alternatives account of knowledge is not *really* an account of knowledge.

It seems that things turn back to the consideration of Goldman regarding envatment. The relevant alternatives theorist will affirm that intuition; while, the skeptic will simply point out that his or her intuition is different. What is needed at this point in the debate is a demonstration by the relevant alternatives theorist that higher-order knowledge is, contrary to my argument, possible.¹³ This would show that, even if the suggested account of ordinary knowledge is too weak, the relevant alternatives theorist hasn't merely refused to answer the skeptic. Also, the relevant alternatives theorist needs to develop a firm response to the worries of Yourgrau and Brueckner that I sketched early in the paper.

The skeptic, on the other hand, needs to develop a positive line of attack, perhaps along the lines that Brueckner and Yourgrau suggest. Failure to argue for positive theses leaves the skeptic with only intuitions, and intuitions never convince anyone of anything.

NOTES

¹My formulation of this problem is an offshoot of that of Thomas Senor of the University of Arkansas. I would like to thank him for helpful discussions leading up to my paper. I am particularly indebted to him for what in my paper is "C" and for the corresponding symbol "K*." These make it clear that I am not concerned with cases in which S doesn't know that p entails q or cases in which S does not reflect on these facts and thus fails to "put things together."

²Attempts have been made at denying P2. Gail Stine's article is one such attempt, as is the infamous "G. E. Moore

Shift." One might also take a positivist line here, denying that the demon hypothesis expresses a possibility.

³Though he does not directly address the question in "Epistemic Operators," I believe that Dretske accepts the closure principle with respect to the relevant alternatives.

⁴See Gail Stine, "Dretske on Knowing the Logical Consequences," *Journal of Philosophy* 57 (May 1971): 249-61; Yourgrau; and my "Epistemological Closure."

⁵As Yourgrau points out, philosophers such as Swain, Harman, and Goldman make use of this notion in varying degrees.

⁶Mark Heller has analyzed the notion of "relevance" by comparing it to the subjunctive conditional approach of Robert Nozick.

⁷As Heller points out, one does not need to link the relevant alternatives approach with reliabilism, but this seems to be both a common and a promising approach.

⁸This and what follows (J and K) are adaptations of Goldman's famous barn case.

⁹For an attempt at responding to Yourgrau's worry, see Brueckner n31.

¹⁰See Alston, "Level Confusions."

¹¹Second-order perceptual knowledge is logically possible, requiring that the perceiver satisfy R, S, and T (as interpreted by V). God, it might be argued, could satisfy these requirements.

¹²It has been suggested that Goldman might argue that one could come to know that a particular belief-forming

process is reliable without having to "do the calculation" herself. One could form the belief that "I know my process is reliable" in a reliable manner and thus (perhaps) know that she knows that "A." This account faces some prima facie circularity problems—the agent would seemingly have to assume that her perceptual processes are reliable to form the belief that it is (she would *read* that it is reliable in a psychology journal for example). For a fuller discussion of this type of difficulty, see Alston, "Epistemic Circularity," reprinted in Alston, *Epistemic Justification*.

¹³Perhaps the consideration raised above in note 12 is the most promising route to explore. This is in fact the direction that I will take in my continuing work on this issue. I believe a route like the one suggested in note 12 will show that second (and higher) order knowledge is in fact possible on the relevant alternatives account, and thus the skeptic will be hard pressed to maintain that the relevant alternatives account is an overly weakened account of knowledge. My aim is to use this as the foundation of an anti-skeptical argument that does not rely on a "contextual" or "linguistic community" account of knowledge.

WORKS CITED

- Alston, William P. "Epistemic Circularity." *Philosophy and Phenomenological Research* 47, no. 1 (September 1986): 1-30.
- . *Epistemic Justification*. Ithaca, NY: Cornell UP, 1989.
- . "Level Confusions in Epistemology." *Midwest Studies in Philosophy* 5 (January 1980): 142-48.
- Brueckner, Anthony. "Skepticism and Epistemic Closure." *Philosophical Topics* 13, no. 3 (1985): 89-117.

- Dretske, Fred. "Epistemic Operators." *The Journal of Philosophy* 24 (December 1970): 1007-23.
- Goldman, Alvin. *Epistemology and Cognition*. Cambridge: Harvard UP, 1986.
- Heller, Mark. "Relevant Alternatives." *Philosophical Studies* 55 (April 1989): 23-40.
- Stine, Gail. "Dretske on Knowing the Logical Consequences." *The Journal of Philosophy* 57 (May 1971): 249-61.
- . "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29 (April 1976): 249-61.
- Warfield, Ted A. "Epistemological Closure." Unpublished.
- Yourgrau, Palle. "Knowledge and Relevant Alternatives." *Synthese* 55 (January 1983): 175-90.