Notes for Week 2 of *Contemporary Debates in Epistemology*

01/28/09

Branden Fitelson

1 Preliminaries: Some Closure Principles

There are various closure principles for knowledge that one might entertain. Here are several:

1. If (i) S knows that p and (ii) $p \vDash q$[1], then S knows that q.
   - This is clearly false. S may not even *believe* q in such a case.
2. If (i) S knows that p and (ii) $p \vDash q$ and (iii) S believes that q, then S knows that q.
   - This is also false. S may believe q, but *not on the basis of* p (*i.e.*, they may not *infer* q from p).
3. If (i) S knows that p and (ii) $p \vDash q$ and (iii) S *infers* q from p, then S knows that q.
   - This is also false. S may infer q from p, but in an *incorrect* (or "incompetent") way.
4. If (i) S knows that p and (ii) $p \vDash q$ and (iii) S *competently infers* q from p, then S knows that q.
   - This is also false. S may (during the inference) obtain some evidence against p (or lose some evidence for p), and thereby lose their knowledge that p. You may also worry about (ii) not being *known by* S.
5. If (i) S knows that p and (ii) S knows that $p \vDash q$ and (iii) S *competently infers* q from p while (iv) *maintaining their knowledge that* p, then S knows that q.
   - Note that Hawthorne uses "deduces" rather than "infers", and he drops (ii) altogether! I prefer to require (ii) and talk about "competent inference". This raises some questions about "competent deduction" and "competent inference". First, why must the entailment hold in order for a deduction to be "competent"? Why can't someone "competently deduce" q from p, even if q doesn't really follow from p (think of a case in which truth *happens to be* preserved, but not *necessarily so*, because all logicians have been wrong in some subtle way about the nature of logical necessity)? More generally, what *else* is built-in to this notion of "competent inference"? It had better not be synonymous with "infers in a knowledge-preserving way". If it were, then the fact that $p \vDash q$ *would* be *irrelevant*. Is there an underlying assumption here that an inference from p to q *preserves knowledge* if (I) some (extensional) relation obtains (or is known to obtain) between p and q (*e.g.*, entailment), and (II) a (suitable) "competent inference" is made by the agent? If that is the recipe, then what is the analogous (I)-relation in the *non*-deductive case (see below for more on this question)? After all, there must be *some* knowledge-preserving *non*-deductive inferences, right (unless we're going to be inductive skeptics)? Indeed, the interesting cases (vis-a-vis skepticism) *are non*-deductive!
6. If S knows that p and S knows that $p \vDash q$, then S has all that it takes, evidentially speaking, to know that q.
   - This is the version of closure that Dretske discusses. Below, I will work with more precise versions of this sort of closure principle. I think this will help us to see what Dretske is worried about.

2 Dretske's Examples

Dretske describes various examples which he thinks are counterexamples to his rendition of closure: (6). I think the following version of closure is a bit more perspicuous with regard to Dretske's worries:

7. If S knows — on the basis of E — that p and S knows that $p \vDash q$, then $\underline{E}$ is a sufficient evidential basis to undergird knowledge (on the part of S) that q.

[1] We are assuming here that $\vDash$ is *classical* entailment — *i.e.*, *necessary truth-preservation*. I will discuss the normativity of logic in week 15. In that context, some authors (*e.g.*, Hartry Field) insist that "entailment" is *inherently epistemic* (and, hence, that $\vDash$ is *not* necessary truth-preservation). That muddies the waters, since it becomes more difficult to understand the *content* of closure principles (which now become *second-order*, epistemic principles). We'll return to these thorny issues about logic in week 15.

One of Dretske's examples (which I will call The Zebra Example) involves an agent S who finds himself at the Zoo, at an exhibit labeled "Zebra". And, in the exhibit, S sees what appears to him to be a Zebra. When asked (by his annoying philosophical companion) whether he knows that (p) the animal in the exhibit is a zebra, S says "yes". But, then, the companion asks whether S knows that (q) it is not the case that (p') the animal is a cleverly disguised horse. At this point, S pauses and says "well, I guess I don't know *that*", even though S is well aware that $p \vDash q$. This contains the ingredients of a *prima facie* counterexample to (7).

Let E be S's evidence for p. Presumably, this will involve perceptual evidence, as well was other sorts of (background) evidence. Let's not worry too much about whether E is sufficient to ground *knowledge* that p (since that already gets us into potential skeptical worries). Instead, let's just take for granted that E supports p to *some* (positive) degree. Now, let's think about p' — the possibility that the animal is a cleverly disguised horse. Let's state this possibility more precisely (to capture the *epistemic* import of "cleverly").

(p') The animal in the exhibit is a horse, but one which is disguised in such a way that it is *indistinguishable from a zebra — relative to evidence E* (which is the basis on which S believes that p).

More precisely, let Za be the claim that the animal a is a zebra, and let Ha assert that a is a horse. Then:

(p') a is such that [Ha, but *E supports Za and Ha equally* (or, at least, that E does not favor Za over Ha)].

Now, I think we can see more clearly what Dretske is worried about. His point is that, by construction, the example seems to be such that $\underline{E}$ does *not* favor p over p'.[2] And, I also think it's rather plausible to suggest that if E does *not* favor p over p', then S is *not* in a position to know — *on the basis of* $\underline{E}$ — that q is true (might one deny this?). Thus, this seems (*prima facie*) to be a rather compelling counterexample to (7). At this point, you're probably thinking that (7) clearly needs to be modified in *at least* in the following way:

7′. If S knows — on the basis of E — that p and S knows that $p \vDash q$, then $E' = \{E +$ the knowledge that $p \vDash q\}$ is a sufficient evidential basis to undergird knowledge (on the part of S) that q.[3]

But, it's difficult (for me, at least) to see how adding the *logical* knowledge that $p \vDash q$ is going to help with the *empirical indistinguishability* problem that undergirds the example. Presumably, the problem with E is *not* that it's lacking the power to distinguish alternative *logical* claims (*i.e.*, giving E the power to favor $p \vDash q$ over $p \nvDash q$ doesn't really seem to help here). This, it seems to me, is a good way to unpack Dretske's worry.[4]

3 Digression: Probability, Support, and *Extensional* Closure Under $\vDash$

Fumerton[5] makes extensive use of the following principle:

(†) To be justified in believing one proposition p on the basis of another proposition E, one must be (a) justified in believing E and (b) justified in believing that E makes probable p.

We're talking about knowledge and not justification here (yet). So, I'm not going to discuss these sorts of principles today. The main thing I want to discuss is the "making probable" concept in clause (b). What Fumerton means here when he says that "E makes probable p" is that some conditional probability $\Pr(p \mid E)$ is high. That is, that $\Pr(p \mid E) > t$, for some threshold value t. **Fact**: this "making probable" relation *is* (extensionally) closed under entailment. That is, if $\Pr(p \mid E) > t$ and $p \vDash q$, then it follows that $\Pr(q \mid E) > t$.

However, there is (plausibly) *another* necessary condition for "justified belief in p *on the basis of* E", and one which *violates* closure (even extensionally). I'll call this (c) E must *probabilistically support* p — that is, E must *raise the probability of* p [*i.e.*, $\Pr(p \mid E) > \Pr(p)$]. Here is a concrete counterexample to the (extensional) closure of "support" under entailment. We're going to sample a card c at random from a standard deck of cards. Let p be the claim that c is the A♠, let E be the claim that c is a black card, and let q be the claim that c is *some* Ace. In this example, E supports p, since $\Pr(p \mid E) = \frac{1}{26} > \Pr(p) = \frac{1}{52}$, and $p \vDash q$, but E *does not support* q, since $\Pr(q \mid E) = \frac{1}{13} = \Pr(q)$. Thus, unlike the "making probable" relation, the "probabilistic

[2]This doesn't follow immediately, but it seems plausible to me. What higher-order epistemic principle(s) does this presuppose?
[3]With this amendment, (7′) is more-or-less equivalent to Dretske's own rendition of closure: (6).
[4]I also think this is the sort of phenomenon that underlies *skeptical* arguments, generally. More on that later in the course.
[5]This is discussed in his book *Metaepistemology and Skepticism*, excerpts of which are on our website. Specifically, see "The Structure of Skeptical Arguments" for his statement of (†), which he calls The Principle of Inferential Justification.

support/relevance" relation is *not* (extensionally) closed under entailment. Intuitively, I think the right thing to say in this sort of case is that (i) E supports p, but (ii) E does *not* support q — *even though* (iii) $p \vDash q$. This gives us at least one salient property of E (and p and q) which could explain some sort of closure failure. The suggestion I am making here is that the ability of E to *favor* p over p' in our Dretske example above depends not only on "how probable E makes p *vs* p'", but also on *how much of a difference* E *makes to the probability of* p *vs* p' (*i.e.*, on *how probabilistically relevant* E is to p *vs* p'). And, as we have just seen, probabilistic relevance is not even *extensionally closed* under entailment. This phenomenon, I suspect, may have some bearing on what is happening in examples of this kind. One final note: Hawthorne (briefly) discusses the following *equivalence* principle [this is the same as principle (5), but with $\dashv\vDash$ rather than $\vDash$]:

8. If (i) S knows that p and (ii) $p \dashv\vDash q$ and (iii) S *competently infers* q from p while (iv) *maintaining their knowledge that* p, then S knows that q.

An "evidentialist" version of such a principle might be:

9. If S knows — on the basis of E — that p and S knows that $p \dashv\vDash q$, then $E' = \{E$ + the knowledge that $p \dashv\vDash q\}$ is a sufficient evidential basis to undergird knowledge (on the part of S) that q.

This is fundamentally different than the previous principles, since it's much more difficult to think of *any* salient (*e.g.*, probabilistic) relations that aren't even *extensionally* closed under *logical equivalence*!

4 Hawthorne's Misprint Example

Hawthorne doesn't seem to be moved by Dretske's example(s). Perhaps this is because he formulates closure as (5) rather than (6), (7) or (7′). His idea seems to be that if you *competently deduce* q *from* p, while maintaining your knowledge that p (say, in The Zebra Example), then you *do* come to know q. From an "evidentialist" [*i.e.*, a (6)–(7′)] perspective, however, it is somewhat mysterious as to *how* your *evidence* could *undergird* knowledge that q in such an case. On the other hand, Hawthorne gives another case, which he thinks raises (*prima facie*) problems — *even for* (5). His Misprint Case runs as follows (my notation added):

> Suppose I learn from a report (E) in a reputable newspaper that (p) Manchester United beat Coventry City. That piece of information, in combination with the fact that the newpaper says that Manchester United beat Coventry City, entails that (q) *it is not the case that* (p') due to a misprint, the newspaper said that Manchester United beat Coventry City when in fact Manchester United lost to Coventry City. Call this proposition "Misprint" (q). On the one hand, it seems odd to suppose that one can know Misprint (q) in advance of reading the newspaper (in advance of learning E); on the other, it seems odd to suppose that one can come to know Misprint (q) by looking at the newspaper (by learning E).

While it has some additional wrinkles, I think this example is not that much different than the gloss of Dretske's original example I gave above (which is explicit about the evidential bases for different beliefs that are featured in the story). As such, I think we can (without loss of generality) focus on The Zebra Example.

5 Vogel on The Zebra Example

Vogel argues that Dretske's examples are not really counterexamples to closure — not even to principles (6)–(7′). I think Vogel's discussion is more on target than Hawthorne's, since he addresses the (alleged) "evidentialist" counterexamples directly. Vogel concedes that your *perceptual* evidence *alone* (at the Zoo) does not undergird knowledge that the animal is not a cleverly disguised horse. But, he claims that your *total* evidence *must* do so — *if* it undergirds knowledge that both p and $p \vDash q$. He says (my parens):

> Your belief that the animal at the zoo is a zebra is justified *in part* by your visual evidence, but it is *also* supported by the background information that counts against the animal's being a disguised mule (*because* it counts *in favor of* or "supports" the animal's being a zebra, which *entails* it's *not* a mule?).

As far as I can tell, Vogel *seems* to be assuming something like the following (where E is S's *total* evidence):

(‡) If E evidentially supports p (for S), and (S knows that) $p \vDash q$, then E evidentially supports q (for S).

But, as we have already seen above, this principle might fail (even *extensionally*), if "evidential support" involves *probabilistic relevance* (or, probably, any type of "difference-making"). It is true that "makes probable" is closed under entailment in this way, but that may not be sufficient to ensure the premises of Vogel's arguments in favor of principles like (6)–(7′). In any case, I suspect that such unclarities in our notion(s) of "evidential support" may have something to do with the present controversy over closure principles.