

Closure NOTES

BT Miller

February 23, 2026

Dretske argues against the Closure Principle:

Closure: if S knows that P is true and knows that P implies Q, then, evidentially speaking, this is enough for S to know that Q is true.

Closure is especially important in discussions of external world skepticism.

Consider the following *inconsistent triad*:¹

1. S knows that they have hands
2. If S can know that they have hands, then S can know that they're not a handless BIV
3. S can not know that they're not a handless BIV

¹ NB this is not an argument, just a set of three propositions that are jointly inconsistent. Of course one could formulate a valid argument by taking any two of the three propositions as premises, and the negation of the third as the conclusion.

At least one of these claims is false.

- Skeptics reject 1.
 - this was Hazlett's approach.
- Mooreans and neo-rationalists reject 3.
 - We'll discuss these views later in the semester.
- If we reject 2 then we can maintain ordinary intuitions about ordinary knowledge (e.g. that I know I have hands) and also about non-knowledge regarding skeptical hypotheses (e.g. BIV, dream, etc).
 - But, that requires that we reject Closure.
 - This is often seen as a very costly move, since Closure is very intuitively plausible.

Softening us up for closure rejection

First point: lots of *propositional attitudes* are not closed under known entailment. Example:

1. S regrets drinking three martinis last night
2. S knows that drinking three martinis last night entails drinking something last night
3. Therefore, S regrets drinking something last night.

This is an invalid argument, because *regret* is not closed under known entailment.

In contrast, *truth* is closed under known entailment². The following is valid:³

1. P is true
2. S knows that P entails Q
3. Q is true

This is meant to show that the closure of *knowledge* under known entailment is not trivial

1. S knows that P
2. S knows P entails Q
3. Therefore, S knows Q

Is this valid? Defenders of closure say yes⁴, Dretske says no.

Second point⁵: sometimes evidential support is not transmitted

- One can *see* that there are cookies in the jar.
- One cannot *see* that mind-independent material objects exist, even if we can come to know it in other ways.⁶
- So there are reasons/ evidence for believing that there are cookies in the jar that are not reasons/ evidence for the existence of mind-independent material objects
- These examples of reason non-transmission are very plentiful
- This does not quite imply that closure is false:
 - It's possible that I know of the existence of mind-independent material objects on other grounds, so this isn't a case where I know P by not the entailed Q.
- But it does soften us up for the rejection of closure.

² Though of course truth is not a propositional attitude.

³ The factivity of knowledge is important here.

⁴ Once elaborate the argument by elaborating on whether S has considered Q, has formed the belief that Q, the basis of that belief, etc.

⁵ Claim?

⁶ Here that *mind-independent objects exist* is an example of what Dretske calls 'protoknowledge': propositions that must be true in order for me to know by perception that P, but which cannot themselves be perceived to be true. NB this account of protoknowledge is perfectly acceptable to those who accept closure. Such a person might hold, for instance, that I perceive that there are cookies in the jar, and then *infer* – not perceive – that material objects exist.

BTM comment: there's some oddness to this argument. Dretske writes:

A claim to have found out, by looking, that there are cookies in the jar is not a claim to have found out, by looking, that there is a material world... one surely isn't claiming to see that there are physical objects in claiming to see there are cookies in the jar. After all, hallucinatory cookies "in" hallucinatory jars can look exactly like real cookies in real cookie jars. So one cannot, not by vision alone, distinguish real cookies from mental figments. One cannot see that the world really is the way it visually appears to be. A way of knowing there are cookies in a jar – visual perception – is not a way of knowing what one knows to be implied by this – that visual appearances are not misleading.

Suppose the possibility of hallucination really does prevent me from 'finding out, by looking' that physical objects exist. Why does that same possibility not prevent me from finding out by looking that there are cookies in the jar? The two cases are different in many ways, and maybe one of those differences makes a difference when it comes to what beliefs can be supported by perception alone. But how can a property common to both cases make that kind of difference?

(end comment)

Abominable Conjunctions

One objection to rejecting closure: it leads to statements like REF:

REF: The refrigerator is empty, but has lots of things in it.

Dretske agrees one would never want to say something like this, but that doesn't mean it's never *true*.

If this is true, that's because we've switched senses of 'empty' and 'full' midway through the sentence:

- Imagine the fridge is empty of food, but full of air molecules.

Is there a real point here, or is this some cheap linguistic trick? Compare:

CON: Fred went to the bank, and it's not the case that Fred went to the bank.

- This sentence appears to assert a straightforward contradiction, and all contradictions are false.

- But 'bank' is ambiguous: it can refer to a financial institution, or to the side of a river.
- If Fred cashed his pay check by didn't go fishing, then CON might express a truth, but only if we use this same 'cheap linguistic trick' to interpret CON
- It would be a mistake argue on this basis that some contradictions are true.

Dretske is not employing a cheap linguistic trick here – he's making an interesting and important claim.⁷

Closure-deniers are accused of being forced to accept abominable conjunction BIV:

⁷ This is not to endorse his claim.

BIV: Fred knows that he has hands, but he doesn't know that he's not a BIV.

Dretske's response is that any interpretation on which this sentence is true will employ the same kind of cheap linguistic trick from above:

The second conjunct introduces possibilities normally assumed to be irrelevant (not counted as possibilities) by someone who asserts the first conjunct. (32)

NB this seems to rely on a Relevant Alternatives account of knowledge.

Later in the paper Dretske expresses some sympathy for that view, but at this point he doesn't seem to want his argument to depend on some further contentious theory.

Heavyweight Implications

Denying closure does not imply denying the possibility of knowledge through inference. Just some inferences. Which ones?

Heavyweight implications vs lightweight implications

Coming to know that P depends on some kind of ‘indicator’ that p is true: a perceptual experience, a piece of heard testimony, etc.

Example: the indicator is the experience of the cookies in the jar.

Lightweight implication: that the jar is not empty, that it’s not full of water, etc.

Heavyweight implication: that a physical world exists, that I’m not the victim of a cookie-oriented evil deceiver, etc.

Key claim:

‘...the relationships we describe by speaking of one thing indicating or carrying information about another condition, P, are relations that do not relate this thing to (all) the known implications of P.’ (36)

I.e., the indicator carries information about the lightweight implications, not the heavyweight implications.

In support:

It is important, therefore, to understand that the relationships we describe by speaking of one thing indicating or carrying information about another condition, P, are relations that do not relate this thing to (all) the known implications of P. Though we rely on a measuring instrument to indicate the value of Q – that it is, say, 5 units – the measuring instrument does not indicate, does not carry information, that Q is not being misrepresented. The position of a pointer (pointing at “5”) on a well functioning meter indicates that the value of Q is 5 units, but it does not thereby indicate that the instrument is not broken, not malfunctioning, not misrepresenting a value of 2 units (say) as 5 units. Even when instruments (and this includes the human senses) are in perfect working order, they do not – they cannot – carry information that what they are providing is genuine information and not misinformation. That isn’t an instrument’s job. The job of a Q-meter is to tell us about Q, not about its reliability in providing information about Q. So even though we know that if Q really is 5 (as the meter says), then Q can’t be 2 being misrepresented (by the meter) as 5, the meter gives us the first piece of information (tells us that Q is 5) without giving us the second piece of information: that Q isn’t really 2 being misrepresented as 5. That isn’t information the Q-meter supplies. It is a heavyweight implication of the information a Q-meter supplies. Anyone who really wanted to know whether the instrument was misrepresenting Q (when it registered “5”) should not look more carefully at the pointer.

A pointer pointing at “5” does not carry that information. The information it carries (when things are working right) is that Q is 5, not that Q doesn’t just appear to be 5.

Hawthorne

Dretske’s distinction between lightweight consequences we can know by inference, and heavyweight consequences we can’t, is intuitive and plausibly tracks ordinary usage.

Objection: closure deniers repeatedly fail to characterize the heavy/lightweight distinction in ways that provide plausible verdicts.

Rough and ready characterization⁸

Heavyweight consequence of P: a consequence of P about which we all have some strong inclination to think that P is not the sort of thing that one can know by the exercise of reason alone and also that P is not the sort of thing that one can know by use of one’s perceptual faculties (even aided by reason)

⁸ NB Hawthorne is not worrying about a vague boundary between heavy- and lightweight consequences, though of course that is another worry. This is good practice: vagueness is a perennial problem for almost everyone, independent of closure, so it’s somewhat weak to claim it presents a specific problem for Nozick.

First example:

Sensitivity

Nozick has argued for Sensitivity as a necessary condition on knowledge:

sensitivity: if P were not true, then S would not believe it

We’ve noted that sensitivity is inconsistent with closure.

- Of course Nozick would be fine with that – he rejects closure anyway.
- But there are lots of cases where the closure failure results in rejecting ordinary, *lightweight* consequences:

Suppose a real dog and a fake cat are in a room, the former keeping the latter from view. I look at the dog and form the belief that there is a dog in the room. From this, I infer there is an animal in the room. Suppose further that if there hadn’t been a dog in the room, I would have seen the fake cat and formed the belief that it was a (real) animal. Then my belief that there is a dog in the room passes the Nozick test but my belief that there is an animal in the room does not.

That there’s an animal in the room is a *lightweight* consequence of my perceptual ‘indicator’.

- The aim is to provide an account on which I can infer lightweight consequences of my evidence, but not heavyweight indicators.
- But *that there's an animal in the room* is definitely a lightweight consequence.
- So this way of rejecting closure fails to satisfy the aim above.
- So Dretske probably doesn't want to follow Nozick in accepting sensitivity as a necessary condition on knowledge.

Conclusive Reasons

Suppose the following is a necessary condition on knowledge:

Conclusive reasons: if I know P [on reasons R], then I would not have reasons R unless P.

Counterexample:

- I'm looking at the zebra area at the zoo, having zebra-like experiences
- on the basis of those experiences, I believe *there's a zebra in the cage, and no other animal*
- In all nearby worlds where I know *there's a zebra in the cage, and no other animal*, there
- But I don't satisfy CR:
 - Suppose I know *there's a zebra in the cage, and no other animal*: then I wouldn't have zebra-like experiences unless *there's a zebra in the cage, and no other animal*
 - But that's false: if it had been a cleverly painted mule, I'd still have those reasons (my experiences)
 - So CR fails
 - So if CR is a necessary condition on knowledge, I don't know that it's a zebra