

S Knows That P

RAM NETA

University of Utah

I. Why give an analysis?

What's the point of trying to give an account of the necessary and sufficient conditions for a subject *S* to know the truth of some proposition *p* (what I'll henceforth call an "analysis" of "*S* knows that *p*")? Some philosophers might say that an analysis should tell us what knowledge is. This answer may be true enough, as far as it goes, but it doesn't place any informative constraints on an acceptable analysis. For all that this answer tells us, we can give an acceptable analysis of "*S* knows that *p*" by saying that *S* knows that *p* if and only if *S* knows that *p*. Why isn't this analysis good enough?

Obviously, this analysis isn't good enough because it isn't informative. But what kind of information should an analysis provide? How about the following analysis: *S* knows that *p* if and only if it's the case that both *S* knows that *p* and $6762/63 = 107.33$? This analysis is informative, but it doesn't provide the kind of information that we want from an analysis. So what kind of information do we want? Some philosophers will claim that we want an analysis to explain how knowledge supervenes upon some particular base of facts, e.g., the physical facts, the natural facts, the non-evaluative facts, or what have you.¹ But why should we want such an explanation? Typically, the attempt to develop such an explanation is motivated by metaphysical and/or practical considerations.

The metaphysical consideration is this: we don't want to accept a metaphysical theory according to which facts of the non-privileged kind can vary even as facts of the privileged kind remain the same. We want facts of the non-privileged kind to supervene upon facts of the privileged kind. But even assuming that this aim is clear and well motivated,² why should we expect *an analysis* to help us to achieve it? We don't expect analyses of "*x* is a table" or "*x* has a social security number" to explain how being a table or having a social security number supervenes upon the privileged base of facts.³

The practical consideration is this: we want guidance in the application of predicates useful for stating facts of the non-privileged kind. For instance, we want guidance in the application of our epistemic predicates to particular cases. But again, why should we expect an analysis to help us to achieve this aim? We don't expect analyses of "x is a table" or "x has a social security number" to give us guidance in the application of these predicates to particular cases. In what follows, I shall not assume that an analysis of "S knows that p" should satisfy either the metaphysical or practical motivations above. This is not to assume that these motivations are unsatisfiable, or even unsatisfiable by analysis. It is merely to admit that I see no reason to expect an analysis to satisfy these motivations, and so I shall not attempt to offer an analysis that satisfies them.

One thing that we should want from an analysis of "S knows that p" is, I suggest, a solution to puzzles that result when we engage in seemingly plain and plausible knowledge-talk. Such puzzles are constituted by sets of statements that are individually plausible but seemingly inconsistent. Consider four such puzzles.

Closure.

- (A1) I know that I have hands.
- (A2) If I know that p, and I know that p entails q, then I can know that q by deducing q from p.
- (A3) I cannot know that I am not a handless brain-in-a-vat ("BIV").⁴

Defeasibility (External World):

- (B1) I know that, at the present moment (t1), I have hands.
- (B2) I might come at some future time (t2) to acquire good evidence that I had not had hands at t1.
- (B3) Were I to acquire such evidence at t2, then at t2 I could not know (even if I continued to believe) that I had had hands at t1.
- (B4) If S knows that p, and S continues to have all of her current evidence for p and continues to believe that p on the basis of this evidence, then S cannot lose her knowledge that p *merely* by acquiring new evidence.⁵

Fallibility.

- (C1) I know that I have hands.
- (C2) If S knows that p, then S can't be wrong about p.
- (C3) The way in which I formed my beliefs that I have hands is not a perfectly reliable way of forming beliefs; any belief so formed can be wrong.⁶

Underdetermination.

- (D1) I know that I have hands.
- (D2) My belief that I have hands is underdetermined by my evidence.
- (D3) If S's belief that p is underdetermined by her evidence, then, for all S knows, not-p.⁷

II. Why give a contextualist analysis?

A solution to one of these puzzles must do two things: it must explain why each of the statements that constitute the puzzle is plausible, and it must explain how to avoid the apparent inconsistency among those statements.⁸ Contextualist solutions attempt to achieve these two aims by claiming that the verb “to know” is context-sensitive in its semantics, that the claims comprising these puzzles are plausible because true in some contexts, but that there’s no context in which all the claims are simultaneously true. Contextualist solutions enjoy one advantage over non-contextualist solutions: they avoid implausible commitment to the falsity of any of the claims comprised in the puzzle. Although the whole contextualist approach has recently been subject to much critical scrutiny,⁹ I have argued elsewhere that the criticisms that have been levelled against that approach can all be met.¹⁰ For present purposes, I shall simply assume that the contextualist approach can meet these criticisms, and that some form of contextualist solution to these puzzles can succeed. The questions I shall address are: which solution? and how can that solution succeed?

Rieber 1998 proposes an account of “S knows that p” that generates a contextualist solution to Closure. In this paper, I’ll argue that Rieber’s account of “S knows that p” is subject to fatal objections, but we can modify it to achieve an adequate account of “S knows that p” that generates a unified contextualist solution to *all four* puzzles. This is a feat that should matter to those philosophers who have proposed contextualist solutions to Closure:¹¹ all of them have motivated their contextualism by appeal to the fact that they can explain the plausibility of each of the statements in Closure taken individually, and they can do this without having to deny that each of those statements is true, at least in the context in which it is plausible. But notice that this consideration would equally well motivate a contextualist approach to the other puzzles.¹² Nonetheless, no contextualist has yet suggested how a contextualist solution to the other puzzles might go.

Indeed, the most influential contextualists are committed to denying that the other puzzles are susceptible to contextualist solution: Cohen 1988 and Lewis 1996 both claim that the evidentiary standards for knowledge are more or less stringent depending upon the context of epistemic appraisal. On their views, then, there are contexts relative to which the evidentiary standards for knowledge are sufficiently lax that S can know that p even though p is underdetermined by S’s evidence. So Cohen and Lewis are both committed to rejecting (D3). DeRose 1995 and Heller 1999 claim that the reliability standards for knowledge are more or less stringent depending upon the context of epistemic appraisal. On their views, then, there are contexts relative to which the reliability standards for knowledge are sufficiently lax that S can know that p even though S’s belief that p was formed by a method that does not track the truth through all possible worlds. So DeRose and Heller are committed to rejecting (C2), interpreted in such a way as to be incompatible with the conjunction of (C1) and (C3).

III. Rieber's analysis and its benefits

According to Rieber, S knows that p if and only if *the fact that p explains why S believes that p*.¹³ But the explanation in question is implicitly contrastive: to say that the fact that p explains why S believes that p is to say that it explains why S believes that p *rather than* q (for some relevant alternative q).¹⁴ Thus, for Rieber, knowledge-ascription is also implicitly contrastive: to say that S knows that p is to say that S knows that p rather than q. The truth-value of any particular ascription of the form "S knows that p" depends upon the alternatives to p that are relevant in the context of ascription.

Here's a sketch of how this view can be applied to solve Closure: Typically, when I think (A1), I'm in a context in which the relevant alternative to "I have hands" is "I don't have hands". But the fact that I have hands explains why I believe that I have hands (rather than that I don't). So, in this context, I can truthfully claim to know (on Rieber's account) that I have hands. Since (A2) is true, it follows that I can know that I'm not a handless BIV, and so (A3) is false. But typically, when I think (A3), I move into a context in which the relevant alternative to "I have hands" is "I'm a handless BIV". And the fact that I have hands doesn't explain why I believe that I have hands (rather than that I am a handless BIV). So, in this new context, I can truthfully claim not to know (on Rieber's account) that I'm not a handless BIV. Since (A2) is true, it follows that I can't know that I have hands. In the first context, then, (A1) and (A2) are true and (A3) is false. In the second context, (A1) is false, and (A2) and (A3) are true. In no context are (A1), (A2), and (A3) all true, but each is plausible when thought because it is true in the context that one is in when thinking it.

Rieber recommends this solution to Closure over the proposed contextualist solutions of DeRose 1995 and Lewis 1996 on the grounds that the solution that Rieber proposes (unlike those proposed by DeRose or Lewis) is plausible independently of its ability to solve Closure. On Rieber's account, the context-sensitivity of knowledge-ascriptions results from the (already widely acknowledged) context-sensitivity of statements about what explains what.

Rieber's account of "S knows that p" also enjoys a significant advantage over many other accounts: it avoids the Gettier problem. Knowledge is generally thought to be true belief plus some "third condition". As Zagzebski 1999 argues, the only way for an informative account of knowledge to avoid the Gettier problem is to have the satisfaction of the third condition entail the satisfaction of the truth condition.¹⁵ Call any such account "Gettier-proof". (Notice that any non-skeptical Gettier-proof account is externalist, in that it imposes conditions on knowledge the satisfaction of which is not ascertainable by introspection alone. We will return to this point below.) Rieber's account is Gettier-proof, since there is no way for the fact that p to explain S's belief that p unless p is true.

Finally, notice that Rieber is not committed to the controversial doctrine that knowledge *is a kind of belief*.¹⁶ Rieber's account gives the conditions under

which someone knows that something is the case, and one of these conditions is that the knower believes that it is the case, but it doesn't follow that the knowledge and the belief are one and the same state. The knowledge may require the belief without being identical to it.¹⁷

IV. Objections to Rieber's analysis

So there is much to be said on behalf of Rieber's account of "S knows that p". But consider how Rieber deals with the following two objections. First, there is the objection concerning mathematical knowledge. Rieber writes:

Does the fact that $5 + 7 = 12$ explain why I believe that $5 + 7 = 12$? We can at least say this: it is not clear that this fact does *not* explain why I believe it. So mathematical knowledge is at least not a clear counterexample to the analysis. In fact, it seems reasonable to argue as follows: since the analysis is plausible in a wide range of ordinary cases and it provides a non-ad hoc solution to the skeptical puzzle, we ought to conclude that the fact that $5 + 7 = 12$ does after all explain why I believe that $5 + 7 = 12$.¹⁸

Rieber seems uncomfortable offering this response, as well he should be. I propose that we modify Rieber's account to avoid this objection.

Second, there is the objection concerning what we might call "deviant-causal-chain explanation". Here, Rieber writes:

An apple falls on Newton's head. By an amazing coincidence, Newton has a brain lesion such that, whatever rate the apple falls, the impact of the apple will (together with the lesion) cause Newton to believe that objects fall at precisely this rate. As it happens, the apple falls at 32 ft./sec.², and Newton walks away with the belief that objects fall at this rate.

Does Newton know that objects fall at 32 ft./sec.²? It seems that he does not. But does the fact that objects fall at 32 ft./sec.² explain why he believes that they fall at 32 ft./sec.²? Not according to Rieber, who claims that this "leaves out what is perhaps the most salient and important part of the explanation, namely that Newton has a very peculiar brain lesion."¹⁹

Is this a satisfying response to the objection? Let's consider the following case, which Rieber doesn't consider:

An apple falls on Newton's head. By an amazing coincidence, Newton has a brain lesion such that, whatever rate the apple falls, the impact of the apple will (together with the lesion) cause Newton to believe that objects fall at precisely this rate *and that he has a brain lesion*. As it happens,

the apple falls at 32 ft./sec.², and Newton walks away with the belief that objects fall at this rate and that he has a brain lesion.²⁰

Here, Newton's conjunctive belief that objects fall at 32 ft./sec.² and that he has a brain lesion is explained by the fact that objects fall at 32 ft./sec.² and that he has a brain lesion. And yet even here Newton doesn't know that objects fall at 32 ft./sec.², and this is something that he must know if he knows the truth of the conjunction. Rieber cannot protest that we have still left out a "salient and important" part of the explanation: we can pack any such part of the explanation into the content of the explained belief as well. We thereby construct a case in which Newton's belief is caused by the fact believed but doesn't count as knowledge.

V. Designing a correct analysis

Since there is much to be said on behalf of Rieber's account of "S knows that p", let's modify it so that it can avoid the foregoing objections. We can do this by distinguishing different ways to explain why someone believes something. Such explanations may appeal to *reasons* for belief, or they may appeal to the *non-rational causes* of belief.²¹ Now, when we explain why Newton believes that objects fall at 32 ft./sec.² by appeal to the fact that they fall at 32 ft./sec.², we are not explaining Newton's belief by appeal to any reasons for which he holds it. In contrast, when we explain why Jane believes that she has hands, we are explaining Jane's belief by appeal to her reasons, for her perceptions and memories give her reason to believe that she has hands, and (we assume) she believes it for those reasons. Indeed, all the cases that Rieber's account handles plausibly are cases in which the explanation is by appeal to the believer's reasons for belief. So let's register this fact by emending Rieber's account as follows:

S knows that p if and only if the fact that p is *the reason for which* S believes that p.

But this account faces a dilemma. Consider the following example:

An apple falls on Newton's head. By an amazing coincidence, Newton has a brain lesion such that, whatever rate the apple falls, the impact of the apple will (together with the lesion) cause Newton to *experience a great desire* to believe that objects fall at precisely this rate and that he has a brain lesion. Indeed, it will cause him to experience this desire so strongly that he will be willing to pay any amount of money if only he could get himself to form the relevant belief. As it happens, the apple falls at 32 ft./sec.², and Newton walks away wanting more than anything to form the

belief that objects fall at this rate and that he has a brain lesion. So he pays a hypnotist lots of money to get him to form this belief.

In this case, Newton's conjunctive belief that objects fall at 32 ft./sec.² and that he has a brain lesion is explained by the fact that objects fall at 32 ft./sec.² and that he has a brain lesion. And the fact explains his belief by providing him with a reason for that belief. But is the fact itself Newton's reason for holding that belief?

Suppose we answer this question in the affirmative. Then, by the present account, Newton knows that objects fall at 32 ft./sec.². But this is wrong: Newton's belief clearly doesn't suffice for knowledge. And it doesn't suffice for knowledge because Newton has the wrong kind of reason for it—he doesn't have any epistemic reason for it.

Suppose that we answer the preceding question in the negative. Then we should say why it is that, in *this* case, the fact believed isn't Newton's reason for holding the belief, but that, in cases of knowledge, the fact believed is the believer's reason for holding the belief. But our account, in its present form, gives us no guidance as to how to do this.

To avoid this dilemma, we must specify the way in which the fact believed must be the reason for which the believer holds the belief, in order for the belief to be sufficient for knowledge. Here, I have a three-stage proposal.

First, let's specify that knowledge requires not just that the believer have a reason for her belief, but that she have an *epistemic* reason for it, i.e., that she have evidence. It may be objected that not all knowledge requires evidence. Here's an argument, adapted from Lehrer 1965, in support of this objection: Suppose that S has evidence *e* for *p*. Now, if S believes that *p* on the strength of *e*, then S must know that *e* is conclusive evidence for *p*. But is this latter knowledge also supported by evidence? If it is, then S must form her belief that *e* is evidence for *p* on the strength of some other evidence *e'*. But then S must know that *e'* is evidence for the proposition that *e* is evidence for *p*. But is this piece of knowledge also supported by evidence? An infinite regress looms, unless we stop it by saying that some knowledge is not supported by evidence.²²

The problem with this argument is that there are other ways of stopping the regress. One way that we can stop it is by allowing that S can know that *p* on the strength of evidence *e* even if S doesn't know that *e* provides evidence for *p*. Although some philosophers might opt for this way to stop the regress,²³ I'm inclined not to do so, for I do not understand how we can acquire knowledge by appealing to evidence that we don't know to be good evidence.²⁴ To acquire knowledge in such a way would be to acquire knowledge by appeal to evidence that just *happened* to be good evidence, and it's not clear to me how this could be a way of acquiring knowledge. I propose another way to stop the regress above. We can stop it by allowing that there are some mental states which are such that merely being in those states suffices to give

one evidence for propositions of the form: this mental state provides evidence for *p*. On this view, at least some of our mental states wear their evidentiary value on their sleeve. For instance, suppose that you believe that your spouse has been secretly unfaithful. Having this belief seems to be a sufficient condition for having evidence for the following proposition: my belief that my spouse has been secretly unfaithful is evidence that my spouse and I do not communicate well.

Specifying that knowledge requires evidence gives us the following account:

S knows that *p* if and only if the fact that *p* is the evidence on the strength of which S believes that *p*.

But this is still inadequate. In order to know, the believer must form her belief *in a reasonable way* on the basis of her evidence.²⁵ So let's emend our account of knowledge once again:

S knows that *p* if and only if the fact that *p* is the evidence on the strength of which S reasonably believes that *p*.

But this account is puzzling. For S's belief that *p* to be based on evidence is for S to possess some evidence and to believe that *p* on the strength of that evidence. But now isn't S's evidence going to have to be distinct from the fact that *p*? S can't reasonably respond to the question "what evidence do you have that *p*?" by asserting *p*.

Still, in order to have a Gettier-proof account of knowledge, we have to arrange for the satisfaction of the other conditions of our account to entail the satisfaction of the truth condition. Lewis 1996 does this by claiming that the truth of *p* is a logically necessary condition of S's *having* the evidence that she has, and I suggest that we follow him on this point. The fact that S has such evidence—call it "conclusive evidence"—implies the truth of *p*.²⁶ Thus, we offer the following account of knowledge:

(K) S knows that *p* if and only if S reasonably believes that *p* on the strength of S's conclusive evidence for *p*.²⁷

Before proceeding, let me say more about what I mean by "conclusive evidence". I've said that to have a conclusive evidence for *p* is to have evidence for *p* that one couldn't have if *p* weren't true. But it doesn't follow from this that *the propositional content* of one's evidence implies *p*. Suppose that a subject S is in mental state M. The propositional content of M is *q*, and *q* does not entail *p*. Can M be conclusive evidence for *p*? Yes—if *the proposition that S is in M* entails *p*. (I provide examples of such evidentiary states below.)

One can have conclusive evidence that *p* without knowing that *p*, and this may happen in at least three ways. First, one may have conclusive evidence

that p , but nonetheless fail to form the belief that p . David has conclusive evidence that Jane is in the room, but he refuses to believe that Jane is in the room because he's been duped by some mad story to the effect that all of his evidence is misleading. In this case, David has conclusive evidence without knowledge.

Second, one may have conclusive evidence that p , and also believe that p , but not believe that p *on the basis of* that evidence. John has conclusive evidence that Smith is Republican, and he also believes that Smith is Republican. But he holds this belief not on the basis of his conclusive evidence, but rather on the basis of his belief that Smith is wealthy. In this case, John has conclusive evidence, and even belief, but he doesn't know.

Third, one may have conclusive evidence that p , and believe that p , and form the belief that p *on the basis of* one's conclusive evidence, but still fail to know that p , because one fails to form the belief *in a reasonable way* on the basis of one's conclusive evidence. This would be the case if, say, one fails to take into account misleading defeaters in one's possession. Mary has conclusive evidence that Tom Grabit stole the book from the library, and, on the basis of this evidence, she believes that Tom stole the book from the library. But she ignores the testimony of Tom's mother, who falsely claims that Tom was out of town but that his thieving identical twin brother Tim was in town. Here again, there is conclusive evidence, and even belief on the basis of that evidence, without knowledge. For the belief is not reasonably formed: relevant evidence is ignored.

We've tailored (K) to avoid the objection from deviant-causal-chain explanation. But it also avoids the objection from mathematical knowledge: the truth of " $2 + 2 = 4$ " is a necessary condition of someone's *intuiting*, or *proving*, that $2 + 2 = 4$, for "intuit" and "prove" are success verbs. So, according to (K), if someone believes that $2 + 2 = 4$ on the strength of her intuiting it or proving it, then she knows that $2 + 2 = 4$. But she doesn't know it if she could be in the very same evidentiary state without " $2 + 2 = 4$ " being true. It might be worrisome that I must appeal to such counterpossible conditionals in my account of our conclusive evidence for necessary truths, for such conditionals have resisted standard treatments of subjunctive conditionals. But, independently of my account, there must be some way for us to understand counterpossible conditionals—or at least to employ them in our reasoning—if we can understand *reductio ad absurdum* proofs. I take it, then, that my appeal to counterpossible conditionals raises no *new* philosophical problems. So long as we know how to reason on the assumption that $2 + 2 \neq 4$, we can evaluate counterpossible conditionals that have this inequality as their antecedent. And, since it's impossible to have "conclusive evidence" for a false proposition, it follows that, if it were the case that $2 + 2 \neq 4$, then no one could have conclusive evidence that $2 + 2 = 4$.

(K) also avoids the Gettier problem. If Smith could have all the same evidence for the proposition that Jones owns a Ford without its being the case

that Jones owns a Ford, then, by (K), Smith doesn't know that Jones owns a Ford. So when can Smith know that Jones owns a Ford? He can know this when he has conclusive evidence—for instance, when he *sees that* Jones owns a Ford, or when he *remembers that* Jones owns a Ford, or when Jones *shows that* he owns a Ford. Many philosophers would object to the idea that such “success” states as seeing that, remembering that, showing that, and so on, can be part of our evidence set. But if we have to be externalists about knowledge in order to avoid the Gettier problem, then why not also be externalists about evidence?²⁸

Some philosophers might object, on the something like the following grounds, to externalism about evidence. What one epistemically ought to believe depends upon what one has evidence for. And “ought” implies “can”—i.e., if one ought to form one's beliefs on the strength of one's evidence, then it follows that one can so form one's beliefs. But to the extent that one cannot appreciate what evidence one has, one cannot form one's belief on the strength of that evidence. So, if one's evidence is to govern what one ought to believe, it follows that one must be able to appreciate what evidence one has. Evidence, therefore, cannot be externalistically constituted.

The problem with this argument is that the conclusion doesn't follow from the premises. Even if one must be able to appreciate what evidence one has, it doesn't follow that one must be able to ascertain it *by introspection alone*. And that is all that externalism denies. One can appreciate what evidence one has in just the same ways that one can appreciate other contingent facts about the world—by looking and seeing. Indeed, as Goldman 1999 argues, if we insist that all evidence must be fully accessible to introspection, then we will not be able to account for the epistemic value of “stored” beliefs.

It may seem as if (K) implies that inductive knowledge is impossible. Aren't our inductively formed beliefs necessarily based on inconclusive evidence? It may seem as if they are. Consider a case: I see a row of 50 switches and, adjacent to it, a row of lights. I press switch A and the light next to it turns on. Then I press switch B and the light next to it turns on. The same happens with switches C and D. On the basis of these observations, I form the belief that each switch controls the light next to it. Do I have *conclusive* evidence for this belief? That depends on what my trials have shown. If my trials *show* that each switch controls the light next to it, then I do have conclusive evidence: my trials can't *show* what isn't so. But it may seem as if my trials do not show this. But then what do they show? Do they show that switches A, B, C, and D control the lights next to them? Or do they show that switches A, B, C, and D control the lights next to them at the moments at which I pressed those switches? Or do my trials not show anything about what controls what, since the noticed correlations may be accidental? Or do my trials not show even that there are any lights or switches, or any external objects? If we do not allow that my trials show that each switch controls the light next to it, then why should we allow that my trials show anything at all about a world external to my mind? The

apparent difficulty in understanding how inductive knowledge is possible on the present account is of the same kind as the apparent difficulty in understanding how any knowledge of the external world (e.g., my knowledge that I have hands) is possible on that account. In order to resolve the former difficulty then, let's address the latter difficulty, as it emerges in our four puzzles.

To solve our four puzzles, we need to bring out a dimension of context-sensitivity in the notion of evidence employed in (K). What's involved in having evidence? I will not attempt to answer this question in full here, but I put forth the following three claims by way of partial answer:

VI. The context-sensitivity of attributions of evidence

First, for *S* to have evidence for *p* is for *S* to have evidence that favors *p* over some alternative(s) that are relevant in the context of epistemic appraisal.²⁹ To illustrate: Richard is drinking Coke. You and I both ask "Does he have evidence that what he's drinking is Coke?" I'm asking this because I suspect that Richard's oft-stated preference for Coke over Pepsi is groundless, for he can't taste the difference. You, however, are asking because you suspect that Richard has had so little experience with carbonated beverages that he might not be able to taste the difference between Coke and Sprite. In that case, it might be that the answer to my question is "no" while the answer to your question is "yes": Richard can have evidence that what he's drinking is Coke rather than Sprite, even if he doesn't have evidence that what he's drinking is Coke rather than Pepsi.

Second, *S* can have evidence that favors *p* over some relevant alternative *q* only if there is some difference between the way that things would be for *S* if *p* were true and the way that things would be for *S* if *q* were true. The notion of "how things would be for a subject" requires sharpening and will receive it below. But for now, notice that there is some plausibility to this crudely stated idea. For instance, if the way that things would be for Richard if he were drinking Coke is the same as the way that things would be for him if he were drinking Sprite (if, say, his sense of taste was dulled by a cold), then he cannot have evidence that he is drinking Coke rather than Sprite.

Third, *S* can have evidence for *p* without knowing that *p*, without believing that *p*, and without knowing that she has evidence for *p*. For instance, Richard can have evidence that he is drinking Coke rather than Pepsi even if he doesn't attend to this evidence, and so fails to form his beliefs on the basis of it.

We can combine these three ideas by claiming that our concept of evidence is context-sensitive in the following way. We can truthfully claim that *S* has evidence that *p* only if *S* is in some introspectible state that she would be in if *p* were true, but that she would not be in if relevant alternatives to *p* were true. What's available to *S*'s introspection thereby indicates that *p*, rather than any of its relevant alternatives, is the case. We shall understand the

evidentiary difference between the way that things would be for S if p were true and the way that things would be for S if relevant alternatives to p were true in terms of this introspectively available indicator. S needn't notice this introspectively available indicator in order for it to occur in her, and to be an indicator.

By sharpening this crude and intuitive conception of evidence, we can develop a solution to our puzzles, a solution according to which what counts as evidence is relative to a context of attribution of evidence.³⁰ Consider an epistemic subject S. Now, let's say that two hypotheses h1 and h2 are "introspectively indistinguishable for S" just in case:

- (a) If h1 were true, then S would be in mental state M1, and
- (b) If h2 were true, then S would be in mental state M2, and
- (c) Any difference between M1 and M2 is not introspectively available to S.

To say that the difference between two mental states is not introspectively available to S is not to say merely that S can't tell which is which. It is rather to say that there's no difference that S can introspectively detect. Let's say that a hypothesis H is "an uneliminated counterpossibility" with respect to S's knowing that p at t just in case (i) H implies that S doesn't know that p at t and (ii) H and "S knows that p at t" are introspectively indistinguishable for S. An appraiser X "raises" an uneliminated counterpossibility with respect to S's knowing that p at t just in case X (seriously and sincerely) treats that counterpossibility as relevant to the appraisal of S's epistemic state, and relevant by virtue of being an uneliminated counterpossibility.

Now we can state our rule:

- (R) When one raises an hypothesis H that is an uneliminated counterpossibility with respect to S's knowing that p at t, one restricts what counts in one's context of appraisal³¹ as S's body of evidence at t to just those mental states that S has, and would have, at t whether or not H is true.³²

Part of the appeal of (R) is that it captures the flexibility that a concept of evidence *should* have. Our concept of evidence, like our concept of knowledge, should track epistemic properties that matter to us. But precisely what matters to us in ascribing evidence (as in ascribing knowledge) varies across contexts of ascription.

The conjunction of (K) and (R) is our account of knowledge. Like Rieber's account, the present account traces the context-sensitivity of knowledge ascriptions to the context sensitivity of ascriptions of another kind: in this case, evidence ascription. Like Rieber's account, it avoids commitment to the controversial doctrine that knowledge is a form of belief.

VII. How this analysis solves our puzzles

The present account of knowledge can do everything Rieber's account can do and more. For instance, it can solve all four of our puzzles:

Closure. (A2) is true (and so plausible) in all contexts. (A1) is true (and so plausible) in contexts in which no uneliminated counterpossibility with respect to my knowing that I have hands is raised. For in such contexts, my body of evidence includes such "success" states as my *seeing that I have hands*. (A3) is true (and so plausible) in contexts in which the BIV hypothesis is raised, since the BIV hypothesis is an uneliminated counterpossibility with respect to my knowing that I have hands. Other skeptical hypotheses (having to do with dreaming or evil geniuses) might do just as well as the BIV hypothesis here. What matters is that they are uneliminated counterpossibilities with respect to my knowing that I have hands.

Defeasibility. (B3) and (B4) are true (and so plausible) in all contexts. (B1) is true (and so plausible) in contexts in which no uneliminated counterpossibility with respect to my knowing that I have hands is raised. (B2) is true (and so plausible) in contexts in which certain uneliminated counterpossibilities are raised, viz. uneliminated counterpossibilities (such as the BIV hypothesis) that imply that my current evidence is misleading.

It may be objected that (B2) can *never* be plausibly denied, for one can never be certain about future contingent propositions. Consider, for instance, my current knowledge that Tom Grabit stole the book from the library—knowledge that I gained by seeing Tom steal the book from the library. Suppose that I meet Mrs. Grabit tomorrow and she falsely testifies that Tom has been out of town but his identical twin brother John has been in town and is a kleptomaniac. Wouldn't I then lose my knowledge? And how can I rule out the possibility that this will happen tomorrow?

This is the problem captured in what Malcolm 1950 calls "the verification argument", i.e., the problem that one can allegedly never be certain of any empirical claim because one can never verify all of its infinitely many empirical consequences.³³ But relative to ordinary contexts of appraisal, something like Malcolm's response to this argument is correct: part of what is involved in regarding my evidence for *p* as *conclusive* is regarding it as impossible that anything that could happen in the future would count as good evidence against *p*. When I do not regard this as impossible, then I do not regard my present evidence as conclusive. For instance, when I regard it as possible that I could come to acquire good evidence of Tom's innocence, then I do not regard my present evidence of Tom's theft as conclusive. When I regard my present evidence of Tom's theft as conclusive, then either I must regard it as impossi-

ble that Mrs. Grabit will offer the aforestated false testimony, or I must regard such testimony as lacking evidentiary value.

Of course, to regard it as impossible that anything that could happen in the future would count as good evidence against *p* is not to regard it as impossible that anything that could happen in the future would cause me to change my mind about *p*, or about my current evidence for *p*.³⁴ If Mrs. Grabit *were* to offer the aforestated testimony, then I might well come to change my mind about Tom's guilt, and about the conclusiveness of my current evidence for Tom's guilt. And I might recognize that fact about myself even now. But I cannot consistently claim to know that Tom stole the book, allow for the possibility that Mrs. Grabit will offer that testimony, and allow that her offering such testimony would be good evidence of Tom's innocence.

Fallibility. (C2) is true (and so plausible) in all contexts. (C1) is true (and so plausible) in ordinary contexts: in such contexts, it is true that my belief that I have hands is reasonably formed in response to such "success" states as my *seeing that I have hands*. If I see that I have hands, then I can't be wrong, so my belief that I have hands is formed by a method that is perfectly reliable. (C3) is true (and so plausible) in contexts in which certain uneliminated counterpossibilities are raised, viz. counterpossibilities (such as the BIV hypothesis) that imply that my belief is formed by some imperfectly reliable method.

Underdetermination. (D3) is true (and so plausible) in all contexts. (D1) is true (and so plausible) in ordinary contexts: in such contexts, it is true that my evidence that I have hands includes such "success" states as my seeing that I have hands. That evidence does not underdetermine the proposition that I have hands, for I cannot have that evidence unless I have hands. (D2) is true (and so plausible) in contexts in which certain uneliminated counterpossibilities are raised, viz. counterpossibilities (such as the BIV hypothesis) that imply that my evidence fails to rule out some alternative hypothesis.

Thus, we achieve an account of "S knows that *p*" that provides a unified contextualist solution to all four puzzles. It also helps to locate the insight in many alternative accounts of knowledge: some indefeasibility analyses are correct in accepting (B4),³⁵ some reliability analyses are correct in accepting (C2),³⁶ and some conclusive evidence analyses are correct in accepting (D3).³⁷ And accounts that do not accept these general epistemological claims commonly attempt to approximate the effect of accepting them by accepting a weaker version of them. Thus, there are indefeasibilists who accept something like (B4) but weaker.³⁸ There are reliabilists who accept something like (C2) but weaker.³⁹ And there are conclusive evidence theorists who accept something like (D3) but weaker.⁴⁰

Our account also locates the insight in the intuitive idea that, when someone knows that *p*, their being right about *p* is not at all accidental:⁴¹ on the present analysis, their being right about *p* is not at all accidental, for it is a result of their reasonably forming their belief that *p* on the basis of conclusive evidence that *p*. Our account doesn't analyze knowledge completely into non-epistemic terms, but we've already rejected the aspiration to a reductive account.

On the present account of knowledge, *S* inductively knows that *p* just in case *S* believes that *p* on the strength of evidence that *shows that p*. But whether *S*'s evidence shows that *p* depends upon what *S*'s evidence includes. What *S*'s evidence includes is relative to a context of ascription, and so too, therefore, is the possibility of *S*'s possession of inductive knowledge. Herein lies the beginning of a contextualist solution to the problem of induction.⁴²

Notes

¹ Cf. Goldman 1979, Sosa 1980, Chisholm 1990.

² For an ingenious argument to the effect that it is neither clear nor well motivated, see Gemes 1987.

³ These are examples of what Haugeland 1983 calls "weak supervenience". I suggest that, for all we have reason to believe, knowledge weakly supervenes upon whatever base of non-epistemic facts we would wish to regard as privileged.

⁴ One or another form of this puzzle is developed and discussed in Unger 1975, Cohen 1988, DeRose 1995. Usually, the puzzle is stated using a stronger closure principle, e.g., "If I know that *p*, and I know that *p* entails *q*, then I know that *q*." This stronger closure principle is rejected by Dretske 1970, Dretske 1971, and Nozick 1981. But even if we reject the stronger closure principle, we can still develop the Closure puzzle with (A2).

Hawthorne 2000 cleverly constructs a counterexample which would tell against (A2) if it tells against any version of closure: "If I actually know I won't ever perform an instance of modus ponens (a reliable oracle once told me) and I know that if I won't ever perform an instance of modus ponens, then I won't instantiate all the classically valid inference rules, I am not in a position thereby to know the consequent." (Hawthorne 2000, 119) Cohen 2000b responds to this example by retreating to the following weakened closure principle: "If I know *P*, and I know that if *P* then *Q*, and I infer *Q* on this basis, then I know *Q*." (Cohen 2000b, 137) We could still generate a version of the closure puzzle by relying on Cohen's weaker closure principle. But I'd rather respond to Hawthorne by saying that it is impossible for someone to know a proposition of the form "I won't ever perform an instance of modus ponens".

⁵ The problem of "knowing less by knowing more" is developed and discussed in Harman 1973, Ginet 1980.

⁶ This is a generalization of the lottery puzzle, which was discovered by Kyburg 1961, and has been widely discussed since.

⁷ Recent discussions of the role of underdetermination considerations in motivating skepticism include Vogel 1990, Yalçın 1992, Brueckner 1994, Cohen 1998a.

⁸ Cohen 1988 develops these two conditions of adequacy on a solution to Closure.

⁹ In Schiffer 1996, Feldman 1999, Hofweber 1999, Klein 2000, Sosa 2000, Prades 2000, and Rysiew forthcoming. An intuitive objection is levelled, and then defused, in DeRose 1999.

¹⁰ See Neta forthcoming b for a reply to Schiffer 1996, Hofweber 1999 and Rysiew forthcoming, Neta forthcoming a for a reply to Feldman 1999 and Klein 2000, and Cohen 2000b for a reply to Prades 2000.

¹¹ Besides Rieber, this category includes Stewart Cohen (see Cohen 1988), Keith DeRose (see DeRose 1995), David Lewis (see Lewis 1996), and Mark Heller (see Heller 1999).

¹² Specifically, by appeal to the implausibility of flatly denying any of (A1), (A2), or (A3).

¹³ Rieber 1998, 194.

¹⁴ *Ibid.*, 195.

¹⁵ A point similar to Zagzebski's has been made by others. For instance, Almeder 1973 suggests that we avoid the Gettier problem by taking the satisfaction of the justification condition on knowledge to entail the satisfaction of the truth condition. A similar suggestion is made in Tienson 1974.

Craig 1990 argues that, if an analysis of knowledge takes the form:

S knows that p iff S is right as to p and S has X,

then either the second condition entails the satisfaction of the first, or else the conditions are not jointly sufficient.

¹⁶ For controversy surrounding this doctrine, see Vendler 1962, Sayre 1997.

¹⁷ Radford 1966 denies even that knowledge requires belief. He argues for this thesis by appeal to examples, but it is not clear what is shown by the examples to which he appeals.

¹⁸ *Ibid.*, 200–1.

¹⁹ *Ibid.*, 201–2. Rieber also says that our intuitions with regard to this case may be the result of our confusing the falsity of a claim with its misleadingness. But he doesn't give any reason to think that we are likely to confuse falsity with misleadingness in *this particular case*.

²⁰ I will follow Rieber, rather than the custom of physicists, in speaking of the rate of acceleration of a falling object as the rate at which it "falls", rather than the rate at which it "accelerates".

²¹ I shall remain neutral on controversies concerning the distinction between these two varieties of explanation. Useful discussions of the distinction can be found in Anscombe 1960, Sellars 1963a, Davidson 1970, McDowell 1994. For discussions of the relation that obtains between belief and the evidence upon which it is based, see chapters 2 and 3 of Harman 1973, Swain 1979, Pappas 1979a.

²² This argument is adapted from an argument in Lehrer 1965 for the conclusion that there must be some justified belief without evidence.

²³ See, for instance, Alston 1988.

²⁴ Although I do not want to put too much trust in my own intuition on this issue, I do believe, for what it's worth, that this intuition is widely shared. An analogous intuition regarding epistemic justification is captured in what Fumerton 1995 calls "the principle of inferential justification": "To be justified in believing one proposition P on the basis of another proposition E, one must be (1) justified in believing E, and (2) justified in believing that E make probable P." (36). Fumerton argues that the compellingness of skeptical arguments derives from this principle. If he is right, then the principle must have very broad appeal.

Also, an analogous intuition regarding a certain variety of epistemic warrant is captured in what Wright 1991 calls the "Proper Execution Principle": "If the acquisition of warrant to believe a proposition depends on the proper execution of some procedure, then executing the procedure cannot give you any stronger a warrant to believe the proposition in question than you have independently for believing that you have executed the procedure correctly." (99).

²⁵ A compelling argument for this second requirement is presented in chapter 5 of Goldman 1986 and in section VII of Wright 1991.

²⁶ See Dretske 1971. I don't accept Dretske's way of spelling out the relevant conditionals of the form "S couldn't have had this reason for believing that p if p hadn't been true". On my interpretation, the "couldn't have" means logical, rather than physical, impossibility. Many philosophers think that skeptical possibilities such as the BIV possibility show that we cannot have such conclusive reasons for believing any contingent *a posteriori* propositions. But contextualists can avoid this conclusion: in one context, you can truthfully be said to have conclusive reason for believing that you have hands (for you can rule out the relevant alternative that you have claws), though in another context, you cannot truthfully be said to have conclusive reason for believing that you have hands (for you cannot rule out the possibility that you're a BIV).

²⁷ (K) is similar to the account of knowledge proposed in Fogelin 1994. Fogelin, though, is not a contextualist, and does not endorse the other conjunct of the present account of knowledge, (R).

(K) is also similar to the account of knowledge proposed in Ginet 1975. But unlike Ginet, I do not think that the “conclusive evidence” requirement can be analyzed into two independent components—one an internal justification requirement and the other an external “undefeatedness” requirement.

²⁸ I offer a fuller discussion of evidence in Neta Forthcoming a. Williamson 1997 develops an account of evidence according to which being in *any* evidentiary state entails the truth of the state’s propositional content. I don’t accept this consequence of Williamson’s view.

²⁹ See Dretske 1970 and Dretske 1972. Unlike Dretske, I want to preserve the closure of evidence under known entailment, and I do so by claiming that alternatives are relevant or not relative to a context of ascription.

³⁰ Here again, I use the phrase “attributions of evidence” to cover attributions of evidence and their denials, as well as questions concerning whether they are true, suppositions that they are true, and so on.

Williamson 1997 argues that a subject’s evidence consists of all and only what the subject knows. If this argument is sound, then it provides a way to argue from the context-sensitivity of the truth-conditions of knowledge attributions to the context-sensitivity of the truth-conditions of evidence attributions. But it does not provide a way to argue for the specific rule of context-shifting specified below. Williamson himself does not endorse contextualism about either knowledge or evidence.

³¹ One’s listeners may refuse to enter into one’s context of appraisal by refusing to treat H as relevant to the appraisal of S’s epistemic state. One can expect this to happen, for instance, when one raises skeptical hypotheses in a courtroom or a scientific laboratory.

³² If H is an uneliminated counterpossibility with respect to S’s knowing that p at t, then wouldn’t S have *all* the same mental states at t whether or not H is true? Williamson 1995 persuasively defends a negative answer to this question. Also see McDowell 1982 and McDowell 1995. To avoid controversy, I shall not attempt to rest my case on their arguments. Rather, I’ll simply say that I’m using the phrase “mental states” in my statement of (R) in such a way as to include states of knowledge.

³³ A *locus classicus* for the verification argument against the indefeasibility of any of our empirical beliefs is Ayer 1940.

³⁴ Malcolm 1952: “what is the meaning of my assertion that I should regard nothing as evidence that there is no ink-bottle here? That assertion describes my *present* attitude toward the statement that here is an ink-bottle. It does not prophesy what my attitude *would* be if various things happened.”

³⁵ Cf. Lehrer and Paxson 1969, Klein 1971, Klein 1981.

³⁶ Cf. Skyrms 1967, Rozeboom 1967, Dretske 1981, Alston 1983.

³⁷ Cf. Tienson 1974, Fogelin 1994.

³⁸ E.g., Harman 1973, Ginet 1975.

³⁹ E.g., Dretske 1971, Dretske 1981, Goldman 1986, and DeRose 1995.

⁴⁰ E.g., Lewis 1996.

⁴¹ Cf. Unger 1968.

⁴² I am grateful to John Barker, Michael Bishop, Jonathan Cohen, Elijah Millgram, and especially Jonathan Schaffer for helpful comments on an earlier draft of this paper.

References

- Almeder, Robert. 1974. “Truth and Evidence.” *The Philosophical Quarterly* 24: 365–8.
 Alston, William. 1983. “What’s Wrong with Immediate Knowledge?” *Synthese* 55: 73–96.
 _____. 1988. “An Internalist Externalism.” *Synthese* 74: 265–83.
 Anscombe, G.E.M. 1960. *Intention*. Ithaca, NY: Cornell University Press.
 Ayer, A.J. 1940. *The Foundations of Empirical Knowledge*. New York: Macmillan.
 Black, Max. 1950. *Philosophical Analysis*. Ithaca, NY: Cornell University Press.

- Brueckner, Anthony. 1994. "The Structure of the Skeptical Argument." *Philosophy and Phenomenological Research* 54: 827–35.
- Chisholm, Roderick. 1990. "The Nature of Epistemic Principles." *Noûs* 24: 209–15.
- Cohen, Stewart. 1988. "How to be a Fallibilist" in Tomberlin 1988.
- . 1998a. "Two Kinds of Skeptical Argument." *Philosophy and Phenomenological Research* 58: 143–59.
- . 1998b. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery" *Australasian Journal of Philosophy* 76: 289–306.
- . 1999. "Contextualism, Skepticism, and the Structure of Reasons" in Tomberlin 1999.
- . 2000a. "Contextualism and Skepticism" in Sosa and Villanueva 2000.
- . 2000b. "Replies" in Sosa and Villanueva 2000.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon.
- Davidson, Donald. 1970. "Mental Events" in Davidson 1984.
- . 1984. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104: 1–51.
- . 2000. "Ought We to Follow Our Evidence?" *Philosophy and Phenomenological Research* 60: 697–706.
- Dretske, Fred. 1970. "Epistemic Operators." *Journal of Philosophy* 67: 1007–23.
- . 1971. "Conclusive Reasons." *Australasian Journal of Philosophy* 49: 1–22 .
- . 1972. "Contrastive Statements." *Philosophical Review* 81: 411–37.
- . 1981. *Knowledge and the Flow of Information*. Cambridge, MA: The MIT Press.
- Feldman, Richard. 1999. "Skepticism and Contextualism" in Tomberlin 1999.
- Fogelin, Robert. 1994. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press.
- . 2000. "Contextualism and Externalism: Trading in One Form of Skepticism for Another" in Sosa and Villanueva 2000.
- French, P., Uehling, T., and Wettstein, H., eds. 1980. *Midwest Studies in Philosophy V: Epistemology*. Minneapolis: University of Minnesota Press.
- Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman and Littlefield.
- Gemes, Ken. 1987. "The World in Itself: Neither Uniform nor Physical." *Synthese* 73: 301–19.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. Dordrecht: Reidel.
- . 1980. "Knowing Less by Knowing More" in French, Uehling, and Wettstein 1980.
- Goldman, Alvin. 1979. "What is Justified Belief?" in Pappas 1979b.
- . 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- . 1999. "Internalism Exposed." *Journal of Philosophy* 96: 271–93.
- Greco, J. and Sosa, E., eds. 1999. *The Blackwell Guide to Epistemology*. Oxford: Blackwell.
- Harman, Gilbert. 1973. *Thought*. Princeton, NJ: Princeton University Press.
- Haugeland, John. 1983. "Weak Supervenience." *American Philosophical Quarterly* 19: 93–103.
- Hawthorne, J. 2000. "Reply to Cohen" in Sosa and Villanueva 2000.
- Heller, Mark. 1999. "Contextualism and Anti-Luck Epistemology" in Tomberlin 1999.
- Hofweber, Thomas. 1999. "Contextualism and the Meaning-Intention Problem" in Korta, Sosa, and Arrazola 1999.
- Klein, Peter. 1971. "A Proposed Definition of Propositional Knowledge." *Journal of Philosophy* 68: 471–82.
- . 1981. *Certainty: A Refutation of Scepticism*. Minneapolis, MN: University of Minnesota Press.
- . 2000. "Contextualism and the Real Nature of Academic Skepticism" in Sosa and Villanueva 2000.
- Kornblith, Hilary. 2000. "The Contextualist Evasion of Epistemology" in Sosa and Villanueva 2000.
- Korta, K., Sosa, E., and Arrazola, X., eds. 1999. *Cognition, Agency and Rationality: Proceedings of the Fifth International Colloquium on Cognitive Science*. Dordrecht: Kluwer.
- Kyburg, H.E., Jr. 1961. *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.

- Lehrer, Keith. 1965. "Knowledge, Truth and Evidence." *Analysis* 25: 168–75.
- Lehrer, Keith and Paxson, Thomas, Jr. 1969. "Knowledge: Undeclared Justified True Belief." *Journal of Philosophy* 66: 225–37.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549–67.
- Malcolm, Norman. 1950. "The Verification Argument" in Black 1950.
- . 1952. "Knowledge and Belief." *Mind* 61: 178–89.
- McDowell, John. 1982. "Criteria, defeasibility and knowledge." *Proceedings of the British Academy* 68: 455–79.
- . 1994 *Mind and World*. Cambridge, MA: Harvard University Press.
- . 1995. "Knowledge and the Internal." *Philosophy and Phenomenological Research* 55: 877–93.
- Neta, Ram. Forthcoming a. "Contextualism and the Problem of the External World." *Philosophy and Phenomenological Research*.
- . Forthcoming b. "Skepticism, Contextualism, and Semantic Self-Knowledge." *Philosophy and Phenomenological Research*.
- Pappas, George. 1979a. "Basing Relations" in Pappas 1979b.
- , ed. 1979b. *Justification and Knowledge*. Dordrecht: Reidel.
- Prades, Josep. 2000. "Scepticism, Contextualism and Closure" in Sosa and Villanueva 2000.
- Radford, Colin. 1966. "Knowledge—By Examples." *Analysis* 27: 1–11.
- Rieber, Steven. 1998. "Skepticism and Contrastive Explanation." *Noûs* 32: 189–204.
- Rozeboom, William. 1967. "Why I Know So Much More Than You Do." *American Philosophical Quarterly* 4: 281–90.
- Rysiew, Patrick. 2001. "The Context-Sensitivity of Knowledge Attributions." *Noûs* 35: 477–514.
- Sayre, Kenneth. 1997. *Belief and Knowledge: Mapping the Cognitive Landscape*. Lanham, MD: Rowman and Littlefield.
- Schiffer, Stephen. 1996. "Contextualist Solutions to Scepticism." *Proceedings of the Aristotelian Society* 96: 317–33.
- Sellars, Wilfrid. 1963a. "Empiricism and the Philosophy of Mind" in Sellars 1963b.
- . 1963b. *Science, Perception, and Reality*. Atascadero, CA: Ridgeview Press.
- Skyrms, Brian. 1967. "The Explication of 'X knows that p'." *Journal of Philosophy* 64: 373–89.
- Sosa, Ernest. 1980. "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge" in French, Uehling, and Wettstein 1980.
- . 2000. "Skepticism and Contextualism" in Sosa and Villanueva 2000.
- Sosa, E., and Villanueva, E., eds. 2000. *Philosophical Issues, Volume 10: Skepticism*. Boston, MA and Oxford: Basil Blackwell.
- Swain, Marshall. 1979. "Evidence and Justification" in Pappas 1979b.
- Tienson, John. 1974. "On Analyzing Knowledge." *Philosophical Studies* 25: 289–93.
- Tomberlin, J., ed. 1988. *Philosophical Perspectives 2: Epistemology*. Atascadero, CA: Ridgeview.
- , ed. 1999. *Philosophical Perspectives 13: Epistemology*. Atascadero, CA: Ridgeview.
- Unger, Peter. 1968. "An Analysis of Factual Knowledge." *Journal of Philosophy* 65: 157–70.
- . 1975. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- Vendler, Zeno. 1962. *Res Cogitans: A Study in Rational Psychology*. Ithaca, NY: Cornell University Press.
- Vogel, Jonathan. 1990. "Cartesian Skepticism and Inference to the Best Explanation." *Journal of Philosophy* 87: 658–66.
- Williams, Michael. 2000. "Is Contextualism Statable?" in Sosa and Villanueva 2000.
- Williamson, Timothy. 1995. "Is Knowing a State of Mind?" *Mind* 104: 533–65.
- . 1997. "Knowledge as Evidence." *Mind* 106: 717–41.
- Wright, Crispin. 1991. "Scepticism and Dreaming: Imploding the Demon." *Mind* 100: 87–116.
- Yalçın, Ümit. 1992. "Skeptical Arguments from Underdetermination." *Philosophical Studies* 68: 106–22.
- Zagzebski, Linda. 1999. "What is Knowledge?" in Greco and Sosa 1999.