

CHRISTOPH JÄGER

SKEPTICISM, INFORMATION, AND CLOSURE: DRETSKE'S
THEORY OF KNOWLEDGE

ABSTRACT. According to Fred Dretske's externalist theory of knowledge a subject knows that p if and only if she believes that p and this belief is caused or causally sustained by the information that p . Another famous feature of Dretske's epistemology is his denial that knowledge is closed under known logical entailment. I argue that, given Dretske's construal of information, he is in fact committed to the view that both information and knowledge *are* closed under known entailment. This has far-reaching consequences. For if it is true that, as Dretske also believes, accepting closure leads to skepticism, he must either embrace skepticism or abandon his information theory of knowledge. The latter alternative would seem to be preferable. But taking this route would deprive one of the most powerfully developed externalist epistemologies of its foundation.

1. EXTERNALISM, INFORMATION, AND THE KK THESIS

Externalist theories of knowledge typically claim to provide an effective response to skepticism. One of the most influential externalist epistemologies has been developed by Fred Dretske. The core idea of his theory is that:

- (K) An epistemic subject knows of some object (or source of information) s that it has the property F if and only if the subject believes s to be F and this belief is caused (or causally sustained) by the information that s is F .

Dretske has spelled out this view in great detail. According to the account of information laid out in his book *Knowledge and the Flow of Information*, a signal carries, relative to a given subject, the information that s is F if and only if the conditional probability of s 's being F , given the signal and the subject's background knowledge, is 1, but, given only the subject's background knowledge, is less than 1.¹

The central feature of this account, which is primarily designed for perceptual knowledge, is that the information relation is veridical:

The appearance of the (shelves in the) refrigerator can only inform you that the refrigerator is empty if it is true that it is empty. The gas gauge cannot carry the information that the tank is half full unless it is half full. Moreover, due to the theory's reference to background knowledge, it may happen that a signal carries some piece of information for me but not for you: if you saw Jenny walking up the stairs to the front door, but I did not, her signature pattern of knocks at the door may inform me that she has arrived, but it does not carry that information for you, since you already know that she is there.

How does skepticism fare in this theory? As Dretske observes, it seems undeniable that skeptical worries undermine our knowledge *that we know* (Dretske, 1981a, p. 128). A crucial feature of his externalism, however, is that – contrary to what Roderick Chisholm, Carl Ginet, and others have argued – knowing that p does not require that the subject possess the higher-order knowledge that she knows that p . The KK thesis, the claim that in order for K to know that p , K must also know that she knows this, is rejected in Dretske's account. Put in terms of information: if K comes to know something by receiving some piece of perceptual information, she need not receive the higher order information *that it is (genuine) information* she receives. I can see that there is beer in the fridge, even if I don't possess any higher-order information about whether – to use Chisholm's idiom – my currently being appeared to beerly is reliable.

The fact that the KK thesis immediately leads into skeptical problems is a good reason for abandoning it.² Yet, can Dretske's account also handle skeptical worries from other directions? Dretske himself was the first philosopher to give clear articulation to the fact that another road to skepticism is the view that knowledge is closed under known logical entailment (see especially Dretske, 1970, 1981b). In a first approximation, the relevant principle of epistemic closure can be sketched as follows:

(PEC) If K knows that p and knows that p entails q ,
 then K knows that q .³

Now, Dretske famously rejects closure as well. His arguments for this conclusion are controversial, but this is not what I want to discuss in this paper.⁴ My question is whether this part of Dretske's epistemology is compatible with his information-theoretic account of knowledge. Is his denial of closure even *consistent* with his idea that generating perceptual knowledge consists essentially in receiving information, given that the information relation is construed in the

probabilistic way sketched above? I argue that it is not. There are deep perplexities regarding the compatibility of these two cornerstones of Dretske's epistemology, i.e. his information-based, probabilistic account of knowledge, and his denial of closure.

2. CONTEXTUALISM AND CLOSURE

Dretske's *Knowledge and the Flow of Information* is a modern classic in externalist epistemology, and his seminal work on closure dates back to the 1970s (see, for example, Dretske, 1970–1972). In his paper “Externalism and Modest Contextualism” he takes some fresh looks at skepticism, closure, and knowledge by comparing his externalism with epistemic contextualism. I shall take my cue from these recent reflections of Dretske's.

Contextualist theories of knowledge, as championed in recent decades by Stewart Cohen, Keith DeRose, and David Lewis,⁵ differ considerably in their details. Yet they share the idea that the truth values of knowledge attributions can vary with the epistemic standards of the attributer's context. The hallmark of this “attributorial” or, as Dretske calls it, “radical contextualism” is its contention that it can properly deal with skeptical puzzles, while at the same time retaining the view that knowledge is closed under known logical entailment. Let K stand for some normal human adult, e for some suitable empirical proposition (such as “This is a zebra,” “Here is a hand,” etc.), and not- h for the negation of some (local or global) skeptical alternative to e . The skeptical puzzle arises from the fact that we would neither want to deny that, in a suitable situation, K knows e , nor that K knows that e entails not- h . If we accept that knowledge is closed under known entailment, it follows that K also knows not- h . But skeptical arguments seem to show that the latter is impossible: we don't *know*, it seems, that skeptical hypotheses are false. The problem is that each of these claims, including PEC (or some embellished version of this principle), appears to be true; but taken together, they are inconsistent. Which one has to give?

Very roughly, the contextualist replies that in contexts in which skeptical hypotheses are salient, the standards for knowledge ascriptions differ from the standards for knowledge ascriptions in everyday contexts. For example, in ordinary contexts the reasons we have for taking appearances to be trustworthy suffice to license the judgement that our perceptual beliefs constitute knowledge. In

skeptical contexts, however, standards have been raised, and relative to these elevated standards knowledge attributions become false. Closure is said to hold in both types of context. Hence, says the contextualist, in ordinary contexts we do know after all that skeptical hypotheses are false, whereas in skeptical contexts we don't know that ordinary empirical propositions are true.⁶

These claims are Dretske's main target. According to conversational contextualism, we can apparently rob someone of his knowledge simply by raising the question whether he is in an epistemic position in which he can exclude skeptical alternatives. Moreover it is possible in this view for a subject, just by contemplating some skeptical possibility and thereby making it a citizen of the realm of relevant alternatives to a given proposition, to create a context in which it is true to say that she *never knew* that proposition. This, Dretske urges, is a most unfortunate result. For if this is an answer to skepticism, it is one that is only available to those who never think about skepticism. Yet, if skepticism is false, shouldn't it be false in the philosophy classroom as well as the grocery store?

I think that here Dretske does put his finger on an important problem of contextualism. I want to ask, however, whether his information theory of knowledge is really any better off. Dretske recommends rejecting closure in order to avoid these unwelcome results. "Abandon closure", he writes (2004, p. 182), "and these unpalatable results vanish." This is right. Rejecting closure is a theoretical option that, if adopted, would solve the kind of skeptical problem outlined above without forcing us to embrace a contextualist position. But giving up closure is a very high price to pay. My question is whether it is a price that Dretske himself is able to pay without having to sell his probabilistic theory of information. I shall now argue that it is not.

Dretske offers two main lines of reasoning against closure. One is a family of indirect arguments to the effect that abandoning closure is unavoidable if we want to put skeptics in their place.⁷ It is doubtful, however, whether these arguments can ultimately cut any ice. To begin with, it is immensely plausible to assume that a body of knowledge can expand by deductive reasoning from known premises. Giving up closure would mean giving up this integral part of our ordinary understanding of knowledge. Secondly, Dretske's *modus ponens* is the skeptic's *modus tollens*: the skeptic will adhere to closure, insist on his claim that we don't know that skeptical hypotheses are false, and thus conclude that we lack knowledge of those homely truths we ordinarily take ourselves to know.⁸

Yet in “Externalism and Modest Contextualism” Dretske also suggests another interesting argument. “I do not, myself, accept closure”, he writes:

I think there are some things we know to be implied by what we know that we do not – perhaps cannot – know to be true. My reason for rejecting closure, however, is not my externalism. It is that I think I have some reasonably clear idea of what kind of evidential relation is required for knowledge (a “conclusive reason” or “information”) and this relation is itself not closed under known implication. We can have, in the relevant sense of “conclusive,” conclusive reasons to believe p is true – we can, that is, get information that p is true – without having conclusive reasons to believe, without having information, that q is true even when we know that p implies q . (Dretske, 2004, p. 176f)

The argument in this passage seems to be this: (perceptual) knowledge requires receiving appropriate information from some signal. But the relation of receiving such information is not closed under known implication; hence (perceptual) knowledge is not closed under known implication. In a forthcoming paper on closure that Dretske has kindly given me the chance to read, he explicitly argues that no signal can carry the information that skeptical hypotheses are false. “There is nothing in the world,” he says, “[. . .] that indicates that there is a material world” (Dretske, forthcoming). Similar remarks can already be found in *Knowledge and the Flow of Information*:

No signal can rule out *all* possibilities if possibilities are identified with what is consistently imaginable. No signal, for instance, can eliminate the possibility that it was generated, not by the normal means, but by some freak cosmic accident, by a deceptive demon, or by supernatural intervention. (Dretske, 1981a, p. 130)

I agree. I wish however to argue, first, that Dretske’s account is – at least with respect to empirical propositions and their antiskeptical consequences – committed to the view that the relation of a signal’s carrying information *is* closed under known entailment. Second, I will show that, given this fact and given that we should work with a refined version of the closure principle, Dretske’s theory is also committed to the closure of knowledge. Contrary to what he believes, his information-theoretic externalism implies, at least for ordinary empirical propositions and their anti-skeptical consequences, that knowledge *is* closed under known entailment.

3. INFORMATION AND CLOSURE

Let us first examine whether *information* is closed under known entailment, i.e. whether it holds that:

- (PIC) If r carries, relative to the subject K , the information that p , and K knows that p entails q , then r carries, relative to K , the information that q .

Recall, first, that in Dretske's account a signal r carrying the information that p (relative to a given subject) implies that the probability of p , given r (and the subject's background knowledge k) is 1. And if K knows that that p entails q , p does entail q . Hence if in addition the conditional probability of p , given some signal r (and k), is 1, the conditional probability of q , given r (and given k), *must* also be 1.

This does not yet amount to an effective counterargument. For as we have seen, Dretske's definition of a signal's carrying information is slightly more complicated: it includes the condition that the prior probability of the proposition in question is less than one. In Dretske's theory, it holds that:

- (I) A signal r carries the information that s is F (relative to a given subject) if and only if $P(F(s)/r \ \& \ k) = 1 \ \& \ P(F(s)/k) < 1$.

Nevertheless, I wish to argue that information in the full sense of this account, at least for the kinds of propositions at issue, is closed, not only under known entailment, but also under entailment *simpliciter*. The propositions at issue are (potential) contents of perceptual beliefs, and their anti-skeptical consequences. Let e stand again for some empirical proposition of an appropriate kind ("This is a zebra," "Here is a hand"), and let not- h represent an appropriate proposition which negates some (local or global) skeptical hypothesis that is incompatible with e ("This is not a cleverly painted mule," "There is a material world"). What needs to be shown, then, is that the following instantiation of the above closure principle regarding information (PIC) is *true*:

- (PIC*) If $P(e/r \ \& \ k) = 1 \ \& \ P(e/k) < 1$, and Knows
 $(K, e \Rightarrow \text{not-}h)$, then $P(\text{not-}h/r \ \& \ k) = 1 \ \& \ P(\text{not-}h/k) < 1$.

In prose: if r carries, relative to K 's background knowledge k , the information e , and K knows that e entails not- h , then r *does* also, relative to k , carry the information that not- h .⁹ I have already indicated the first step of the argument for this claim: if we assume that a signal r carries the information e (which implies that the conditional probability of e , given the signal and k , is 1), and that K knows that e entails not- h (and hence that e does entail not- h), we must conclude

that the conditional probability of not- h , given r and k , is also 1. Put in the lingo of the probability calculus:

- (1) $P(e/r \ \& \ k) = 1 \ \& \ P(e/k) < 1$,
- (2) $e \Rightarrow \text{not-}h$.

Therefore:

- (3) $P(\text{not-}h/r \ \& \ k) = 1$.

But can it also be shown that the antecedent probability of not- h is less than 1, i.e. that $P(\text{not-}h/k) < 1$?

At this point, we might ask whether this second condition is acceptable. Does not this requirement take us too far away from our intuitive understanding of the concept of information? For this requirement has the consequence that, if K already knows that p , it is not possible that there still be any signal for K that carries the information that p . You see Jenny approaching, and a few moments later you also hear her voice. Should we not say that this auditory signal does carry the information for you that she is here, despite the fact that you already know this? Perhaps. At least one may be inclined to say so if the counterfactual holds that, had you not already seen (and thus known) that she is there, hearing her speak to you would have been sufficient for you to generate that knowledge. However, dropping Dretske's condition about the antecedent probability of informational content would amount to a substantive revision of his account. Fortunately such a move is not necessary to bring our point home. For with respect to negations of *skeptical hypotheses*, the requirement that their antecedent probability be less than 1 must, in the framework of Dretske's epistemology, clearly be regarded as fulfilled. The reason is that Dretske takes skepticism seriously. "Skepticism," we even hear him saying, "is true" (Dretske, 2004, p. 174). Now by this, I take it, he does not mean to say that skeptical hypotheses are actually true, but rather that the skeptic is right in claiming that we are not entitled to be certain that they are false. The assumption, in other words, is that (given what we know about the world) the probability of skeptical hypotheses being true is not zero. This means of course that the probability of their negations – i.e. of assumptions of *not* being brains in vats, about the existence of a material world, etc. – is less than 1. So from this part of Dretske's epistemology we get:

- (4) $P(\text{not-}h/k) < 1$.

And hence we finally arrive at:

- (5) $P(\text{not-}h/r \ \& \ k) = 1 \ \& \ P(\text{not-}h/k) < 1$.

(5) says that the signal r carries the information that not- h , and thus that our skeptical hypothesis h , which is incompatible with the empirical proposition e , is false. In summary, then, it has emerged so far that, at least for the kind of propositions here at issue, Dretske's information relation *is* closed under (known) logical entailment: if there is a perceptual signal that carries the information for you that you have hands, and if you know that if you have hands, there is a material world, then that signal also carries the information for you that there is a material world.

What should be concluded from this? Dretske must either take on board the idea that information is closed under (known) entailment, or give up his theory of information. However, each alternative has shattering consequences.

Consider first what would happen if we adopted a Dretske-style theory of knowledge and information, thereby accepting that information is closed under (known) entailment. Someone may want to defend Dretske's overall attack on skepticism along the following lines: "All right," it may be responded, "you have shown that, for perceptual propositions and their anti-skeptical consequences, the *information relation* is closed under known entailment. But what we are really interested in with regard to skeptical puzzles is whether *knowledge* is closed under known entailment. So even if you are right about information, you are barking up the wrong tree. Your argument has little impact on what ultimately is at issue." But this objection would be misguided for several reasons.

First, Dretske seems to think that (1) if information is not closed under known entailment, then neither is knowledge. He further seems to think that (2) the antecedent of this condition is fulfilled (and that therefore knowledge is not closed). But if what has been said so far is right, this argument does not go through since it relies at least on one false premise, namely (2). It is of course another question whether the first premise of this argument is true. If this is false, the abstract possibility remains for Dretske to retain the view that closure fails for knowledge, although (as I have shown) it does not fail for information. It may be noted in this context that Dretske says in the above quotation that we can "*get* information that p is true without *having* information that q is true even if we know that p implies q ." This might be interpreted as referring to a relation other than a signal's merely carrying information. However, could it be a reasonable position within Dretske's epistemological framework to accept that the relation of "carrying information" is closed under known entailment, while denying this for knowledge? I will return to this question

shortly, but first I would like to stress that full-blown knowledge skepticism is in any case not the only problem that we must come to grips with in our epistemological endeavors.

Our discussion thus far already shows that the theory under consideration leads to skepticism about *perceptual information*. And that is bad enough. Consider for instance global skeptical hypotheses. How *could* a perceptual signal relay the information that such hypotheses are false? How could a view of Jenny, nice as it may be, carry the information that there is a material world? If there is one lesson to be drawn from skeptical arguments, it is that perceptual experience certainly cannot teach us that skepticism is false. Put in terms of information: no perceptible signal can carry information about the falsity of skeptical hypotheses. We thus have arrived, to begin with, at a *reductio* of Dretske's theory of information. If we concede, as we should, that perceptual experience is neutral with regard to skeptical scenarios, any account that yields a result to the contrary must have gone astray. I have already indicated that Dretske himself seems to share that view. He speaks of anti-skeptical implications of empirical propositions as "heavyweight implications," and he underlines that:

Ordinary things we come to know by perception always have heavyweight implications that are out of range: we cannot see (hear, smell, or feel) that they are true. I can see that there are cookies in the jar, but I cannot see that there is an external world. [...] This is true of all indicators, all sources of information. That is why there is nothing in the world – either mental or material – that indicates that there is a material world. Nothing in the present that indicates there is a past. (Dretske, forthcoming)

Exactly. But if what I have said is on target, Dretske's information theory of knowledge commits him to precisely these consequences which he himself declares untenable. The problem is that we must either enter the den of skepticism with regard to perceptual information, or develop a theory of perceptual information that differs considerably from the one Dretske proposes. This is an important epistemological result on its own.

4. KNOWLEDGE CLOSURE

Let us now return to knowledge closure. A person knows that *p*, we are told, if and only if her belief that *p* is caused (or causally sustained) by the information that *p*. How can an abstract entity like information act in the world of causes? The metaphysics of causation is an area of persisting controversies. Yet, according to the standard

view, as advocated for instance by Donald Davidson, David Lewis, Jaegwon Kim, and also Dretske himself,¹⁰ causes must be events. Dretske adds that they are events which are effective in virtue of having certain properties (Dretske, 1981a, p. 88). Regarding the relation between perceptual signals and beliefs, we may say that perceptual signals, construed as events, cause beliefs in virtue of the fact that they carry certain pieces of information. Thus, Dretske's account of knowledge may also be represented as follows:

- (K*) An epistemic subject K knows of some (perceptual) object s that it has the property F if and only if there is a signal r which carries the information that s is F , and r causes (or causally sustains), in virtue of carrying that information and via a non-deviant causal chain, K 's belief that s is F .

The qualification "via a non-deviant causal chain" is needed to exclude cases in which a signal causes K to hold a true belief but produces that belief in the wrong way, as for instance when some neurological instrument, triggered by some signal, produces the belief in K by direct brain stimulation.¹¹ Causal theories of knowledge and belief are faced with the task of providing an account of non-deviant causal chains that lead from (potential) sources of epistemic attitudes to the appropriate beliefs. This has proven to be a difficult task. But let us assume, for the sake of the argument, that a satisfying account can be worked out.

Next it should be noted that there are important independent reasons for adding a condition to the simple formulation of the closure principle PEC that we (and Dretske) have been working with so far ("If K knows that p and knows that p entails q , then K knows that q ").¹² It is widely acknowledged that, as it stands, PEC is problematic simply because it is doubtful whether *belief* is closed under known (or believed) entailment. On a common-sense understanding of the notion of belief, for instance, it is not thus closed: it is certainly not a conceptual truth that people believe every logical consequence of what they believe, even if they see that they follow. However, what we are after when investigating *epistemic closure* (as opposed to what may be called "doxastic closure") is an illuminating analysis of conditions specific to the concept of knowledge. The principle to be examined should therefore not be vulnerable on account of incorporating a dubious assumption about belief.

The obvious way to protect the principle from this lapse is to add a condition to PEC which says that the subject at least also *believes* q .¹³ However, this manoeuvre does not yet exclude that the belief that q is held for some reason that has nothing to do with the subject's knowledge about p and p 's consequences. Hence we should strengthen the antecedent of the relevant closure principle to the effect that the belief that q be a *result of*, or be *based on*, the subject's knowledge (and beliefs) that p and that p entails q . In summary, then, the epistemologically interesting question is not whether PEC, but whether a principle of epistemic closure along the following lines is true:

(PEC*) If K knows that p , and knows that p entails q , and believes q on the basis of knowing (and hence believing) p and knowing (and hence believing) that p entails q , then K knows that q .

I believe this principle *is* true, but I shall not argue for this claim here. Stephen Hales (1995) has argued that it is even *trivially* true. Advocates of such a principle, he claims, face the task of showing that it is nevertheless not philosophically empty. The argument I will present now meets this requirement. My argument shows that if the principle is true, certain externalist theories of knowledge fall prey to skepticism. This, I take it, is not a philosophically trivial result. Moreover, the important point in the present context is that, whatever we may think about PEC*, Dretske's epistemology is committed to this principle. And if this is true, and if we accept, with Dretske, that the skeptic is right in claiming that we don't know that skeptical hypotheses are false, Dretske's theory also falls victim to knowledge skepticism.

So let us feed into PEC* the relevant parameters. First, it seems clear that in Dretske's account the notion of believing q "as a result" or "on the basis" of believing p must be interpreted in causal terms. His theory of knowledge clearly suggests a causal interpretation of the epistemic basing relation. Let us tacitly understand that non-deviant causal-chain conditions are fulfilled. Let e (as an instantiation of p) again be some suitable empirical proposition and not- h (as an instantiation of q) one of e 's anti-skeptical consequences. Given a causal interpretation of the basing relation, the antecedent of PEC* then gives us that K 's belief that e , which is part of what constitutes K 's knowledge that e , is one of the causes of K 's belief that not- h . But if K 's belief that e is an instance of perceptual belief, it must, according to Dretske, have been caused by some perceptual signal r which carries the information e . It thus follows that r is also among the causes of K 's

belief that not-*h*. Furthermore, as I have shown in Section 3, if *r* carries the information that *e*, and *K* knows that *e* entails not-*h*, then *r* must also carry the information that not-*h*. Hence *K*'s belief that not-*h* has been caused by a signal that carries the information not-*h* – and isn't that precisely what suffices, according to Dretske's definition, for that belief to be an instance of knowledge?

Prima facie, there is still a fly in the ointment. It may be objected that, strictly speaking, *r* can only be regarded as a partial cause of *K*'s belief that not-*h*. For PEC* rules (in its causal interpretation) that this belief must be caused, via a non-deviant causal chain, by the belief that *e* and the belief that *e* entails not-*h*. Would Dretske nevertheless say that *K*'s belief that not-*h* amounts to knowledge? This is the story he has to tell. I think he would in fact admit that we are dealing with a case of knowledge here. More importantly, he certainly *should* say so. For what could conceivably prevent such a belief from deserving the title of knowledge? By hypothesis, the belief that not-*h* is caused (a) by a signal that carries the information that not-*h*, and (b) by the subject's true belief that her belief that not-*h* is entailed by a true belief she holds (i.e. her belief that *e*). What could be safer than that? What else could be demanded in an externalist theory such as Dretske's?

5. CONCLUSION

I take these reflections to show that it will be preferable for Dretske to opt for the second alternative mentioned at the end of Section 3 and give up his probabilistic theory of information. At least he needs to revise that theory substantially if it is to handle the problems discussed in this paper. Until an alternative is on the table, however, his theory, which is one of the most detailed and most powerfully developed externalist accounts of knowledge, has lost its foundation. Moreover, it cannot be regarded as a viable alternative to contextualism. Dretske is one of the great pioneers of externalism in epistemology, and much of the force and attraction of his externalism has been due to its carefully crafted account of information. We now realize that, as it stands, this theory does not escape skepticism.

ACKNOWLEDGMENTS

For helpful comments and discussions I would like to thank Elke Brendel, Mark Siebel, an anonymous referee for *Erkenntnis*, and

especially Fred Dretske. I cannot resist the temptation to say that Dretske has agreed, in correspondence and in his reply to my presentation of this material at the conference in Mainz, that the arguments developed in this paper are on target. He seems to be fairly optimistic, however, that his theory can be repaired without giving up too much of its central idea. I do not think that the prospects for this are too rosy; indeed I believe that the issues I have discussed here point to a more general and fundamental problem of externalist accounts of knowledge: they avoid higher-order skepticism by rejecting the KK thesis, but they are committed to closure (and hence invite skeptical arguments that work with closure). However that may be, it is always much easier to criticize a theory than to come up with a relevant alternative. I wish I had one.

NOTES

- ¹ Dretske (1981a, p. 65). See also Dretske et al. (1983, p. 57). For his definition of knowledge see (1981a, p. 86), and (1983, p. 58). Instructive reconstructions and discussions of Dretske's information theory of knowledge can be found in Foley (1987), the "Open Peer Commentary" in Dretske (1983), and in the papers collected in McLaughlin (1991).
- ² For various epistemological problems the KK thesis creates see for example Greco (2000, pp. 181–184).
- ³ As I shall discuss below, this simple formulation of the closure principle is vulnerable to objections that have nothing to do with skepticism. I will embellish this formulation in Section 4, but for a start it will suffice to work with PEC. Dretske's argument that closure invites skepticism will be sketched in the next section.
- ⁴ That Dretske's famous Zebra Case does not in fact constitute a genuine counterexample to the closure principle has for example been argued by Jonathan Vogel (1990). Yet, Vogel concedes that close cousins to Dretske's example (such as the Car Theft Case) may undermine closure. The problem is that such examples appear to have features that cannot be exploited by arguments for (global) skepticism. Another supposed counterexample to closure has been presented by Robert Audi (1988, p. 77). For a critical discussion of Audi's argument see Feldman (1995). Mark Heller (1999) proposes an "Expanded Relevant Alternatives Principle" that, as he argues, avoids the problems of Dretske's original account. Peter Klein (1995) thinks that Dretske's argument against closure is correct regarding "externally situated evidence," but that it fails to realize that closure can be defended on internalist grounds. The reason, he says, is that the principle "does not require that the source of justification for the entailed proposition is anything other than the entailing proposition" (p. 221). My argument in this paper shows that if the notion of "externally situated evidence" is spelled out in terms of a Dretske-style theory of information, his argument does not even work for external evidence.

- ⁵ For groundbreaking work in the area see for example Cohen (1988), DeRose (1995), and Lewis (1996).
- ⁶ For a clear and representative statement of this approach see Cohen (2000, p. 103).
- ⁷ Sometimes Dretske is most explicit about this: “The *only* way to preserve knowledge of homely truths, the truths everyone takes themselves to know, is [...] to abandon closure” (Dretske forthcoming, my emphasis).
- ⁸ For a helpful discussion of these indirect arguments see for instance Williams (1996, pp. 330–336). Williams argues in chapter 8 of *Unnatural Doubts* that virtually all externalist attempts to show that closure fails are unsuccessful.
- ⁹ Put in the terminology Dretske introduces in *Knowledge and the Flow of Information*, we may say that, if *K* receives the information that *e*, and knows that *e* entails not-*h*, then the information that not-*h* is *analytically nested* in the information that *e*. Cf. Dretske (1981a, p. 71).
- ¹⁰ Cf. Dretske (1981a, p. 32): “Causality is a manifestation of a regular, lawlike, succession between events of type *C* and events of type *E* under relevantly similar conditions.”
- ¹¹ That Dretske’s theory falls victim to such counterexamples has been argued by Alvin Plantinga. See Plantinga (1993, pp. 195–197).
- ¹² For more detailed examinations of various epistemic closure principles see for example Brueckner (1985), Hales (1995), Luper (2001), or Barke (2002, pp. 26–43).
- ¹³ This is an embellishment considered for example by Brueckner (1985, p. 91). He argues that only a closure principle that has been altered along such lines will do justice to the skeptic, for “his target are not the careless epistemics” who fail to believe what they know follows from what they know. Instead, “the target knowers are ones who know that certain propositions are skeptical counterpossibilities to what they claim to know and who believe that these possibilities do not obtain” (1985, pp. 91f.).

REFERENCES

- Audi, R.: 1988, *Belief, Justification, and Knowledge*, Wadsworth, Belmont, CA.
- Barke, A. C. J.: 2002, *The Closure of Knowledge in Context*, Mentis, Paderborn.
- Brueckner, A.: 1985, ‘Skepticism and Epistemic Closure’, *Philosophical Topics* **13**, 89–117.
- Cohen, S.: 1988, ‘How to be a Fallibilist’, *Philosophical Perspectives* **2**, 91–123.
- Cohen, S.: 2000, ‘Contextualism and Skepticism’, *Philosophical Issues* **10**: *Skepticism*, edited by E. Sosa and E. Villanueva, 94–107.
- DeRose, K.: 1995, ‘Solving the Skeptical Problem’, *The Philosophical Review* **104**, 1–52. Reprinted in DeRose and Warfield 1999.
- DeRose, K. and T. Warfield (eds.): 1999, *Skepticism – A Contemporary Reader*, Oxford University Press, New York, Oxford.
- Dretske, F. I.: 1970, ‘Epistemic Operators’, *The Journal of Philosophy* **67**, 1002–1023. Reprinted in Dretske 2000, 30–47.
- Dretske, F. I.: 1971, ‘Conclusive Reasons’, *Australasian Journal of Philosophy* **49**, 1–22. Reprinted in Dretske 2000, 3–29.
- Dretske, F. I.: 1972, ‘Contrastive Statements’, *The Philosophical Review* **81**, 411–430.

- Dretske, F. I.: 1981a, *Knowledge and the Flow of Information*, Basil Blackwell, Oxford.
- Dretske, F. I.: 1981b, 'The Pragmatic Dimension of Knowledge', *Philosophical Issues* **40**, 363–378. Reprinted in Dretske 2000, 48–63.
- Dretske, F. I.: 1983, 'Précis of Knowledge and the Flow of Information', 'Open Peer Commentary', and 'Author's Response', *The Behavioral and Brain Sciences* **6**, 55–90.
- Dretske, F. I.: 1991, 'Dretske's Replies', in McLaughlin 1991, 180–221.
- Dretske, F. I.: 2000, *Perception, Knowledge, and Belief – Selected Essays*, Cambridge University Press, Cambridge.
- Dretske, F. I.: 2004, 'Externalism and Modest Contextualism', *Erkenntnis* **61**, 173–186.
- Dretske, F. I.: Forthcoming, 'The Case Against Closure', in E. Sosa and M. Steup (eds.), *Contemporary Debates in Epistemology*, Basil Blackwell, Oxford.
- Feldman, R.: 1995, 'In Defence of Closure', *Philosophical Quarterly* **45**, 487–494.
- Foley, R.: 1987, 'Dretske's 'Information Theoretic' Account of Knowledge', *Synthese* **70**, 159–184.
- Greco, J.: 2000, *Putting Skeptics in Their Place*, Cambridge University Press, Cambridge.
- Hales, St. D.: 1995, 'Epistemic Closure Principles', *The Southern Journal of Philosophy* **XXXIII**, 185–201.
- Heller, M.: 1999, 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* **77**, 196–208.
- Klein, P.: 1995, 'Skepticism and Closure: Why the Evil Genius Argument Fails', *Philosophical Topics* **23**, 213–236.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567. Reprinted in DeRose and Warfield 1999.
- Luper, St.: 2001, 'Epistemic Closure Principles', *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/closure-epistemic/>.
- McLaughlin, B. P. (ed.): 1991, *Dretske and His Critics*, Basil Blackwell, Oxford.
- Plantinga, A.: 1993, *Warrant: The Current Debate*, Oxford University Press, Oxford.
- Vogel, J.: 1990, 'Are There Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism*, Kluwer Academic Publishers, Dordrecht.
- Williams, M.: 1996, *Unnatural Doubts – Epistemological Realism and the Basis of Skepticism*, Princeton University Press, Princeton.

Department of Philosophy
 University of Aberdeen
 Old Brewery, High Street
 Aberdeen, A24 3UB, UK
 E-mail: christoph.jaeger@abdn.ac.uk