

2006

Dretske on Knowledge Closure

Steven Luper

Trinity University, sluper@trinity.edu

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



Part of the [Philosophy Commons](#)

Repository Citation

Luper, S. (2006). Dretske on knowledge closure. *Australasian Journal of Philosophy*, 84(3), 379-394. doi: 10.1080/00048400600895862

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.

Dretske on Knowledge Closure

Fred Dretske has long criticized the claim that empirical knowledge is closed under (known) entailment. He rejects any closure principle, however qualified, implying that, by knowing commonsense claims (such as *I am in Boston*) we are positioned to know that skeptical scenarios do not hold, even if the former entail the latter. Among the closure principles Dretske rejects is the following:

K: If a subject S knows p , and S believes q because S knows that q is entailed by p , then S knows q .

Dretske developed his main case against K (and restricted versions of K) in his early essays [1970, 1971]. In recent work [2003] he offers fresh reasons for rejecting closure, and revives aspects of his previous case, saying that given a proper understanding of perceptual knowledge we will reject K.

Dretske's main case against knowledge closure is well understood. I will briefly review it and his recent objections, but my main concern here is to examine Dretske's attempt to indict knowledge closure on charges he also levels against the closure of perception and perceptual knowledge. A good deal in his case is correct and important. In fact, I will argue, understood one way, perceptual knowledge is not closed under entailment; however, once we understand why, we are left with no reason to reject knowledge closure.

So the plan is this: first we glance at Dretske's old critique of closure, then his recent objections, then we examine the case he bases on his account of perceptual knowledge.

The Old Case

Dretske's main case against K has two components. The first rejects K on the basis of a relevant alternatives account of knowledge. The other says rejecting K is the way to resist skepticism.

First argument. On a relevant alternatives account 'knowledge that p requires one (or one's evidence) to exclude not all, but only all relevant, alternatives to p .' According to Dretske, such an account leads "naturally" but "not inevitably" to a failure of knowledge closure. For it "commits" us to saying that

the evidence that (by excluding all relevant alternatives) enables me to know my wife is on the sofa does not enable me to know that it is not a cleverly disguised imposter since, in most circumstances at least, this is not a relevant alternative [2003: 112-3].

As Dretske acknowledges [2003: 112], the relevant alternatives account is a weak basis for denying K, for there are ways to reconcile the two [e.g., Stine 1976; Cohen 1988; Lewis 1996]. One way this can be done is as follows [Luper 1984; 1987].

First, we say that S's knowing p is, roughly, S's arriving at the belief p on the basis of a reason R such that:

SI: If R held, p would be true.

When condition SI holds, let us say that R is a *safe indicator* that p is true [for refinements, see Luper 1987; 2003; for a similar approach, and the introduction of the term 'safety,' see Sosa 1999; 2003; for a defense against recent objections, see Luper 2006]. Condition SI says p holds throughout the R worlds (worlds in which R holds) that

are close to the actual world. The actual world is S's situation as it is at the time S arrives at the belief p via R [for clarification of 'near possible world,' see Stalnaker 1968; Lewis 1973; Nozick 1981: 680 note 8]. In adopting the safe indicator analysis, we position ourselves to endorse K, since R safely indicates p only if R also safely indicates q , where q is anything entailed by p .

Second, we say that an alternative to p , A , is *relevant* (relative to S and S's situation) if and only if:

RA: In S's circumstances, A might hold (i.e., it is false that: given S's circumstances A would not hold).

On this view, no alternative to p that is remote is relevant.

Third, we consider A ruled out on the basis of R when and only when:

SIR: were R to hold A would not hold.

This way of understanding relevance dovetails with the safe indication account without raising any challenges to K. Suppose S knows p via R. Then not only can S rule out p 's alternatives, relevant or otherwise, but q 's too, where p implies q . For assuming S knows p via R, and p entails q , then if R held, the alternatives to q would not.

Moreover, our way of understanding relevance gives us intuitive results, as we can illustrate using the Ginet-Goldman [1976] papier-mâché barn case. In this example, I believe *barn*: there is a barn in front of me. However, as I look at the (real) barn in front of me, the possibility that I am confronting a fake is relevant, since I am surrounded by papier-mâché barns that look just like the real thing: in close possible worlds in which I have the visual impressions which led me to believe a barn is there, I am seeing a fake. By contrast, when I am in a part of the world that does not feature fake barns and the like,

and I see a barn, the possibility that I am confronting a fake is not relevant, since in no nearby possible world am I confronted with a fake. So I know in the latter, un-Gettierized situation and not in the former, Gettieresque situation.

On Dretske's own view, an alternative A to p is relevant if and only if:

If *not- p* , A might hold.

Accordingly, alternatives to p that are remote are *almost* never relevant. For example, in our un-Gettierized situation when I believe *barn* the possibility that I am seeing a fake is not relevant. But for Dretske *not- p* is automatically a relevant alternative to p no matter how remote *not- p* is. Hence his account of relevance is not consistent in ruling remote possibilities as irrelevant. Being more consistent in this respect counts in favor of understanding relevance in terms of RA, and since RA, combined with the safe indication account of knowledge, preserves closure, it is hard to see that Dretske can plausibly claim that the relevant alternatives account leads 'naturally' to a failure of knowledge closure.

(But consider the set of impressions, E , that, in the un-Gettierized barn case, constitutes my evidence for believing *barn*. Is it plausible to say that E is also evidence for thinking *not-fake*: it is false that I am confronting a fake barn and no real barn is present at all? If not, doesn't this suggest, as Dretske thinks, that evidence is not closed under (known) entailment: we can have good evidence for a claim without having good evidence for each of its consequences? And doesn't this, in turn, suggest that we ought to reject K ? Admittedly, it does seem dubious to say that evidence is closed under entailment. I will not try to settle this issue here. But why say that knowledge is closed only if evidence is closed? A traditional foundationalist line of thought might appeal to some version of the following principle:

I know q on the basis of p only if I am warranted in believing p on the basis of a given set of evidence, and that evidence also warrants me in believing q .

If this foundationalist principle were true, knowledge really could be closed only if evidence (or rather warrant) were. But Dretske is an externalist, and rejects such traditional principles. Moreover, he provides no reason whatever for the view that knowledge closure depends on evidence closure. In fact, he appears to reject such a dependence when he says that ‘perceptual knowledge, if we have it, derives from the circumstances in which one comes to believe, not one’s justification for the belief’ [2003: 105]. What enables me to know *barn* is the fact that E is a conclusive reason for *barn*, not that E is good evidence for *barn*. Thus, on his view, coming to know q , where q is a consequence of *barn*, is not a matter of E’s being good evidence for q .)

Second argument. Dretske also argued, much as Robert Nozick did later [1981], that his account of knowledge allows us to come to terms with skepticism, but only if we reject K and similar principles. Dretske said, roughly, that knowing p is a matter of having a conclusive reason for believing p , where R is a conclusive reason for believing p if and only if:

CR: were p false, R would not hold.

Accepting the conclusive reasons account allows us to make a concession to skepticism—namely, we do not know whether typical skeptical scenarios hold—without succumbing to it: we still know ordinary knowledge claims, such as that I am in Boston, since nothing stops us from having the requisite conclusive reasons. But if knowing is a matter of having conclusive reasons, K fails, since having a conclusive reason for believing something does not guarantee having a conclusive reason for believing p ’s

consequences. For example, my sensory information gives me a conclusive reason for believing I am in Boston, but not for believing that my sensory input is not being fed to me by aliens who have taken my brain to the planet Crouton and put it in a vat.

This second line of attack has met with considerable resistance. Several theorists [e.g., Bonjour 1987; Fumerton 1987; Feldman 1999] have claimed that *any* argument against K should be rejected because its conclusion is absurd. Theorists [Cohen 1988; Lewis 1996] have also pointed out that there are good strategies for resisting skepticism without rejecting K. One example [Luper 1984; 2003] is this. If we accept the safe indication account (that is, if we grant that it provides sufficient conditions for knowledge), we will also accept K, primarily because something can safely indicate p only if it also safely indicates the truth of p 's consequences. Moreover, the appeal of skepticism is easily understood: the safety of a reason is much like the conclusiveness of a reason; in fact, SI is the contraposition of CR [Luper 1984]. So CR and SI are easily confused, and even if we accept SI, we might believe that we do not know we are not brains in vats on far Crouton because we notice we do not meet CR, and fail to notice that we *do* meet SI. On other occasions, when we embrace K, we will have SI in mind, which sustains closure, not CR, which does not. Accepting the safe indication account gives us the advantages of the conclusive reasons view—we explain (away) the appeal of (one form of) skepticism without succumbing to it—without the disadvantage of having to reject closure.

The Recent Objections

So much for Dretske's old critique of closure. His newer strategy is summed up in the following passage:

Not only is [perception closure] false, none of our nonperceptual ways of coming to know, none of our ways of preserving knowledge, and none of our ways of extending it are closed under known implication. . . .If all this is so, if none of our ways of knowing, extending knowledge or preserving knowledge are closed, it seems odd to suppose that knowledge itself is closed. How is one supposed to get closure on something when every way of getting, extending and preserving it is open [2003: 113-4]?

Not only is perception not closed, many other items relevant to knowledge are not closed either; Dretske goes on to provide a 'small sample' of these: testimony, proof, memory, indication, and information. I will discuss perception closure in the next section. Putting it aside for now, what should we make of Dretske's suggestion that when we focus on the many items by which knowledge is gained, preserved or extended, and notice that, individually, they do not sustain closure, 'it seems odd to suppose that knowledge itself is closed'?

About at least two of the items he lists, information and indication, Dretske's case is unconvincing: he does not establish that these do not sustain closure on an individual basis. Consider information. Dretske maintains that, normally, when a thermometer registers 32 degrees on a Fahrenheit scale, it carries the information that *temp*: the temperature is 32 degrees Fahrenheit. (A thermometer that registers 0 degrees Celsius also carries the information that *temp*, though via a different scale.) Now, *temp* entails that the following proposition, *broken*, is false:

The temperature is 100 degrees Fahrenheit and the thermometer is broken.

Yet Dretske denies that the thermometer carries the information that *not-broken* holds. This is probably because Dretske understands information in terms of conclusive reasons, so that R carries the information that *p* is so only if, were *p* false, R would not hold. However, there is no reason to account for information in terms of conclusive reasons, which does not sustain closure, rather than safe indication, which does. On the latter option, we say, roughly, that R carries the information that *p* when and only when, were R to hold, *p* would also hold [Luper 1987], and we conclude that if a thermometer carries the information that *temp* holds, it also carries the information that *not-broken* holds. (It is true that a thermometer does not make it obvious that *not-broken* holds, but on both accounts, conclusive reasons and safe indication, something can carry information without making the information readily apparent: a book-length document might carry a bit of information hidden deep in its pages, and intergalactic rays can carry information about distant phenomena we cannot yet fathom.)

Not only does Dretske give us no reason to prefer his account of information, he inadvertently supplies (in 1981) a reason to prefer the safe indication account. According to Dretske, an account of information is adequate only if it upholds the following condition: a signal carries the information that A is F for a person S if and only if the conditional probability that A is F given S's background knowledge is 1 (but less than 1 given S's background knowledge alone—I set this parenthetical qualification aside). Dretske's claims about information are at odds with this condition (Jager 2004). Assume that, given S's (non-Gettierized) circumstances, barn-type percepts constitute a conclusive reason for S to believe *barn*. On Dretske's analysis, *barn*'s conditional

probability is 1, and S's percepts carry the information *barn*. Now, if the conditional probability of a proposition p given r is 1, and p entails q , then the conditional probability of q given r is 1. Since *barn* entails *not-fake*, *not-fake*'s conditional probability given S's percepts is 1, and, by Dretske's adequacy condition, it had better turn out that S's barn-type percepts carry the information *not-fake*. Yet they do not, if we insist that they must be conclusive reasons for *not-fake*. The safe indication account of information conforms to the adequacy condition, if we put aside the parenthetical qualification, which we will want to do, so as to ensure that knowing p on the basis of R coincides with R's carrying the information that p .

Now consider indication. As we have already seen, on the safe indication account, indication *is* individually closed: R indicates that p is true only if R indicates that p 's consequences are true. Yet Dretske claims that tracks in the snow indicate that a deer is in the woods and not that (say) all the deer in the woods have not been replaced with simulacra. Presumably Dretske says this because he thinks that indication should be understood in terms of conclusive reasons, thus: R's indicating p amounts to R's not holding if p were false. Yet analyzing indication in terms of safety is at least as plausible as analyzing indication as Dretske would.

But suppose Dretske were correct in saying that all or some of his items fail to sustain closure on an individual basis. Suppose we follow him in rejecting the following principles:

- If S has *received testimony* that p and (believes q because) p entails q , then S has received testimony that q .
- If S has *proven* p and (believes q because) p entails q , then S has proven q .

-- If S *remembers* p and (believes q because) p entails q , then S remembers q .

As Dretske knows, the falsity of these principles would not constitute good grounds for rejecting K. In the matter of knowledge closure, the issue is not whether these principles are true, but rather whether knowledge by testimony, proof or memory *positions* us to know the consequences of the things we know. What is relevant is whether the following principle, suitably qualified, is true:

T: If S knows p via testimony, or proof, or memory, and believes q because p entails q , then S knows q .

This principle is not only intuitive, it can be supported if knowledge is safe indication. If, for example, your testimony is, under the circumstances, a safe indicator that p is so, then it, with or without the deduction of q from p , is also a safe indicator that q is so, where q is any consequence of p .

Certainly, T will need qualification, and some necessary modifications may indicate the need to qualify K. But it is worth reminding ourselves that Dretske's central project is to explain why our knowledge of commonsense claims does not position us to know that skeptical scenarios do not hold; he means to reject any closure principle, however qualified, that implies that we may arrive at antiskeptical knowledge given our knowledge of commonsense claims. To that end, it is not enough for Dretske to show that K (or T) must be qualified, for even on various qualified versions of K we may arrive at antiskeptical knowledge.

The Case Based on Perception

According to Dretske [2003: 108], ‘perception teaches us’ that the following principles (with or without the parenthetical qualifications) are false (Dretske himself formulates the former, but he clearly means to reject the latter as well):

P: If S perceives p , and (S believes q because) S knows that p entails q , then S perceives q .

PK: If S knows p via perception, and (S believes q because) S knows that p entails q , then S knows q via perception.

Perception also teaches us to reject K, the closure of ‘factual’ (empirical) knowledge [2003: 112]. In its broadest outlines this case against P, PK, and K has two main steps.

First. Dretske thinks he can show that we do not perceive or know the truth of certain sorts of propositions he calls ‘limiting propositions’ [2003: 112]. This set of claims is not clearly delineated, but many of Dretske’s ‘limiting propositions’ refer to situations with a feature we might call *elusiveness*. A situation sk is elusive to me when the following is true: were sk not to hold, I would still have the experiences I have now. A familiar example is that I am not a brain in a vat being fed my present sensory input. Another example arises when I look at a cookie jar in my kitchen. The example is *not-rayed*: I am not standing in a jar-free kitchen undergoing a cookie jar hallucination caused by a ray from a passing alien spaceship.

Second. Dretske also thinks he can show that we *do* perceive and know thereby all sorts of ordinary contingent claims that entail ‘limiting propositions’. An example he gives is *jar*—I am confronting a jar containing cookies—which entails *not-rayed*.

If we do not perceive or know elusive claims, but we do perceive and know ordinary claims that entail elusive claims, we have to reject P, PK and K. Recapping:

- I. We perceive and know that various ordinary contingent claims (such as *jar*) hold.
- II. We neither perceive nor know that elusive claims (such as *not-rayed*) hold.
- III. Elusive claims are entailed by some of the ordinary claims we perceive and by some of the ordinary claims we know.
- IV. So P, PK, and K are false.

So much for the main structure of Dretske's case; now let us consider the details.

To support I, Dretske relies on intuitions about what we perceive and know through perception. He thinks it is intuitively plausible to say that we perceive, and know thereby, the truth of ordinary contingent claims.

Dretske also defends II on intuitive grounds. He thinks our intuitions support both of the following claims:

- IIA. We do not perceive that elusive claims hold.
- IIB. We do not know via perception that elusive claims hold.

Furthermore, in underwriting IIB, our intuitions suggest something stronger, for IIB, he thinks, supports the following claim:

- IIC. We do not know that elusive claims hold.

So far, Dretske's case for I and II is based on intuitions and the suggestion that IIB supports IIC. But I and II can be further supported if we can find an analysis of perception and perceptual knowledge that underwrites the intuitive plausibility of IIA and IIB. Dretske thinks the conclusive reasons account is precisely what is needed.

According to Dretske, in *S's perceiving that p*, and in *S's knowing p via perception*, the key element is, roughly, *S's* possessing a set of experiences *E* that constitute a conclusive reason for *p* (Dretske adds refinements we can ignore for the sake

of simplicity). On this account of perception and perceptual knowledge, S perceives, and knows via perception, that *jar* holds: if *jar* were false, some situation that is every bit as mundane as *jar* would hold, such as the situation in which S is confronting an ordinary jar containing no cookies, and in that case S would not have the experiences E associated with *jar*. Dretske's account also allows us to say that S's perceiving, and knowing perceptually, that *jar* holds does not require that S have experiences S would lack if *rayed* held. For even if *jar* were false *rayed* would not hold; some mundane situation would hold, and S would lack E. Finally, Dretske's account of perception and perceptual knowledge backs what we said about elusive possibilities: on his view, perceiving, and knowing via perception, that *not-rayed* holds, while I examine the jar of cookies in front of me, requires that I meet the following condition: if *rayed* held, I would not have my cookie-jar-in-front-of-me experiences. But I would.

We now have before us Dretske's case against P, PK and K, which rests finally on his arguments for IIA, IIB, and IIC. Let us examine each of these three arguments in turn.

IIA. Dretske says we do not perceive that elusive possibilities hold. His explanation rests on the acceptability of the conclusive reasons account of perception, according to which perceiving *p* requires having experiences E that constitute a conclusive reason for *p*. His thought is that our experiences will not be conclusive reasons for believing elusive claims. However, Dretske's account is no better than a safe indication account of perception, according to which perceiving *p* is roughly having experiences that safely indicate that *p* is true. On this latter theory, in principle there is nothing to stop us from perceiving that elusive claims hold. For example, in non-

Gettieresque circumstances, my jar-before-me experiences safely indicate not just the truth of *jar*, but also of *not-rayed*.

Nonetheless, there are reasons to doubt that we perceive the truth of elusive claims. We might accept *naïve perceptual representationalism*, which says that we perceive *p* only if we form a perceptual representation of the way things are if *p* holds. On this view, we might well doubt that we perceive the truth of elusive claims on the grounds that we rarely if ever form perceptual representations of elusive situations. It is plausible to say one can form (say) a visual representation of the fact that *jar* holds: one comes to have a visual impression of a jar, together with an impression of one's body and its spatial relationship to the jar, and so forth. But forming a visual representation that *not-rayed* holds, or that any other elusive possibility holds, is far from straightforward.

However, it seems possible to make sense of perceiving that elusive claims hold. We have only to adapt Dretske's own distinction between primary (or direct) perception and secondary (or indirect) perception [1969: 78-88, 153-163]. Imagine a physicist looking at a tract of bubbles in a cloud chamber. She cannot see electrons, so she cannot directly see that electrons are moving through the medium within the instrument. But she sees the bubbles, and that a series of bubbles is appearing in the instrument; furthermore, it is reasonable to think that she can see that the electrons are moving *by* seeing that the bubbles are forming in a specific sort of way. She sees the former indirectly by seeing the latter directly. Let us say that S has *primary* perception that something A has feature F roughly when:

1. S perceives A
2. S's belief that A is F is caused by S's experiences E

3. E safely indicate that A is F.

Let us further stipulate that S has *secondary* perception that something B (distinct from A) has feature G roughly when:

4. S has primary perception that A is F
5. The fact that A is F is a safe indicator that B is G
6. S's belief that B is G is caused by S's perceiving primarily that A is F.

Accordingly, a physicist might (1) perceive bubbles; (2) have experiences E' that bring about the belief that the bubbles have certain features F', under conditions given which (3) E' safely indicate that the bubbles have F'. Satisfying (1)-(3), she satisfies (4): she has primary knowledge that the bubbles have F'. The fact that the bubbles have F' is a safe indicator that electrons are moving in the cloud chamber, as (5) demands. So, assuming that her belief about the electrons is caused by her primary perception that the bubbles have F', as (6) requires, she perceives secondarily that electrons are moving in the chamber.

There is no reason in principle why secondary perception should not grade off into tertiary perception, and perhaps even higher levels. Thus, for example, the physicist might see that electrons are moving through her cloud chamber even though she is watching from an adjoining room using a camera trained on her device. Hence we will want to say that S perceives *q* on level *n* if and only if either $n = 1$ and S primarily perceives *p*, or $n > 1$ and the following conditions are met:

7. S perceives *p* on level $n-1$
8. The fact *p* is a safe indicator that *q*
9. S's belief *q* is caused by S's perceiving *p* on level $n-1$.

The upshot is clear: indirectly, we can perceive that some elusive claims hold. Having primary perception of the fact that *jar* holds positions us to have secondary perception of the fact that *rayed* does not hold. What is more, a fairly straightforward closure principle applies to perception, namely:

PS: If S perceives *p*, and S is caused to believe *q* by perceiving *p*, and *p* entails *q*, then S perceives *q*.

We have shown that, by replacing Dretske's conclusive reasons account of perception with the safe indication account, we can make good sense of perceiving the truth of elusive claims. But before we accept either account, let us register one reservation: it may be that *both* accounts of perception are too strong. Recall the papier-mâché barn example. I believe there is a barn in front of me because, looking right at it, I see it is there, yet, unbeknownst to me, the neighborhood is full of papier-mâché barns that look just like the real thing. In Goldman's scenario, it is clear that I fail to *know* a barn is there, but it seems just as clear that I *see* that a barn is in front of me. Yet on the conclusive reasons and safe indication accounts, I *fail* to see that the barn is there, for, given the proximity of the fakes, I might I have had my barn-in-front-of-me experiences without a barn being present.

Our intuitions about Goldman's example suggest that it involves genuine perception, which, in turn, suggests rejecting our account (and Dretske's as well). However, there is another option. It also seems intuitively plausible to say that perceiving (or at least seeing) *p* entails knowing *p*, as several theorists, including Dretske [2003: 108, note 3; 1969: 124; and Williamson 2002], have suggested. If the entailment thesis is true—if perceiving *p* entails knowing *p*—then there is perception failure in

Goldman's example after all (since there is clearly knowledge failure). If we embrace our intuitions about the entailment thesis, we must reject our intuitions about perception failure in the Goldman example. Rejecting the latter intuitions seems preferable, inasmuch as embracing the former allows us to unify our accounts of perception and perceptual knowledge: we can explain both in terms of safe indication. For the safe indication account of perception is entirely consistent with the entailment thesis.

Thus, we are led back to the safe indication account of perception, and on that account it is possible to perceive that elusive claims hold, contrary to IIA.

IIB. We turn to Dretske's view that we do not know via perception that elusive possibilities hold. Now, it is critical to notice that 'we know p via perception' can mean more than one thing. Consider the following notions:

S has *noninferential perceptual* knowledge that p if and only if, in knowing p , S relies strictly on perception (perception that is not supplemented by any form of inference, deductive or nondeductive).

S has *inferential perceptual* knowledge that p if and only if, in knowing p , S infers p from something (or some things) of which S has noninferential perceptual knowledge.

S has *perceptually based* knowledge that p if and only if S has either noninferential or inferential perceptual knowledge that p .

When we say a person S knows p via perception, we might mean that S has noninferential perceptual knowledge that p . But we might also mean that S has perceptually *based* knowledge that p . Therefore, IIB might mean either of the following:

IIB1 We lack noninferential perceptual knowledge of elusive claims.

IIB2 We lack perceptually based knowledge of elusive claims.

To evaluate Dretske's position, then, we will need to ask whether he can support each claim the way he defends IIB. That is, we must ask whether he can defend IIB1 and IIB2 by noting that each is intuitively plausible, and suggesting that the intuitiveness of each is best explained in terms of the conclusive reasons account of perceptual knowledge. Let us start with IIB1.

IIB1 *is* intuitively plausible. But the explanation has nothing to do with the conclusive reasons account of perceptual knowledge. Strict (wholly noninferential) perceptual knowledge presumably has a perception component and a knowledge component. That is, S knows *p* strictly via perception when and only when the following conditions hold:

- (a) S perceives *p* (the perception component)
- (b) S knows *p* (the knowledge component)
- (c) S's perceiving *p* is what positions S to know *p* (i. e., S's meeting condition (a) is what positions S to meet condition (b)).

To explain why it seems difficult to achieve noninferential perceptual knowledge of elusive claims we have only to point out two things. First, knowing *p* strictly via perception entails *perceiving* that *p* holds. Second, as noted earlier, it is tempting to accept naïve perceptual representationalism, according to which perceiving *p* entails forming a perceptual representation of the fact *p*. Together these two points suggest that we do not perceive elusive claims, for it is hard to accept that we form perceptual representations of elusive situations.

So much for why IIB1 strikes us as intuitively plausible. Now recall that we have rejected naïve representationalism in favor of the safe indication account of perception, and on that theory nothing stands in the way of strict perceptual knowledge of elusive claims. Appearances to the contrary, then, IIB1 is false.

Even if we analyze the knowledge component of noninferential perceptual knowledge in terms of safe indication rather than conclusive reasons, we still must reject closure as applied to knowledge strictly via perception. That is, the following principle (with or without the parenthetical qualification) is false:

PKS: If S knows p strictly via perception, and (S believes q because) S knows that q is entailed by p , then S knows q strictly via perception.

PKS is *trivially* false. It is false because no inferential belief *counts* as ‘noninferential perceptual knowledge.’

As for IIB2: I suggest that Dretske’s explanation of its plausibility fails. For an alternative explanation is at least as good: IIB2 is false, but for reasons that are easily overlooked. Suppose we accept the safe indication account of perceptual knowledge. Then we know many different elusive claims (such as *not-rayed*) indirectly, by inference from other claims we know strictly on the basis of perception, such as *jar*. IIB2 is then false. Yet it is easy to see how we might be led to accept IIB2. We might find ourselves confused in one of two ways. First. On the safe indication view, we know an elusive claim via perception only if we satisfy condition SI. It is easy to confuse SI, which we satisfy, with CR, which we fail to satisfy. Second. As we said earlier, naïve representationalism suggests we do not know that elusive claims hold *via noninferential perception*. If this kind of knowledge failure were actual, it would be easily confused

with a second kind of knowledge failure: our inability to know elusive claims even if we supplement strict perceptual knowledge with inference.

Finally, let us add that the safe indication account of perceptual knowledge has an important advantage over the conclusive reasons account: the former, unlike the latter, positions us to accept the closure of perceptually based knowledge. It enables us to endorse something like the following principle:

PKI: If S has perceptually based knowledge that p (i.e., S has either noninferential or inferential perceptual knowledge that p), and S believes q because S knows that q is entailed by p , then S has perceptually based knowledge that q .

IIC. Dretske thinks that *IIC*, the assertion that elusive claims are not known, is supported by *IIB*, the view that we cannot know elusive claims via perception. However, *IIB*, we have said, can mean *IIB1* (we lack noninferential perceptual knowledge of elusive claims) or *IIB2* (we lack perceptually based knowledge of elusive claims). We must take this complication into account in interpreting Dretske. Does he mean that *IIB1* supports *IIC* or does he mean that *IIB2* supports *IIC*?

Now, *IIB2* really would support *IIC*. For *IIB2* implies not just that we do not know elusive claims strictly on the basis of perception, but also that we do not know elusive claims even if we help ourselves to what we can infer from things we know strictly on the basis of perception. It is very hard to see how else we can know elusive claims. However, we have already shown that *IIB2* is false. So it is no good to defend *IIC* on the basis of *IIB2*.

In any case, in saying that IIC is supported by IIB, it is fairly clear that Dretske did not mean that IIB2 supports IIC. Like many philosophers, Dretske adopts a definition, more or less by stipulation, by which ‘perceptual knowledge’ is equated with what we have called ‘noninferential perceptual knowledge’ [1969: 159]. In asserting that IIB supports IIC, then, he clearly means to say that IIB1 supports IIC. Unfortunately, it is not at all clear that IIB1 does support IIC. Of course, the main difficulty is that IIB1 is false, as we argued earlier. But suppose it were true: suppose we did not perceive that elusive claims hold. It would follow that we did not know that they held *by* perceiving that they held. But why should this lead us to rule out knowing elusive claims by some other means? Why, in particular, should we rule out knowing elusive claims by deducing them from other claims we know strictly via perception? Unless we conflate IIB1 with IIB2, it is hard to see how IIB1 provides any support whatever for IIC, much less grounds for rejecting K.

In closing, let me briefly discuss one other argument for IIC that Dretske hints at. At one point he says it would be bizarre to think we can know a ‘limiting proposition’ by deducing it from one of the ordinary claims we know, but he does not say why [2003: 112]. However, I suspect he is completing a thought he began much earlier in his essay:

there are always things my knowledge depends on, facts without which my beliefs would be false, that I cannot justify. So the knowledge, if I have it, must be the product of things I need not know or be justified in believing. [2003: 106]

Among the things that help produce knowledge is the truth of various ‘limiting propositions’:

One doesn't have to be justified in thinking that the world was not created in the way Russell imagines in order to remember—hence, know—facts about the past. It is the fact that Russell's hypothesis is false, not one's justification for thinking it false, that enables one to remember what one had for breakfast. [2003: 106]

Dretske thinks that, to know an ordinary claim p , we must meet certain conditions, among them the requirement that the 'limiting propositions' that entail p are true. But we meet this requirement, and all other requirements for knowing p , without being positioned to know, or warranted in believing, the 'limiting propositions.' In a sense, he seems to think, in knowing the ordinary claim p , we merely *presuppose* the truth of the 'limiting propositions.' And under such conditions, we cannot come to know the 'limiting propositions' by inferring them from the ordinary claim p .

But why not? Is it that we lack justification ('justification' in the traditional sense) for p ? That seems unlikely, since Dretske is a thoroughgoing externalist who thinks that knowledge does not require justification.

Perhaps he thinks we cannot come to know any claim q by deducing it from another claim p which we know if our knowing p depends on q 's truth. Maybe this pattern of inference (call it *pseudocircular* reasoning) strikes him as objectionable. But if so, Dretske owes us an explanation. And in explaining himself he will be unable to fall back on his view that knowledge requires conclusive reasons. For that view does not rule out knowledge via pseudocircular reasoning. It does not even rule out knowledge of 'limiting propositions' via pseudocircular reasoning. At best, it rules out knowledge of *elusive* claims via pseudocircular reasoning, simply because we lack conclusive reasons for believing elusive claims whether our reasoning is pseudocircular or not. (Even this is

an overstatement, for nothing in the nature of conclusive reasons stops us from casting things we know *in the role of* conclusive reasons which position us to know further things, which would allow us to know elusive claims through applications of K. For example, we can start with our knowledge that we are not in vats, and conclude that we know we are not brains in vats. Our visual percepts give us conclusive reason to believe we are not in vats, and our not being in vats is conclusive reason to believe we are not brains in vats.) By way of illustration, consider two examples. In both we have a conclusive reason for believing something, and we arrive at that belief via pseudocircular reasoning. The second is a ‘limiting proposition,’ while the first is not, as far as I can tell (again: Dretske’s term is not precise). Neither is an elusive claim.

First example: I can have a conclusive reason for believing I am alive after inferring it from my belief (whose truth I know) that I am typing. If I were not alive, I would not have the experiences that led me to believe I am typing, nor the belief that led me to conclude that I am alive. Yet I must be alive if I am to know I am typing.

Second example: Dretske thinks *There are physical objects* is a ‘limiting proposition.’ However, I can have a conclusive reason for believing that there are physical objects, even if I infer it from my belief (whose truth I know) that there is a jar of cookies in front of me. If there were no physical objects, I would not have my belief, my jar-of-cookies-in-front-of-me experiences, or any other experiences, for that matter, since I would not exist.

In view of these examples, one could easily cite the conclusive reasons view in support of the possibility of knowledge via pseudocircular reasoning. If he really means to reject the latter anyway, Dretske needs to tell us why.

Let us summarize the discussion. First, we have defended the safe indication account of perception as against Dretske's conclusive reasons account. On our approach we can perceive, indirectly, that elusive claims hold, and perception is closed in the following sense:

PS: If S perceives p , and S is caused to believe q by perceiving p , and p entails q , then S perceives q .

Second, we rejected Dretske's conclusive reasons account of perceptual *knowledge* in favor of the safe indication account. If, like Dretske, we go on to define 'perceptual knowledge' to mean *noninferential* perceptual knowledge, we still allow for the perception of elusive claims, but we must deny that perceptual knowledge is closed under (known) entailment. That is, the following version of PK is trivially false:

PKS: If S knows p noninferentially via perception, and S believes q because S knows that q is entailed by p , then S knows q noninferentially via perception.

However, if we understand 'perceptual knowledge' to mean perceptually *based* knowledge, which embraces things we know by inference from claims known strictly via perception, then we know elusive claims via perception, and perceptually based knowledge is closed under entailment:

PKI: If S has perceptually based knowledge that p , and S believes q because S knows that q is entailed by p , then S has perceptually based knowledge that q .

Finally, if we adopt the safe indication account of factual knowledge, K is true.

Apparently, nothing about perception teaches us otherwise.

References

- BonJour, Lawrence 1987. Nozick, Externalism, and Skepticism, in Luper 1987a, 297-314.
- Cohen, Stewart 1988. How to Be a Fallibilist, *Philosophical Perspectives* 2: 91-123.
- Dretske, Fred 1969. *Seeing and Knowing*, Chicago, IL: University of Chicago Press.
- 1970. Epistemic Operators, *Journal of Philosophy* 67: 1007-23.
- 1971. Conclusive Reasons, *Australasian Journal of Philosophy* 49: 1-22.
- 1981. *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.
- 2003. Skepticism: What Perception Teaches, in Luper 2003a, 105-119.
- Feldman, Fred 1999. Contextualism and Skepticism, *Philosophical Perspectives* 13, 91-114.
- Fumerton, Richard 1987. Nozick's Epistemology, in Luper 1987a, 163-182.
- Goldman, Alvin 1976. Discrimination and Perceptual Knowledge, *Journal of Philosophy* 73: 771-791.
- Jaeger, Christoph 2004. Skepticism, Information, and Closure: Dretske's Theory of Knowledge, *Erkenntnis* 61: 187-201.
- Lewis, David 1973. *Counterfactuals*, Cambridge: Harvard University Press.
- 1996. Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-67.
- Luper(-Foy), Steven, 1984. The Epistemic Predicament, *Australasian Journal of Philosophy* 62: 26-48.

- (ed) 1987a. *The Possibility of Knowledge*, Totowa: Rowman & Littlefield.
- 1987b. The Possibility of Skepticism, in S. Luper 1987a, 219-241.
- (ed) 2003a. *The Sceptics*, Aldershot: Ashgate Publishing Limited.
- 2003b. Indiscernability Skepticism, in Luper 2003a, 183-203.
- 2006. Restorative Rigging and the Safe Indication Account, *Synthese*, forthcoming.
- Nozick, Robert 1981. *Philosophical Explanations*, Harvard: Belknap Press.
- Sosa, Ernest 1999. How to Defeat Opposition to Moore, *Philosophical Perspectives* 13:
141-53.
- 2003. Neither Contextualism nor Skepticism, in Luper 2003a, 165-183.
- Stalnaker, Robert, 1968. A Theory of Conditionals, in *Studies in Logical Theory*, ed. N.
Rescher, Oxford: Basil Blackwell, 1968.
- Stine, G. C. 1976. Skepticism, Relevant Alternatives, and Deductive Closure,
Philosophical Studies 29: 249-61.
- Williamson, Timothy 2000. *Knowledge and Its Limits*, Oxford: Oxford University Press.