

Metaphysics and Epistemology: Question 1

Wylie Breckenridge

Word count: 5184

What is the epistemological significance of the question whether knowledge is closed under known entailment?

1. Introductory

It would be difficult to deny that it is a significant epistemological question whether or not knowledge is closed under known implication (entailment), that is, whether or not the following principle is true:

Closure

If S knows that p , and S knows that if p then it follows that q , then S knows that q .¹

For on the one hand, we take the truth of Closure to be a fundamental fact about knowledge, and an important way to acquire it. How do I know that today is a weekday? Because I know that today is Friday, and I know that if today is Friday then it follows that today is a weekday.²

But on the other hand, it seems that the truth of Closure is incompatible with the truth of two other claims, each of which is intuitively attractive:

- (H) I know that I have hands
- (B) I do not know that I am not a (handless) brain in a vat (BIV).

For suppose that I know that I have hands. I also know that if I have hands then it follows that I am not a BIV.³ So, by Closure, I know that I am not a BIV. But wait, I don't know that!

What is going on? Are (H) and (B) a counterexample to Closure? Or is it that at least one of (H) or (B) is false? In either case, why do we find the false claim(s) intuitively attractive? Or are we wrong to think that there is any incompatibility here? If so, why? Call this family of questions, Problem.⁴

It is not immediately clear how we should respond to Problem. It is no good to simply deny one of Closure, (H) and (B) in order to maintain the other two. Why, for example, should I maintain Closure and (H) at the expense of (B), or maintain (H) and (B) at the expense of Closure, when my belief in each of the three seems to be equally strong? Any response that issues verdicts as to the truth values of Closure, (H), and (B) must at least be supported by appeal to an account of knowledge that can explain those truth values, and, just as importantly, by an account of why each is intuitively attractive.⁵ And any response that denies there is any incompatibility to be concerned about (we will be looking at one) must at least explain why there is not, and why it appears that there is.

¹ It is difficult to formulate the principle in a way that accurately captures the intuitive notion. I will follow the majority in not quibbling over the details. See Nozick ([1981], 170) and Forbes ([1984], 49-50) for a discussion of the issues involved.

² Forbes ([1984], 49-52) attempts an argument for Closure based upon general considerations about knowledge. But Brueckner [1985] presents a devastating response.

³ I will follow the majority in taking this as given.

⁴ Perhaps I should call it 'Paradox'. But I don't want to prejudge the appropriate response against anyone who wants to deny that there *is* a paradox. Thus, 'Problem'.

⁵ I set aside, then, the Moorean response, and its slightly more sophisticated version Modus Ponens Fallibilism. Cohen ([1999], 64-5) has a good discussion of each.

Most of us take the truth of Closure to be a fundamental fact about knowledge. And yet there is lingering doubt about whether or not it is true. For, as I hope to show, if you claim that Closure is true, then, given the current state of epistemology, you will be unable, when pressed, to give a satisfying response to Problem. But if you claim that Closure is false then you will not be able to either. That being the case, it would be wise for you to entertain some doubt about whether or not Closure is true. The epistemological significance of the question is that, while we believe and very much want Closure to be true, there is still some doubt about whether or not it is. It is a key unsolved problem for epistemology.

My aim, then, is to argue for what I claimed in the previous paragraph: that anyone who wants to deny Closure is unable to give a satisfying response to Problem, and that anyone who wants to maintain Closure is unable to as well. I will do so by considering, in Section 2, what seems to me the most promising approach to Problem for anyone who wants to deny Closure, and considering, in Section 3, what seems to me the most promising approach for anyone who wants to maintain Closure. I will argue that neither approach has yet achieved a satisfying response.

2. Responding to Problem without Closure: Nozick

If one wants to respond to Problem by denying Closure, then I think the most promising way on the market is to appeal to Robert Nozick's [1981] counterfactual analysis of knowledge.⁶ In its first incarnation, Nozick's account is that *S* knows that *p* if and only if

- (1) It is true that *p*
- (2) *S* believes that *p*
- (3) If it were not true that *p*, then *S* would not believe that *p*
- (4) If it were true that *p*, then *S* would believe that *p*.⁷

(We can take the truth conditions of the counterfactuals to be: it is true in world *w* that, if it were the case that *p* then it would be the case that *q*, just in case it is true that *q* in all the closest worlds to *w* in which it is true that *p*).⁸ Conditions (3) and (4) are intended to capture the idea that *S* truly believes that *p*, not by accident, but because her belief that *p* is connected to the fact that *p* in the right kind of way: she 'tracks' the truth.

We can appeal to this theory to explain why I believe (H) and (B). I do so in each case because I understand and apply the conditions (1) - (4).⁹ First, I believe (H), that I know that I have hands, because I believe that each of the following is true: it is true that I have hands; I believe that I have hands; if it were not true that I have hands, then I would not believe that I have hands (the closest worlds in which I do not have hands are all ones in which it is clear to me that I don't have hands); and if it were true that I have hands, then I would believe that I have hands (the closest worlds in which I do have hands are all ones in which it is clear to me that I do have hands). Second, I believe (B), that I do not know that I am not a BIV, because I believe that it is

⁶ One alternative is to appeal to Fred Dretske's [1970] 'relevant alternatives' account of knowledge. He claims that according to his account Closure is false, but it's not entirely clear that he is right (see Stine [1976] for a discussion of this issue). Furthermore, it seems that most theorists who appeal to the idea of relevant alternatives now do so as part of a contextualist response to Problem, and we will consider such responses in Section 3. For these reasons, I will not consider Dretske's account here.

⁷ For reasons that we need not be concerned with, Nozick modifies condition (4) to: if it were true that *p*, then *S* would believe that *p* and would not believe that not-*p* (178).

⁸ We need to deviate from the standard Lewisian possible worlds semantics and require that the set of closest worlds to *w* in which it is true that *p* includes some worlds distinct from *w* (otherwise condition (4) is made redundant by condition (1)). Appealing to possible worlds semantics is not an essential part of Nozick's account, but a helpful way to understand it. Nozick himself appeals to it.

⁹ Perhaps not explicitly.

not true that I am a BIV. Moreover, according to Nozick's theory, it is possible that I *correctly* believe both (H) and (B): it is possible, after all, that the world is just as I believe it to be.¹⁰

So we can use Nozick's account to explain why I believe both (H) and (B) - I just run the counterfactuals¹¹ and accept the results. And we can say that my concern that I cannot correctly believe both (H) and (B) is needless - it is based upon a belief in Closure, and Closure is false.

Nozick's account of knowledge has intuitive appeal, and it can be used to make a good start on responding to Problem. But two questions remain. First, how can we explain the fact that Closure is so intuitively attractive? Second, how well does Nozick's account fare as a *general* theory of knowledge, applicable outside the context of Problem?

The answer to the second question is: as it is formulated above, not so well. It seems that it is both too permissive (it attributes to us some knowledge that intuitively we don't have), and too restrictive (it does not attribute to us some knowledge that intuitively we do have).

Too permissive: a machine displays a hologram vase in front of a screen inside a box, just in case there is a real vase behind the screen inside the box; I see the hologram vase and come to believe truly (but do not know) that there is a real vase inside the box; if there were no real vase inside the box then I would not believe that there is, because there would be no hologram vase either; and if there were a real vase inside the box then I would believe that there is, because there would be a hologram vase there too.¹² Too restrictive: the screen is removed and the machine is now designed to display a hologram vase inside the box just in case there is no real vase inside the box; I see a real vase inside the box, and hence come to know that there is a real vase inside the box; but had there not been a real vase inside the box there would have been a hologram vase inside the box, and I would still have believed that there is a real vase inside the box (condition (3) fails).^{13,14}

Nozick, of course, recognises and discusses extensional problems with his theory, and acknowledges that his account requires some 'refinements and epicycles'. I cannot do justice to the issues involved here, but I, for one, cannot help but feel that the more epicycles are added to its basic formulation, the more Nozick's account loses its intuitive appeal, and the more it becomes vulnerable to counterexamples that prey upon the new (and existing) epicycles.

In any case, here I am more concerned about the apparent conceptual inadequacy of Nozick's account.

According to Nozick's account, whether or not I know that I am not a BIV depends upon what I would believe in the counterfactual situation in which I am a BIV. But Brueckner ([1994], 828-30) suggests that the focus on what we would believe is misplaced. For suppose that if I were a

¹⁰ It is a mistake to think that Nozick's account shows that I *do* correctly believe both (H) and (B). It only shows that I *can* correctly believe both (H) and (B). If I actually am a BIV, then I do not correctly believe (H): I do not know that I have hands, because it is not true that I have hands. (But I do correctly believe (B); some solace, at least. Assuming, that is, that if I were to utter 'I am not a BIV' in the vat then I would express the same proposition as I actually do. See Putnam's 'Brains in a Vat' in Putnam [1981].)

¹¹ Again, perhaps not explicitly.

¹² Nozick mentions a case like this, but merely notes that it is 'somewhat counterintuitive' (p. 190). Forbes ([1984], 45) points out that it is in fact a Gettier case that Nozick's account gets wrong.

¹³ Nozick ([1981], 179) complicates his analysis to accommodate a similar kind of case (his grandmother case).

¹⁴ These are, admittedly, artificial cases, and perhaps our intuitions about them are not so strong or reliable. But there are troubling everyday cases. See Brueckner [1984] and Vogel [1987] for some particularly interesting ones.

BIV, then I would have sensory evidence exactly similar to my actual sensory evidence, but my belief-forming mechanisms would be different from my actual ones in such a way that I would not mistakenly believe that I am not a BIV. Then it seems that I still do not know that I am not a BIV, even though according to Nozick's account I do. This is closely related to a concern raised by DeRose ([1995], 196-7). DeRose points out that, according to Nozick, I do not know that I am not a dog incorrectly thinking that I have hands; but, intuitively, I *do* know that: I know that I am not a dog incorrectly thinking that I have hands, because I know that I am not a dog. Brueckner suggests that the real issue when it comes to what we know is the quality of our actual evidence for what we actually believe, not what we would believe in certain counterfactual situations.¹⁵ If Brueckner is right, and I think that he is, then Nozick's counterfactual approach to analysing knowledge is fundamentally mistaken, and it becomes less plausible to think that with suitable modification it can ever be made extensionally correct.

The conceptual inadequacy of Nozick's account can also be seen in its inability to do the one thing remaining to do in response to Problem: explain why we believe that Closure is true, and in fact are reluctant to deny it even when faced with Problem. As far as I know, Nozick has little if anything to say about this. And I can't see what he might plausibly say. If I grasp the conditions for knowledge well enough to form the right beliefs about (H) and (B), then why do I not grasp them well enough to form the right belief about Closure (i.e. that it is false)? At least, once I recognise that I meet the conditions for knowing that I have hands, but not the conditions for knowing that I am not a BIV, then why does it not at least *then* become apparent to me that Closure is false? Instead, I get the feeling that perhaps I *don't* know that I have hands, or that I *do* know that I am not a BIV. I do not doubt that I understand what knowledge is, and I do not doubt the truth of any relevant counterfactuals, but I do doubt whether or not I meet the conditions for certain items of knowledge. That suggests that my grasp of the conditions for knowledge does not consist in my grasp of the truth-conditions for certain counterfactuals.

I think the same problem will be faced by any response to Problem that appeals to any analysis of knowledge according to which Closure is false. If we appeal to our grasp of the analysis to explain why we correctly believe both (H) and (B), then it seems that we cannot explain why we are reluctant to deny Closure, even when we recognise that (H) and (B) are an apparent counterexample to it. So I think that any response to Problem that appeals to Nozick's account, or any other account that denies Closure, will not be satisfying: if it can explain why we are intuitively committed to (H) and (B), then it will not be able to explain why we are intuitively committed to Closure.

3. Responding to Problem with Closure: Contextualism

If one wants to respond to Problem without denying Closure, then I think that the most promising approach is a contextualist one.¹⁶ According to contextualists, an utterance of the sentence '*S* knows that *p*' in context *c* is true just in case:

- (1) It is true that *p*
- (2) *S* believes that *p*
- (3) *S* meets standard *N* with respect to the proposition that *p*, where *N* is the epistemic standard determined by *c* (the context of *utterance*).

They do not specify conditions under which *S* knows that *p*, because, they believe, '*S* knows that *p*' is a context-sensitive expression: in one context it will express one proposition, in

¹⁵ I take it that Brueckner assumes that the quality of one's evidence is not to be measured in counterfactual terms, or else he might well end up agreeing with Nozick.

¹⁶ Prominent contextualists include Stine [1976], DeRose ([1995], [1999]), Cohen ([1988], [1991], [1999]), and Lewis ([1979], [1996]).

another, another (even though the sentence being uttered is the same). We might say, somewhat loosely, but I think without danger of confusion, that ‘*S* knows that *p*’ in context *c* expresses the proposition that *S* knows that *p* relative to standard *N*, where *N* is the epistemic standard determined by *c*.¹⁷

Their claim that knowledge ascriptions are context-sensitive is motivated by situations like the following. Smith, looking at his flight itinerary, utters, ‘I know that the flight has a layover in Chicago’. Mary, having to make an important business contact in Chicago, and aware that flight schedules sometimes change at the last minute, utters, ‘Smith does *not* know that the flight has a layover in Chicago.’ But, plausibly, Smith and Mary are not in disagreement. Smith is in a context in which epistemic standards are low. When he utters, ‘I know that the flight has a layover in Chicago’, he expresses the proposition that Smith knows that the flight has a layover in Chicago relative to standard Low. Mary, on the other hand, is in a context in which epistemic standards are high. When she utters, ‘Smith does not know that the flight has a layover in Chicago’, she expresses the proposition that Smith does not know that the flight has a layover in Chicago relative to High. It may be that both Smith and Mary are right.¹⁸

The above is the general contextualist idea; the details can be filled out in various ways. One might adopt a Dretske-style relevant alternatives approach, and say that an epistemic standard is a set of relevant alternatives to the proposition that *p*, that *S* meets the standard *N* with respect to the proposition that *p* just in case *S* can rule out the members of *N*, and that the mechanism by which the context determines *N* includes a ‘Rule of Relevance’, according to which mentioning the proposition that not-*p* makes the proposition that *p* a relevant alternative.¹⁹ Or one might adopt a Nozick-style counterfactual approach, and say that an epistemic standard is a region of possible worlds around the actual world, that *S* meets the standard *N* with respect to the proposition that *p* just in case *S* tracks the truth of the proposition that *p* throughout *N*, and that the mechanism by which the context determines *N* includes a ‘Rule of Sensitivity’, according to which mentioning the proposition that not-*p* widens *N*, where necessary, to include the worlds in which it is true that *p*.²⁰ (It is one of the virtues of the contextualist approach that it can accommodate key features of Dretske’s relevant alternatives and Nozick’s counterfactual accounts.)²¹ Here I am just concerned with the general contextualist approach.

By contextualist lights, what the closure principle says, when fully spelled out, is something like the following:

Closure (contextualised)

If *S* knows that *p* relative to standard *N*, and *S* knows that if *p* then it follows that *q* relative to standard *N*, then *S* knows that *q* relative to standard *N*.²²

¹⁷ When I use the expression ‘*S* knows that *p* relative to standard *N*’, ‘knows’ is to be understood as *not* being a context-sensitive word. If there were danger of confusion then we could use ‘knows*’ instead. But I don’t think there is danger of confusion.

¹⁸ This example is taken from Cohen ([1999], 58-60), where he argues in some detail for the claim that ascriptions of knowledge are context-sensitive. One of his arguments is that since justification (of the appropriate kind) is a component of knowledge, and since ascriptions of justification are context-sensitive (because, like being flat and being rich, justification comes in degrees), ascriptions of knowledge are bound to be context-sensitive as well. Prior ([2001], 97) points out that context also seems to play a role in whether or not we think a certain case is a Gettier case, suggesting that knowledge ascriptions are context-sensitive.

¹⁹ DeRose [1995] contains a good discussion.

²⁰ DeRose offers an account along these lines. See DeRose ([1995], 206-7) and Schiffer ([1996], 323-5) for a nice presentation.

²¹ There are other ways to fill out the details. Cohen [1999] has a good account.

²² One might quibble about the standard to which *S* must know that if *p* then it follows that *q*: perhaps it should be Highest. But I don’t think that we need to quibble here.

Consider, now, the argument of Section 1 that purported to show that Closure, (H), and (B) are incompatible. We can take it to be the argument expressed by the following sentences:

Sentences

- (S1) 'I know that I have hands'
- (S2) 'I do not know that I am not a BIV'
- (S3) 'Closure' (I mean, the sentence used above)
- (S4) 'I know that if I have hands then it follows that I am not a BIV'
- (S5) 'I do not know that I have hands'

Call the propositions expressed by (S1) - (S5), '(P1)' - '(P5)', and the argument expressed 'Argument'. It is natural to think that (P5) follows from (P2), (P3) and (P4), and contradicts (P1), and hence that Argument shows there to be an incompatibility between (P1), (P2) and (P3). But the contextualist would urge us to be more cautious. Allowing that the context might vary between utterances, Sentences express:

Argument

- (P1) I know that I have hands relative to standard *A*
- (P2) It is not the case that I know that I am not a BIV relative to standard *B*
- (P3) Closure
- (P4) I know that if I have hands then it follows that I am not a BIV relative to standard *C*
- (P5) I do not know that I have hands relative to standard *D*

The contextualist points out that for it to be the case that Argument is valid and that (P5) contradicts (P1), it must be that *A*, *B*, *C*, and *D* are the same. For if *B*, *C* and *D* are not the same, then (P5) does not follow from (P2), (P3), and (P4); and if *A* and *D* are not the same, then (P5) does not contradict (P1).

The contextualist now denies that Argument, as expressed by me, shows there to be an incompatibility between (P1), (P2) and (P3), by claiming that Argument, as expressed by me, is invalid. Suppose that there are just two epistemic standards, Low, the standard of our everyday contexts, and High, the standard of our more philosophical contexts. The contextualist claims that when I utter (S1), I am in a context of standard Low, and express the proposition that I know that I have hands relative to standard Low.²³ By the time I utter (S5), however, I am in a context of standard High, and express the proposition that I do not know that I have hands relative to standard High. The argument is invalid, because there has been a change of context and hence epistemic standard along the way. (P5) does not contradict (P1), and Argument does not show that (P1), (P2) and (P3) are incompatible.

Why, then, do I think that the argument I express is valid? Because, the contextualist says, I fail to notice the change in context that renders it invalid, and I mistake the proposition that I *actually* express by uttering (S5), for the proposition that I *would* express were I to utter that sentence in the same context in which I utter (S1). In short, I do not realize that (P5) is *not* the negation of (P1).

This is a nice result: contextualists have an account according to which Closure, the proposition that I express using 'I know that I have hands', and the proposition that I express using 'I do not know that I am not a BIV', can each be true, *and* they have an explanation for why I am led by Argument to mistakenly think that they cannot be. They seem to have just what we want.

²³ Assume that when I start the argument I am not thinking philosophically. If you like, take it that I uttered the first sentence four hours ago in an everyday context.

But there are problems. Here I will just look at one that I think is particularly concerning.²⁴

It has to do with the contextualist account of why I mistakenly think that Argument is valid. According to the contextualist, when I finish Argument in standard High, I mistake the proposition that I did express with ‘I know that I have hands’ in standard Low, for the proposition that I would have expressed with those words in standard High, and I do so because I fail to appreciate (at least explicitly) the context-sensitivity of knowledge sentences, and I am unaware of any relevant shift in context. Schiffer ([1996], 325-8) argues that any semantic account of the context-sensitivity of knowledge sentences will render this ‘error theory’ implausible: competent speakers would be aware of this context-sensitivity, and would not make the mistake that the contextualist claims he does. Cohen ([1999], 79), in reply, argues that, on the contrary, it is no less plausible than claiming that ascriptions of flatness are context-sensitive, even though competent speakers can fail to realize this, and, he argues, this is very plausible.

Although I’m inclined to agree with Schiffer, I will grant that the contextualist’s error theory has some plausibility. What concerns me more is one of its consequences. According to the contextualist, when I am in standard High, I cannot discriminate the proposition that I did express with ‘I know that I have hands’ in standard Low, from the proposition that I would have expressed with those words in standard High. By the same token, I think she must say that when I was in standard Low, I could not discriminate the proposition that I do express with ‘I know that I have hands’ from the proposition that I would have expressed with those words in standard High, for the very act of discrimination would raise the standard to High.

Nothing unusual about that: if I am in a train that I cannot see out of and cannot tell whether or not is moving, then I cannot discriminate the proposition that I express with ‘It is pretty here’ in a London context, from the proposition that I would have expressed with ‘It is pretty here’ in an Oxford context. There is a sense in which, when I utter ‘It is pretty here’, I do not know what I am saying: I do not know whether I am expressing the proposition that it is pretty in London, or the proposition that it is pretty in Oxford, or a host of other propositions. Nonetheless, I realize that which proposition I have expressed depends upon where I am, and I can take steps to find out which proposition that is (I can ask the conductor where I am).

If the contextualist is right, then when I utter ‘I know that I have hands’ when expressing Argument, there is a sense in which I do not know what I am saying: I cannot discriminate the proposition that I have expressed from the proposition that I know that I have hands relative to Low, or the proposition that I know that I have hands relative to High, or a host of other propositions (one for each epistemic standard). But in this case, unlike in the train case, (a) I might not realize that ‘I know that I have hands’ is context-sensitive, and hence might not realize that in a certain sense I do not know what I have said, unless some contextualist suggests to me that it is, and (b) even then I seem to be unable to find out what I have said. For even when the contextualist tries to reassure me that I started the argument in a context of standard Low, and finished the argument in a context of standard High, I do not then feel relieved to learn that I have not contradicted myself. Rather, there is a lingering feeling that I have. If the contextualist is right, and the propositions that I express are not contradictory, then I do not know that they are not, and seem unable to come to know. When I utter, ‘I know that I have hands’, there is a sense in which I do not know, and am in no position to know, what I have said.

²⁴ Cohen ([1999], 77-83) has a good discussion of problems that face the contextualist. One significant one that I will not consider is the availability of so-called ‘warranted assertability maneuvers’. DeRose ([1995], 210-4; [1999], 195-203) has a good discussion of this.

The same goes not just for my utterance of ‘I know that I have hands’, but for a great many other of my utterances of the form ‘I know that p ’: I can always go on to entertain an argument similar to the one above, and from the contextualist response to that argument it follows that in a certain sense I do not know, and cannot come to know, what I said when I uttered that sentence. Being told the contextualist story about ‘ S knows that p ’, is like being told that we are all on that train above, using the sentence ‘It is pretty here’, without realizing that which proposition we thereby express depends upon where the train is, and, in any case, not being in a position to know where the train is. It is difficult not to find this unsettling: the contextualist has an account of Problem that allows to be true all of the relevant knowledge sentences that I take to be true, but in so doing she robs me of knowledge of what those sentences express.²⁵ For this reason at least, I think the contextualist response is unsatisfying.

4. Concluding Remarks

What I hope to have shown is that any epistemologist who takes Closure to be false cannot as yet give a satisfying response to Problem, and nor can any epistemologist who takes Closure to be true. Since Closure is at the heart of Problem, epistemologists ought to keep an open mind about the question whether or not Closure is true. It is an unsolved problem for epistemology, and that, I suggest, is its epistemological significance.

REFERENCES

- Brueckner, A. (1984), ‘Why Nozick is a Sceptic’, *Mind* **93**, pp. 259-64.
- Brueckner, A. (1985), ‘Transmission for Knowledge Not Established’, *Philosophical Quarterly* **35**, pp. 193-5.
- Brueckner, A. (1994), ‘The Structure of the Skeptical Argument’, *Philosophy and Phenomenological Research* **54**, pp. 827-35.
- Cohen, S. (1988), ‘How to be a Fallibilist’, in J. Tomberlin (ed.), *Philosophical Perspectives, Vol. II: Epistemology* (Astacadero: Ridgeview Publishing Company), pp. 91-123.
- Cohen, S. (1991), ‘Skepticism, Relevance, and Relativity’, in B. McLaughlin (ed.), *Dretske and His Critics* (Cambridge, MA: Blackwell), pp. 17-37.
- Cohen, S. (1999), ‘Contextualism, Skepticism, and the Structure of Reasons’, in J. Tomberlin (ed.), *Philosophical Perspectives* **13**, pp. 57-90.
- DeRose, K. (1995), ‘Solving the Skeptical Problem’, *Philosophical Review* **104**, pp. 1-52.
- DeRose, K. (1999), ‘Contextualism: an Explanation and Defense’, in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Malden, MA: Blackwell), pp. 187-205.
- Dretske, F. (1970), ‘Epistemic Operators’, *Journal of Philosophy* **67**, pp. 1007-23.
- Feldman, R. (1999), ‘Contextualism and Skepticism’, in J. Tomberlin (ed.), *Philosophical Perspectives* **13**, pp. 91-114.
- Forbes, G. (1984), ‘Nozick on Skepticism’, *Philosophical Quarterly* **34**, pp. 43-52.
- Lewis, D. K. (1979), ‘Scorekeeping in a Language Game’, *Journal of Philosophical Logic* **8**, pp. 339-59.
- Lewis, D. K. (1996), ‘Elusive Knowledge’, *Australasian Journal of Philosophy* **74**, pp. 549-67.
- Nozick, R. (1981), *Philosophical Explanations* (Cambridge, MA: Harvard University Press).
- Pryor, J. (2001), ‘Highlights of Recent Epistemology’, *British Journal for the Philosophy of Science* **52**, pp. 95-124.
- Putnam, H. (1981), *Reason, Truth and History* (New York: Cambridge University Press).
- Schiffer, S. (1996), ‘Contextualist Solutions to Skepticism’, *Proceedings of the Aristotelian Society* **96**, pp. 317-33.
- Stine, G. C. (1976), ‘Skepticism, Relevant Alternatives, and Deductive Closure’, *Philosophical*

²⁵ In a way, the contextualist response is like Putnam’s, for Putnam’s response can be understood as, in a certain sense, robbing us of knowledge of what we are saying. See Putnam’s ‘Brains in a Vat’ in Putnam [1981].

- Studies* **29**, pp. 249-61.
- Vogel, J. (1987), 'Tracking, Closure, and Inductive Knowledge', in S. Luper-Foy (ed.), *The Possibility of Knowledge: Nozick and His Critics* (Totowa: Rowman and Littlefield), pp. 197-215.
- Vogel, J. (1990), 'Are There Counter-Examples to the Closure Principle?', in M. D. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism* (Dordrecht: Kluwer), pp. 13-27.
- Vogel, J. (1999), 'The New Relevant Alternatives Theory', in J. Tomberlin (ed.), *Philosophical Perspectives* **13**, pp. 57-90.
- Warfield, T. A. (2004), 'When Epistemic Closure Does and Does Not Fail: A Lesson from the History of Epistemology', *Analysis* **64**, To appear.
- Williams, M. (1999), 'Skepticism', in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Malden, MA: Blackwell), pp. 35-69.