

Sosa, Safety, Sensitivity, and Skeptical Hypotheses

Keith DeRose

Fortunately for those of us who work on the topic, Ernie Sosa has devoted much of his (seemingly inexhaustible) intellectual energy to the problem of philosophical skepticism. And to great effect. With the three exceptions of Peter Unger, whose 1975 Ignorance: A Case for Scepticism is a grossly under-appreciated classic of epistemology; Timothy Williamson, whose 2000 Knowledge and its Limits is, I hope, on its way to being a less underappreciated classic; and Thomas Reid, I have benefitted more from Sosa's wrestlings with skepticism than from anyone else's work on the topic.

Though I am an advocate of a particular kind of "contextualist Moorean" response to skepticism, I still have strong sympathies with "straightforwardly Moorean" responses of the type Sosa favors. If I were forced to abandon the contextualist approach, I would myself adopt a straightforwardly Moorean position. Sosa's work, then, represents an exploration of what, for me, is the path not taken. I am very happy to have such an expert traveler exploring that path so fruitfully.

In two very recent papers — "How to Defeat Opposition to Moore" (HDOM) and "Skepticism and Contextualism" (S&C)¹ — Sosa spends considerable space explicitly comparing his solution to the skeptical problem with contextualist solutions in general, and with mine in particular. I will here continue that conversation. Sosa has not convinced me to change paths, nor do I expect to convince him (though I will claim that in certain ways our approaches are very similar, and perhaps more similar than Sosa thinks). Nevertheless, I welcome the opportunity to compare notes. To properly keep the focus on Sosa's work, I will avoid as much as possible spelling out the details of my own and other alternative approaches, and will endeavor to explain them only so far as is needed to make the comparative points I will argue for.

Key to Sosa's recent efforts (in HDOM and S&C) is his advocacy of a "safety" approach to knowledge and skepticism. Fred Dretske, Robert Nozick, and (in a quite different way) I make use of the concept of the "sensitivity" of beliefs in our response to skepticism. Sosa admits the "undeniable intuitive attractiveness" of the sensitivity approach, but claims to be able to co-opt its benefits by means of his substitute notion of safety, which he claims is easily confused with sensitivity, but which produces much better results when applied to the problem of skepticism (HDOM, p. 143). But it turns out that Sosa, too, uses the notion of sensitivity in his account in a way that renders his account as susceptible to the main problems he raises as is my account — and perhaps as susceptible as are the other sensitivity accounts. Or so I will argue. For this and other reasons, I will dispute Sosa's claim to have produced a superior account.

1. Sensitivity Accounts — Direct and Indirect

A variety of cases elicit from us a strong and surprising intuitive pull toward saying that the subjects of the case don't know the propositions in question. Thus, in the relevant familiar cases, there is a strong pull toward saying that I do not know that

- E1. I've lost the lottery
- E2. My newspaper isn't mistaken about whether the Cubs won yesterday
- E3. Those animals are not just cleverly painted mules
- E4. I'm not a BIV

Sosa points out (HDOM, p. 147), and I agree, that it isn't nearly as intuitively clear as some would make it out to be that there is no knowledge of things like E4 — and I'd say the same about our other Es. (More on this in section 9, below.) Indeed, it is my position that, in an important way, I in fact do know all of the above in the relevant cases.² Nevertheless, as I trust even those who sympathize with that position of mine will agree, there is at least a strong intuitive pull toward the verdict that I don't know each of the above. Of course, there are many

propositions which I intuitively seem not to know. What's surprising about the above? Well, each of the above E's can be paired with another proposition, which we'll in each case label "O" about which there are powerful intuitions to the effect that a) I do know that O and, b) If I don't know that E, then I don't know that O. Consider these Os, which can each be paired with the similarly numbered E, above:

- O1. I won't be able to repay my loan by the end of the year
- O2. The Cubs won yesterday
- O3. Those animals are zebras
- O4. I have hands

In the case of E2/O2, we suppose that my only source of information about the result of the game is my newspaper, which didn't carry a story about the game, but just listed the score under "Yesterday's Results." Intuitively, if the newspaper is a normally reliable one, and, of course, if the Cubs did in fact win, it seems that I know that they won. Yet, in the imagined circumstances (my newspaper is my only source of information about this game), this conditional also seems intuitively correct: If I don't know that my paper isn't mistaken about whether the Cubs won yesterday, then I don't know that they won — if I don't know that E2, then I don't know that O2. These two fairly strong intuitions, if correct, would seem to point to the conclusion that I know that E2. That's why it's surprising that there's such a strong intuitive pull toward saying that I don't know E2. Similar points would apply to our other case pairs. What accounts for this intuitive pull toward saying that the likes of E1-E4 are not known?

It's here that many appeal to the notion of sensitivity. Roughly, a subject S's true belief that p is sensitive iff if p were not the case, then S would not have believed that p. Given the natural understanding of the relevant cases, E1-E4 seem not to be sensitive beliefs, while O1-O4 do seem to be sensitive. (Thus, to continue using the E2/O2 pair, If the Cubs had not won, I would not believe that they had won seems true, while If my paper had been mistaken about yesterday's game, I would not believe it wasn't mistaken does not seem true — it seems that if my paper were mistaken, I'd have believed as strongly as I in fact do that it wasn't.) Sensitivity

explanations appeal to this insensitivity of beliefs E1-E4 to explain why they seem not to constitute knowledge.

The direct way to do this is to follow Dretske and Nozick in supposing that sensitivity is a necessary condition for knowledge. If our concept of knowledge were simply that of true, sensitive belief, it would be no surprise that we tend to judge that insensitive beliefs are not knowledge. And, of course, that point will hold also for more complicated theories of knowledge, so long as they make sensitivity a condition for knowledge.

I also appeal to the insensitivity of E1-E4 to explain why those beliefs can seem not to be pieces of knowledge, but I do not take the above direct approach. Mine is an indirect sensitivity account — one that appeals to the insensitivity of E1-E4 in explaining why they seem not to be knowledge, but does not do so by building a sensitivity condition (or anything like a sensitivity condition) into the very concept of knowledge.

Both direct and indirect sensitivity accounts appeal to the insensitivity of E1-E4 in their explanations of why these beliefs seem not to be knowledge. Both types of accounts then seem to depend on some claim to the effect that we have at least a fairly general — though not necessarily exceptionless — tendency to judge that insensitive beliefs are not knowledge. Without some such assumption, the insensitivity of E1-E4 would not do the explanatory work assigned to it. So both types of account utilize what in “Solving the Skeptical Problem” (SSP),³ I called the “Subjunctive Conditionals Account” (“SCA”) — in the relevant cases, they explain why S seems not to know that p by means of the following two claims:

SCA

1. S’s belief that p is insensitive, and
2. We have some at least fairly general — though perhaps not exceptionless — tendency to judge that insensitive beliefs are not knowledge

Where direct and indirect sensitivity accounts diverge is in their further account of why (2) holds. Direct sensitivity accounts hold that this is so because:

- a. Sensitivity is a necessary condition for knowledge

Indirect sensitivity accounts, then, utilize SCA, but have some explanation other than the one based on (a) for why (2) holds. (In section 4, below, I'll mention a type of sensitivity account — a “modest direct sensitivity account” — that uses a variant of (a) that's nonetheless close enough to (a) that the account should still be labeled “direct”.)

2. The Attack by Counter-Example on Sensitivity Accounts — And Why SCA Seems on the Right Track Nonetheless

As I've noted, Sosa hopes to advance his own safety account as preferable to sensitivity accounts. What's wrong with the sensitivity accounts? Sosa's main attack — at least the main attack that is not limited to targeting only direct sensitivity accounts (Sosa and I seem to be in agreement about what's wrong with direct sensitivity approaches) — is one of counter-examples: he presents cases in which we intuitively judge that a subject knows that p, despite the fact that S's belief that p is, and seems to us to be, insensitive (HDOM, pp. 145-146). If our intuitions about such cases are correct, the cases are counter-examples to the theories of knowledge on which direct sensitivity accounts are based, since they show that (a) is false. But such cases also provide exceptions to the generalization, (2), utilized by even indirect sensitivity accounts, for about such cases, we are not inclined to judge that the subject doesn't know, despite the apparent insensitivity of the subject's belief. Of course, I've formulated (2) so that it is perfectly compatible with there being exceptions to the tendency it posits. (Indeed, I've formulated it in a way that positively anticipates exceptions.) Still, it would be better to explain by means of a generalization that has absolutely no exceptions, so counter-examples like Sosa's are damaging to indirect sensitivity accounts.

Sosa does not present his counter-examples as something new. He is citing an old problem for sensitivity accounts to set the stage for his new, alternative account. Sensitivity theorists have long been aware of counter-examples like the two Sosa presents⁴ — and some other types of cases, too. How have they (we) responded?

For one thing, by suggesting modifications to their accounts. When Nozick first presented his own brand of direct sensitivity theory, he did so along with counter-examples to the simple version — prominently including cases where the subject does know that *p*, despite the fact that she would have believed that *p* even if *p* had been false. As is well-known, Nozick suggested complications to his account involving methods of belief formation to handle the problem cases. In my presentation of an indirect sensitivity account in *SSP*, I discussed several kinds of exceptions to (2), and discussed various ways that generalization might be modified in an attempt to handle such cases (see *SSP*, pp. 19-23). Still, no sensitivity theorist, to my knowledge, has even pretended that all the cases have been successfully dealt with.

But I also argued that though the SCA generalization is not ideally precise, there is good reason to think that it is on the right track, and that it can be used in good explanations. I will repeat the essence of that argument here.

First, and obviously, I pointed out that I was using (2) to explain why we seem not to know in various cases, and the generalization needn't be exceptionless to play that explanatory role. The exceptions perhaps show that the generalization can be refined and improved in certain ways, and may even point us in hopeful directions toward finding some such refinements (some of which are no doubt important and will significantly advance our understanding, and, indeed, some of which I explore), but heaven help us if we have to wait until the generalizations we use in philosophy (or elsewhere) have to be perfectly Chisholmed and exceptionless before we can put them to explanatory work!

But why think the sensitivity account is even on the right track? Why think the exceptions reveal only the need for further tinkering, rather than for a completely different account? Without repeating the case variants I discuss (see *SSP*, pp. 23-27), the reason is that where the account works, it works so impressively well. First, that a subject would have believed *p* even if *p* were false does intuitively seem like a good reason to think the subject

doesn't know that p. And, secondly, and more impressively, when we take cases like the familiar specifications of the situations in which our current Es are usually placed, and then start imagining the most natural ways of modifying the situation in question so that the subject does seem to know the relevant proposition, we will find in an imposingly impressive array of case variants that the very changes needed to make the subject seem to know also render the subject's belief sensitive. As I conclude in SSP, "Again and again, SCA posits a certain block [the insensitivity of the belief] to our judging that we know, and the changes that would clear the way for our judging that we know also remove this block. This makes it difficult not to believe that SCA is at least roughly correct" (p. 25).

Are we to suppose that it's just a coincidence that these Es seem not to be pieces of knowledge when they are in their usual settings, where they are insensitive beliefs, but that they no longer give this "no-knowledge" appearance in the modified situations in which they are sensitive — that the very changes needed to make the appearance of no-knowledge fade away also render the beliefs in question sensitive? Perhaps someone will devise a good explanation, having nothing to do with sensitivity, for why our Es seem not to be knowledge in their usual settings (in which they're insensitive beliefs), and will also allow us to see why they seem to be knowledge in the modified situations (where they are sensitive). Perhaps. But I'm not holding my breath.

3. Sosa's Safety Account

Sosa, though, is not a hit-and-run counter-exampler. He has an alternative account, based on his notion of safety, for why E4 can seem not to be knowledge, and his account seems to be generalizable to cover our other Es as well. Using '-->' for the subjunctive conditional, Sosa explains his notion of safety, and its relation to sensitivity, as follows:

A belief is sensitive iff had it been false, S would not have held it, whereas a belief is safe iff S would not have held it without its being true. For short: S's belief B(p) is sensitive iff $\sim p \rightarrow \sim B(p)$, whereas S's belief is safe iff $B(p) \rightarrow p$. These are not equivalent, since subjunctive conditionals do not contrapose. (HDOM, p. 146)

Sosa claims that safety is (while sensitivity is not) a requirement for knowledge (p. 147). One's belief that one is not a BIV is safe, according to Sosa, and so one does know that one is not a BIV.⁵ The appearance that E4 is not known is an illusion (p. 147). What accounts for this illusion? Here we're asking for an explanation of the same phenomenon that SCA is designed to explain — why E4 can seem not to be knowledge. In short, Sosa's explanation is that the belief that one is not a BIV is an insensitive belief, and, though sensitivity isn't required for knowledge, it is easily confused with safety, which is a requirement for knowledge. In Sosa's own words:

Safety and sensitivity, being mutual contrapositives, are easily confused, so it is easy to confuse the correct requirement of safety...with a requirement of sensitivity. It is easy to overlook that subjunctive conditionals do not contrapose. (p. 148)

Sosa's account can handle our other Es (in their familiar settings) as well: Because, in their familiar settings, beliefs in E1-E3 are insensitive, Sosa's safety account will apply to them as well as it does to E4.

Furthermore, Sosa's account can explain why it so often happens that the very changes to the examples that remove the "no-knowledge" appearance also have the result that the beliefs in question are now sensitive: Once the beliefs in question are rendered sensitive, Sosa's explanation for the appearance of no-knowledge no longer applies to them.

So it can appear that Sosa has indeed produced a non-sensitivity account which can deliver at least many of the main advantages of sensitivity accounts.

4. Sosa's Account as a Sensitivity Account — and His Counter-Examples

But wait! Isn't Sosa's own account a sensitivity account? And isn't it precisely in being a sensitivity account that it is able to deliver those advantages? Sosa's explanation for why E4 can seem not to be known clearly employs the first of SCA's two claims:

1. S's belief that p is insensitive

He combines this with the following in his explanation:

- b. Safety is a necessary condition for knowledge
- c. Because it is the contrapositive of sensitivity, we easily confuse safety with sensitivity

In fact, we can construe Sosa's account as an indirect sensitivity account, which uses (b)-(c) as its explanation for why (2) holds: It's because of (b)-(c) that we so often think that insensitive beliefs are not knowledge.

So far, this is just a matter of classification. It's no surprise that some broad category of explanations can be constructed which includes both Sosa's account and the accounts he is opposing. Do the shared features in virtue of which I'm grouping Sosa's account in with the others bear on the advantages/disadvantages of the accounts?

Yes, it seems. It's in virtue of partaking of the form of SCA that Sosa's account is able to achieve the advantages of the other sensitivity accounts. But, in virtue of that same form, Sosa's account seems to be rendered as vulnerable to his own counter-examples as are the accounts he attacks. At least it's hard to see why Sosa's account, as a sensitivity account, isn't as susceptible to the counter-examples he presents as are other sensitivity accounts. I've noted that Sosa's

account seems to apply to E1-E3 (in their usual settings), as well as to E4, and it should be given credit for being able to explain the appearance of no-knowledge in the cases of those other insensitive beliefs. But Sosa's account would also apply equally well to his own counter-examples. Sosa's counter-examples involve insensitive beliefs that do not give the no-knowledge appearance given by our Es. Since these are insensitive beliefs, Sosa's explanation — built on (1), (b), and (c) — would apply to them as well as to our Es. Sosa's explanation, then, would lead us to expect the appearance of no-knowledge also in the case of his counter-examples, but, of course, no such appearance materializes.

So we should ask Sosa: Why don't we get that appearance of no-knowledge in the case of the counter-examples? Ideally, of course, Sosa will refine his explanation so that it will apply where and only where it should — so that it will explain the appearance of no-knowledge where and only where it occurs. (Perhaps in the case of his particular form of indirect sensitivity account, the refinement will consist in an articulation of the circumstances under which we're not likely to confuse sensitivity with safety.) In the meantime, Sosa can, like me, point out that (2) needn't be exceptionless to do the explanatory work to which it's put, and he can have the same grounds I have for thinking that indirect sensitivity accounts are on the right track. (Though I won't presume that he will endorse that argument that I rest on.)

In short, as far as I can see, Sosa is in the same boat I'm in with respect to the exceptions to (2). And I don't even see that he has any important advantage over direct sensitivity accounts here — provided that those accounts are sufficiently modest. Suppose a direct sensitivity theorist uses the phenomena I've cited as reasons for thinking that SCA is on the right track as support for the view that knowledge is roughly true, sensitive belief (with perhaps some other necessary conditions for knowledge also thrown in, to taste). Such a theorist may think there is an important property of beliefs, **sensitivity***, which is necessary for knowledge, and of which the conditionals in terms of which sensitivity is defined give a good first approximation. Such a modest direct sensitivity theorist will replace (a) with

a'. Sensitivity approximates a necessary condition for knowledge

and will take the counter-examples to (a) as illustrating some of the differences between sensitivity and **sensitivity*** and as pointing to some of the directions one might take in refining one's account of **sensitivity***. (a') provides a foundation for (2), and, as far as I can see, the modest direct sensitivity theorist should be no more embarrassed by her need for approximation here than Sosa should be by the fact that he can only claim that we often confound safety with sensitivity, and yet he uses our (fairly general, but not exceptionless?) tendency to so confound in his account.

Imagine a new character, Immodest Sosa, who gives basically the same account as does our actual Sosa, but who claims that safety correctly and with ideal precision articulates the needed necessary condition for knowledge, and who claims that safety is always — with absolutely no exceptions — confounded with sensitivity, and so predicts that, without exception, all insensitive beliefs will give an appearance of no-knowledge. Our actual, more modest, and more sensible, Sosa will think that Immodest Sosa is on the right track — except for his foolish pretensions to precision. And our actual Sosa, I presume, thinks that there are good reasons for accepting his modest version of the safety account, despite the fact that it is not ideally precise. I don't see why our modest direct sensitivity theorist isn't in the same position as Sosa with respect to the counter-examples Sosa presents.

So, it does seem that, with respect to his own counter-examples, Sosa is in the same boat as me — and also as the direct sensitivity theorist, so long as the direct sensitivity theorist is sensibly modest.

I should add that I don't think this is a reason for Sosa to abandon his theory. The sensitivity boat is a fine boat to be in. As far as I can see, no non-sensitivity account even seriously competes with SCA in terms of successfully explaining the relevant phenomena. But it is to say that Sosa doesn't have the advantage here he thinks he has over the other sensitivity accounts. At least, I can't see why he should be counted as any less damaged by his own examples than I am — or even than the (appropriately modest) direct sensitivity theorist is.

5. Safety and the Problem of True/True Subjunctives

So far, I've only argued that Sosa's account doesn't have the main relative advantage he thinks it has. Are there any relative disadvantages to his safety account?

As we've seen, Sosa proposes safety as a necessary condition for knowledge, where S's belief that p is safe iff $B(p) \rightarrow p$. In perhaps the most straightforward way to convert those symbols into (semi-)English, what safety requires is this: that if the subject had believed that p, p would have been the case. I find such a condition for knowledge very problematic.

In particular, we are faced here with the problem of "true/true subjunctives". Consider: Tom is tall; Tom is a pharmacist; now, true or false:

If Tom had been tall, he would have been a pharmacist?

My main reaction here is that there is something gravely wrong with using the above conditional: "What are you saying — 'If Tom had been tall'?! We've just been told that Tom is tall. Didn't you hear?". The conditional seems to pre-suppose that Tom is not tall, but it has been stipulated in the example that he is. But what kind of supposition is this? Is it a truth-condition of the conditional that its antecedent is false? It isn't clear that what's wrong with the conditional here implies that the conditional is false. It also isn't clear that the conditional isn't false. As far as truth-value goes, nothing is clear here. All that's clear is that the conditional is somehow wrong. Generally, in treating such conditionals, namely subjunctive or counterfactual conditionals that have a true antecedent and a true consequent,⁶ it is recognized that there are no strong intuitions about truth-value here to answer to. The use of the subjunctive or counterfactual form of conditional does in some sense pre-suppose that the antecedent is false, and so, when we encounter one with a true antecedent, we are overwhelmed by a sense of impropriety, but it isn't clear that the wrongness involved is such as to render the conditional false.

The most common treatment of true/true subjunctives holds that they are all true. This view takes the wrongness we sense to be some form of conversational inappropriateness that does not imply falsehood. This standard view is generally arrived at by applying a theory of subjunctive conditionals that is taken to be well-motivated by its handling of other conditionals where there are strong intuitions and accepting that theory's results here, in the realm of true/true subjunctives, where the intuitive situation is considerably more murky. I find that standard view quite attractive, but, partly because I've never seen a convincing case for thinking that wrongness of these conditionals is merely one of conversational inappropriateness, I am also attracted to the rival view that no true/true subjunctives are true. This rival view can claim either that true/true subjunctives are all false or that they are all neither true nor false. (Or, I suppose, one can hold that some of them are false, while others lack truth-value, but I don't see why someone would hold to such a hybrid view.) What's essential to the rival view is that no true/true subjunctives are true.

If either the standard view or the rival account, in either of its forms, of true/true subjunctives is correct, Sosa's safety-based account of knowledge is in trouble. For, presumably, that S believes that p, and that p is true are requirements for S to know that p. Those requirements are standardly assumed, and Sosa doesn't indicate that he is denying them, so I take him to be proposing safety as a necessary condition for knowledge, to go alongside the standard conditions that p is true and that S believes that p, and perhaps some other necessary conditions, in his analysis of S knows that p. We can then use the standard conditions to narrow our focus to only true beliefs, which means that in all the important cases — the ones that satisfy the standard conditions by involving true beliefs — to which we can apply Sosa's condition of safety, the relevant conditional, $B(p) \rightarrow p$, will be a true/true subjunctive.

Now, suppose that the standard view of true/true subjunctives is correct: all true/true subjunctives are true. That renders Sosa's safety condition redundant: Every case that satisfies the standard conditions for knowledge will automatically pass Sosa's safety test as well. Safety disastrously fails to eliminate from the realm of the known any case of true belief.

The rival view of true/true subjunctives is even more fatal for Sosa's theory. If no true/true subjunctives are true, then any belief that satisfies the standard requirements for

knowledge will automatically fail Sosa's safety test. Thus, the rival treatment of true/true subjunctives, together with an account of knowledge which requires safety along with the standard conditions of truth and belief, will imply that nobody knows anything — a result Sosa is at great pains to avoid.

What Sosa needs is a third type of treatment of true/true subjunctives — one on which they're not all true, and on which they're not all false. Only then can safety be a necessary condition for knowledge that's both independent of and consistent with the standard conditions for knowledge. And that such a third type of theory might be correct certainly is not out of the question. Nozick posits, as a fourth condition for knowledge (in addition to sensitivity and the two standard conditions), a conditional, $p \rightarrow B(p)$, which is the reverse of Sosa's safety conditional, and so faces the same problem of true/true subjunctives that confronts Sosa.⁷ Nozick presents an argument for thinking that some true/true subjunctives are true, while others are false.⁸ I'm certainly not convinced by Nozick's case, and I'm not alone. The standard view (that true/true subjunctives are all true), I think, is still the standard view. And, for me, the third type of account that Nozick and Sosa need doesn't even take second place: In terms of how likely I think it is that the theories are correct, I'd put this third type of account behind both the "standard" and the "rival" account. And, of course, even if the third type of account is correct, that just gives Sosa's theory of knowledge a chance of being right. It's important for Sosa, as it is for Nozick, that some true/true subjunctives are true and others false, but he also needs the right ones to be true and the right ones to be false; they must divide into true and false in ways that render the safety account of knowledge sensible.

Even if, against my better judgment, the third type of account of true/true subjunctives is correct, I will be quite unable to check the safety account of knowledge against examples to determine if it is right. For such a task will consist of taking examples where S has a true belief that p, and then asking of them whether it is also true that if S had believed that p, p would have been true, and then checking those results against intuitions about whether S knows that p in the cases. But it's hopeless, at least for me, to try to carry out this process. Wherever S does believe that p, and p is true, the conditional If S had believed that p, p would have been the case will just strike me as weird and wrong; there's just no way that I'll be able to get stable and

discriminating reactions that might match up with intuitions about whether the subject knows in the various cases: “Well, here the [true/true subjunctive] conditional seems true, but here it seems false.” Apparently I’m not alone, given the typical attitude taken toward true/true subjunctives by students of conditionals. Thus, I find accounts of knowledge like Sosa’s (and Nozick’s) that put so much weight on true/true subjunctives highly problematic.

6. Other Formulations of Safety

Now, all of the above is based on phrasing the safety conditional, $B(p) \rightarrow p$, in “...had...would have been...” terms: If S had believed that p, p would have been the case. As I wrote, that strikes me as the most straightforward way to express $B(p) \rightarrow p$ as a subjunctive or counterfactual conditional and as the contrapositive of the sensitivity conditional. But Sosa expresses his safety condition in other ways — ways that don’t induce the “weird and wrong” reaction in cases where S believes that p and p is true. Sosa writes:

Call a belief by S that p “safe” iff: S would believe that p only if it were so that p.
(Alternatively, a belief by S that p is “safe” iff: S would not believe that p without it being the case that p; or, better, iff: as a matter of fact, though perhaps not as a matter of strict necessity, not easily would S believe that p without it being the case that p.) (HDOM, p. 142)

So while I’ve been formulating safety as

S1. If S had believed that p, p would have been the case,

Sosa suggests these formulations:

S2. S would believe that p only if it were so that p

S3. S would not believe that p without it being the case that p

S4. As a matter of fact, though perhaps not as a matter of strict necessity, not easily would S believe that p without it being the case that p.

Suppose Tom correctly believes that I ate breakfast this morning. When I try to evaluate for truth the conditional,

T1. If Tom had believed that I ate breakfast this morning, I would have eaten breakfast this morning,

(or: If Tom had believed that I ate breakfast this morning, it would have been the case that I ate breakfast this morning), I just find it weird and wrong, and there's no hope that in some ways of further specifying the example I'll find the conditional true, while where the details of the case are specified in other ways, I'll find the conditional false. Things are quite different with:

T2. Tom would believe that I ate breakfast this morning only if it were so that I ate breakfast this morning,

T3. Tom would not believe that I ate breakfast this morning without it being the case that I ate breakfast this morning,

and

T4. As a matter of fact, though perhaps not as a matter of strict necessity, not easily would Tom believe that I ate breakfast this morning without it being the case that I ate breakfast this morning.

Here there is hope for discriminating reactions. For each of T2-T4, as opposed to T1, there are some ways of filling in the case (keeping constant that I did in fact eat breakfast and that Tom believes that I did) such that the conditional will seem true, and other ways which will make the conditional seem false. What's more, at least when we check fairly obvious cases of knowledge and fairly obvious cases of no-knowledge, the initial results of the safety account of knowledge, where safety is understood here in terms of T2-T4, as opposed to T1, seem promising. T2-T4 do seem true in the obvious cases that first come to mind when we try to imagine a case where Tom seems clearly to know that I ate breakfast. For example, suppose that Tom himself observed me eating breakfast, he was sure it was me, etc. And things indeed are different in the obvious cases of no-knowledge (despite the presence of true belief) that first come to mind. Suppose, for example, that Lying Larry told Tom 20 things this afternoon, including that I ate breakfast this morning. Tom believed everything Larry told him and has no basis for thinking I ate breakfast other than that Larry told him so. Though Larry thought everything he was telling Tom was false, in fact only 19 of the 20 things Larry told Tom were false, with my having eaten breakfast being the only truth in the bunch. Here Tom seems not to know that I ate breakfast, and T2-T4 seem false.

So where safety is construed in one of the alternative (S2-S4) ways, then, a safety account of knowledge seems promising. However, where safety is so construed, I think its claim to being $B(p) \rightarrow p$, and to being the contrapositive of sensitivity is jeopardized. Much here depends on some very fine points about conditionals, but the sensitivity condition is:

If p had not been the case, S would not have believed that p ,

and it seems that it's S1, and not any of S2-S4, that is the contrapositive of that. And as we've seen, it is essential to Sosa's account of why it seems that we don't know that skeptical hypotheses are false that safety is the contrapositive of sensitivity. S2 and S3 would be better candidates for being the contrapositive of the sensitivity conditional if S2's "would believe" and S3's "would not believe" were changed to "would have believed" and "would not have

believed”, respectively — and perhaps if S2’s “only if it were so” were changed to “only if it had been so” and S3’s “without it being the case” were changed to “without it having been the case”.

But once those quite significant changes are made, the new versions of S2 and S3 would have the same problem that S1 has in being applied to cases where p is the case and S has a true belief that p. S4 is so different from the obvious way of contraposing the sensitivity conditional that its claim to being that contrapositive is extremely tenuous.

7. Safety and Strength of Epistemic Position

As Sosa notes (HDOM, p. 144), my own solution (in SSP) to the skeptical problem emerges from the interplay of two quite different ways of evaluating beliefs. First, there’s the issue of whether the belief is sensitive or not. Second, there’s the matter of how strong an epistemic position the subject is in with respect to the thing believed. Sosa presents his safety condition as something of a substitute for sensitivity — something that can be easily confused with sensitivity, but which does a better job in a theory or picture of knowledge. But when I compare the picture of knowledge that informs Sosa’s treatment with my own, it strikes me that the really interesting comparison is between Sosa’s safety condition, as least in its S2-S4 formulations, and especially in the case of S4, with my notion of strength of epistemic position.

Sosa’s gloss on my notion of strength of epistemic position is good for our current purpose:

One’s epistemic position with respect to P is stronger the more remote are the least remote possibilities wherein one’s belief as to whether p does not match the fact of the matter. (HDOM, p. 144).

Compare that with the S4 version of Sosa’s notion of safety:

S4. As a matter of fact, though perhaps not as a matter of strict necessity, not easily would S believe that p without it being the case that p.

S4-safety seems to be a matter of degree, its formulation prompting the question, “how easily?” — or more precisely (though much less grammatically): “how not easily?”. And the matter of how S4-safe is one’s belief seems very close to the issue of how strong an epistemic position one is in with respect to p. Strength of epistemic position is also a matter of how easily S could have gone wrong with respect to p, this being measured by how remote are the least remote possibilities in which one does go wrong. Applying the same way of measuring “ease” to S4, yields:

One’s belief that p is safer the more remote are the least remote possibilities wherein one believes that p without it being the case that p.

Sosa himself gives such a treatment of his safety in the last sentence of this passage:

Here is the striking result: if we opt for safety as the right requirement then a Moorean stance is defensible, and we avoid skepticism. That is to say, one does satisfy the requirement that one’s belief of not-H be safe: after all, not easily would one believe that not-H (which is not to say that not possibly could one believe that without it being true). In the actual world, and for quite a distance away from the actual world, up to quite remote possible worlds, our belief that we are not radically deceived matches the fact as to whether we are or are not radically deceived. (pp. 146-147)

So S4-safety and strength of epistemic position are very similar notions. The main difference is that S4-safety, if I’m understanding it correctly, is troubled only by worlds (especially very nearby worlds) in which one believes that p but p is not the case (ways in which S would easily believe that p without it being the case that p), while strength is disturbed both by

that and by the presence of worlds (especially very nearby worlds) in which p is true but S disbelieves it. Which notion is more appropriate to an account of knowledge is a very difficult matter, which I won't go into here — though, of course, my preference is for the notion of strength. (Well, let me quickly mention one relative advantage of strength: Its application to beliefs in necessary truths. We can, of course, believe but fail to know necessary truths. The problem for safety here is that there can be no nearby worlds in which we believe a necessary truth, but in which it isn't true, since, being necessary, there are no worlds in which it isn't true. Thus, there can be no worlds to disturb the safety of a belief in a necessary truth. But the strength of such a belief can be upset by the presence of nearby (in the relevant way) worlds in which one disbelieves the necessary truth in question.)

If I'm right that these two notions are quite close, then the picture of knowledge that informs Sosa's treatment of skepticism is very close to my own, for Sosa's rough account of knowledge is that of safe enough true belief, while mine is that of strong enough true belief. (Here, I ease exposition by simply referring to beliefs themselves as "strong," by which I mean that the believer is in a strong epistemic position with respect to the proposition that is the object of the belief.) This may be obscured by Sosa's grouping me together with direct sensitivity theorists. Unlike those direct theorists, I do not picture knowledge as being anything like sensitive true belief. Knowledge for me is (roughly) strong enough true belief, where this notion of strength seems to be quite close to Sosa's notion of safety, at least as Sosa sometimes formulates it. It is one of my main tasks, as an indirect sensitivity theorist, as it is one of Sosa's, to explain why insensitive beliefs so often give the appearance of no-knowledge, even though sensitivity is no part of the concept of knowledge, and even when the beliefs are quite strong. Correctly viewed, I'm no more a sensitivity theorist than is Sosa — and am no more damaged by his counter-examples than he is. We both appeal to certain beliefs' insensitivity to explain why they give an appearance of no-knowledge, but we both do so in an indirect way: We do not picture knowledge as being anything like sensitive true belief, but in other ways (strikingly similar ways, it turns out), and then we each try to explain (in quite different ways) why insensitive beliefs so often seem not to be knowledge, even when they have the real knowledge-making property (safety, strength) to a high degree.

So, Sosa and I are operating with very similar pictures of knowledge — knowledge as (at least roughly) safe/strong enough true belief . I take this picture in a contextualist direction: For me, the matter of how safe/strong a belief must be for it to constitute knowledge is a context-sensitive matter. In some conversational contexts, “S knows that p” can be true only if S is in a very strong epistemic position with respect to p, while in other contexts, S need only be in a moderately strong epistemic position with respect to p for the same knowledge-attributing sentence to be true of her. Sosa is no doubt quite sympathetic to taking our picture of knowledge in such a contextualist direction, since he is not opposed to epistemic contextualist per se — and in fact seems to accept it. But he does seem skeptical about the usefulness of contextualism in addressing traditional epistemological concerns — and, in particular, in providing a solution to the problem of philosophical skepticism.⁹ Here, our disagreement is sharp. I believe that a contextualist version of the picture of knowledge under consideration can provide a powerful explanation of why insensitive beliefs can so often give an appearance of no-knowledge, even when they are quite strong, and can thereby provide a powerful solution to the particular skeptical paradox that both Sosa and I wrestle with, and I take its ability to provide this solution to be a strong reason — though certainly not the only reason — for accepting contextualism.

This is not the place to explain my account of why insensitive beliefs typically seem not to be knowledge. (See SSP, esp. pp. 35-38.) But once one is operating with a picture of knowledge as S4-safe/strong enough true belief, and is inclined to take this picture in a contextualist direction, the materials required by my account seem all in place, and the account I give becomes almost irresistible — at least to me. So, given the problems with the safety account that I outlined above in sections 5-6, and given that the counter-examples Sosa uses against my account seem to give no advantage to his own indirect sensitivity account over my indirect sensitivity account, as I argued in section 4, I’m inclined to stick to my own account, and not accept Sosa’s safety theory — at least in its current form. Of course, there’s plenty of room for Sosa’s safety theory to be developed in various different ways, and there’s no telling at this point how successful those developments will be.

Sosa, I suspect, will weigh up the relative advantages and disadvantages quite differently. I will close by discussing two closely related issues that Sosa discusses and that bear on this

weighing — the attractiveness of contextualist solutions to skepticism, and the question of just how “intuitively correct” it is that we don’t know that skeptical hypotheses are false.

8. Contextualist Solutions to Skepticism

Sosa and I, along with many other current epistemologists, spend much time and energy addressing a particular form of skeptical argument, which I have called the “Argument from Ignorance” (AI). Where ‘O’ is a proposition about the external world that one would ordinarily think one knows (e.g., I have hands) and ‘H’ is a suitably chosen skeptical hypothesis (e.g., I am a bodiless brain in a vat who has been electrochemically stimulated to have those sensory experiences I’ve had, henceforth a ‘BIV’), AI, in its simplest form, proceeds as follows:

1. I don’t know that not-H.
2. If I don’t know that not-H, then I don’t know that O.

So, C. I don’t know that O.¹⁰

AI’s premises are both highly plausible — though more on this in section 9, below. The negation of its conclusion, however, is also highly plausible. AI thus presents us with a puzzle: 1, 2, and not-C (I do know that O) each seems individually plausible, but they can’t all be true. The most attractive straightforward (non-contextualist) options for dealing with this puzzle are (following Sosa’s succinct formulation)¹¹:

Skeptic:	1, 2, C
Nozick, et al.:	1, ~C, ~2
Moore:	2, ~C, ~1 (HDOM, p. 144)

Sosa is a straightforward Moorean, taking the third option above. As a Moorean, he takes his counter-intuitive (but see section 9, below) stand on the issue of 1, claiming, with Moore,¹² that we do indeed know that not-H. As I've urged (see SSP, esp. pp. 3, 42), to successfully follow this path the Moorean must explain why 1 seems true — why it seems that we don't know that not-H. Sosa accepts this requirement (see HDOM, esp. p. 147), and seeks to meet it in the way outlined above in section 3 and critiqued in sections 4-6.

I endorse a non-straightforward, contextualist solution to the puzzle: I accept a contextualist theory of knowledge attributions;¹³ I accept that at some very unusually high standards for knowledge (which we'll here call the "absolute" standards) we don't count as knowing that we have hands; I claim that we do know that we have hands according to the much lower standards for knowledge that typically govern most of our ordinary conversations; and I seek to explain the persuasiveness of the skeptic's argument, at least in part, by claiming that the presentation of the skeptic's argument has at least some tendency to put into play the very "absolute" standards at which we don't count as knowing that we have hands.

Sosa, who is happy enough with contextualist theory of knowledge attributions, in various ways wants a more stridently anti-skeptical solution to our puzzle than contextualist solutions provide. And the contextualist solution is indeed more conciliatory to skepticism than is the straightforward Moorean position — and in the end may be too concessive for Sosa's tastes. But I will here address one of Sosa's worries in a way that may make the contextualist solution a bit more palatable.

Sosa's most extreme worry is that the contextualist solution is simply irrelevant to traditional epistemological reflection on skepticism. Sosa does not claim that it is irrelevant, but is led to wonder what the relevance is (S&C, pp. 3-4). He compares

(a1) People often utter truths when they say "Somebody loves me."

(a2) Does anybody love me?

with

(c1) People often utter truths when they say, “I know there are hands.”

(c2) Do people ever know that there are hands?,

adding: “c2 is presented as a question we might pose in philosophical reflection, in a philosophy journal or conference.” A distraught person led to ask a2 is hoping to be reassured by the answer that indeed somebody does love them; a1 miserably fails to provide the needed comfort. In philosophical reflection on skepticism, in philosophy journals, and in philosophy conferences, many have been moved to ask c2. Why is c1 any more relevant to c2, asked in a philosophical setting, than a1 is to a2?

Let’s assume that philosophical discussions of skepticism are invariably governed by the “absolute” standards according to which nobody knows that there are hands. That’s a highly debatable assumption, but it helps to make the question of relevance more pointed. Under that assumption, when c2 is asked in philosophical discussion of skepticism, the truthful answer to it, according to contextualist solutions to skepticism, is the distressing “no”. Why is c1 any comfort here — any more than a1 is in the distressing situation in which the answer to a2 is the distressing “no”?

The important difference between a1/a2 and c1/c2 is that the relevant “context-sensitivity” in a1/a2 — that “me” refers on each occasion of use to the speaker, and so refers to different people as it is spoken by different people — is clear to all, and so nobody is likely to be misled into thinking that a1 implies that the answer to a2 is positive. By contrast, if the contextualist treatment of skepticism is correct, it is only controversially so, and it is far from being clear to all. Under contextualist analysis, the skeptic raises the standards for knowledge, and our sensitivity to her standards-raising maneuvers inclines us to at least some extent to give negative answers to questions about whether there is knowledge in various cases. At the same time, however, we are sensitive to the fact that it’s usually correct to say that people do know all sorts of things — i.e., we realize that c1 . Not realizing that those usual affirmations are compatible with skeptical denials of knowledge, we get confused.

So far, that sounds quite skeptical. The correct answer to c2, as we’ve been asking it in philosophical discussions of skepticism, is “no”, and our tendency to think otherwise is based on

a confusion — the confusion of thinking that c1 implies that c2 should receive a positive answer. But the confusion goes both ways. Not only does c1 lead us into thinking that c2 should be answered positively, but we can also easily be misled into thinking that a negative answer to c2 implies c1 is false. That, I think, is a large part of why a negative answer to c2 can seem so menacing and so important. At least much of the sting of skepticism comes from the thought that we've been wrong all along in thinking and saying that we "know" various things. If it turns out that almost all our thoughts and assertions to the effect that we "know" various things — including very serious thoughts and assertions made in very serious settings where appropriately high standards (but not yet philosophical, absolute standards) for knowledge hold sway — are really correct, skepticism loses most of its sting, at least for me. Therein lies the comfort of the contextualist response.

Is that sufficient comfort? How much skeptical distress remains behind? That, it seems, depends on how important it is that we know according to "absolute" standards. If one thinks that is all-important, one will find almost no comfort in the contextualist solution. But for my part, once the skeptical strategy is seen to have no tendency to show that any of my claims to know — except those very rare ones made in settings governed by "absolute" standards — are in any way wrong, and once I start to get a clear look at what it would take to "know" according to the skeptic's absolute standards, I find the distress caused by my failure to meet those standards to be minimal at best — perhaps to be compared with the "distress" produced by the realization that I'm not omnipotent.

c1 does not tell us how c2, interpreted according to absolute standards, should be answered — any more than a1 tells one how to answer a2. If traditional epistemological reflection on skepticism concerned only how to answer c2, interpreted absolutely, the contribution of the contextualist solution to that project would be to tell us that the skeptic's answer to the question of traditional inquiry about skepticism is the right one. Any tendency we might have to reject the skeptic's answer that's based on facts like c1 is misguided, according to the contextualist's solution. The contextualist solution might also, in that case, lead us to believe that traditional inquiry into skepticism has not been concerned with a very important question. But insofar as philosophical inquiry into skepticism concerns, not just the question of how c2,

interpreted absolutely, should be answered, but also concerns the importance of that question, so interpreted; insofar as it addresses the relation between philosophical skepticism and ordinary thought about knowledge; and especially insofar as philosophical inquiry into skepticism seeks to address the truly menacing thought that c1 might not be correct after all, the contextualist solution has a lot to offer, and is far from irrelevant.

9. Intuitive Complexity: Do We Know That We're Not Brains in Vats?

Finally, I will close by agreeing with Sosa on an important point, and briefly exploring the implications of that point.

Though I'm not a "straightforward Moorean" like Sosa, I am, like most who advocate contextualist solutions to skepticism, a "contextualist Moorean": I believe that according to ordinary, everyday standards of knowledge, we do know that we're not BIVs.¹⁴ As a Moorean, I, like Sosa, have to explain why it can seem that we don't know that the skeptical hypothesis is false. (In short, my answer is that the unusually high standards at which one doesn't know that the hypothesis is false are the very standards that tend to get put into play when the skeptical hypothesis is brought up. But for more of the story, see SSP.) So, in SSP and now in earlier sections of this paper, I claim that premise 1 of AI is plausible, and focus much effort on explaining why it is so plausible. Sosa insightfully responds:

Consider, moreover, the need to explain how the skeptic's premise — that one does not know oneself not to be radically misled, etc. — is as plausible as it is. That requirement must be balanced by an equally relevant and stringent requirement: namely, that one explain how that premise is as implausible as it is. To many of us it just does not seem so uniformly plausible that one cannot be said correctly to know that one is not at this very moment being fed experiences while envatted. So the explanatory requirement is in fact rather more complex than

might seem at first. And given the distribution of intuitions here, the contextualist and the Nozickean, et al., still owe us an explanation. (HDOM, p. 147)

Though I've found the distribution of intuitions a bit more tilted toward the skeptic-friendly verdict that we don't know we're not BIVs than Sosa's experience leads him to believe,¹⁵ I certainly agree that the intuitive situation is complex, and, in fact, am myself among those to whom it is not "uniformly plausible" that we don't know that we're not BIVs. Since I first encountered philosophical skepticism in the form of AI, I personally have been fairly strongly inclined to think that I do know the various skeptical hypotheses to be false. Nevertheless, I at the same time felt the appeal of the skeptic's claim that we don't know this, and have found that most others feel this appeal more strongly than I do. Upon questioning, I find that almost everyone feels a conflict of intuitive forces pulling them in different directions on the question of whether they do or don't know this about themselves. My focus has been on explaining that strong pull here in the skeptic's favor.

But Sosa asks a very fair — and widely neglected — question: What about the opposing intuitive pull against the skeptic?

According to Sosa, one's belief that one is not a BIV is a safe, but not a sensitive belief. Our concept of knowledge, on Sosa's account is (roughly) that of safe true belief, but we easily confuse safety with sensitivity. Thus, because the belief in question is safe, it will strike us, at least to some extent, as being a case of knowledge, but because it is insensitive, it will, at least to some extent, strike us as not being a case of knowledge.

How might the contextualist meet the challenge Sosa registers at the end of the above indented passage? Here, I'll speak only for myself. I think two aspects of my account that it has in common with Sosa's account provide a good start toward meeting this challenge. First, according to my solution, we are in a very strong epistemic position with respect to I am not a BIV¹⁶ — which is quite close to Sosa's claim that this belief is "quite safe" (HDOM, p. 142). Second, knowledge is, on my account, (roughly) strong enough true belief. These together make our tendency to think that we do know that we're not BIVs quite unsurprising: On my account, as on Sosa's account, one's belief that one is not a BIV has a lot of that property that's needed to

convert true belief into knowledge. Indeed, that one is not a BIV one of the things we know best — though it's peculiarly difficult to truthfully say that we know it. In light of all this, it's the reverse of surprising that we have a tendency to think we do know it, and that we're quite conflicted about the issue.

Notes

1. "How to Defeat Opposition to Moore," Philosophical Perspectives 13 (1999): 141-153; "Skepticism and Contextualism," Philosophical Issues 10 (2000): 1-18.

2. I hold that, according to ordinary standards for knowledge, one does know each of the above. The reason it can seem otherwise is that the extraordinary high standards at which one doesn't know them are precisely the standards that tend to be put into play when those items are brought up.

3. Philosophical Review 104 (1995): pp. 1-52; reprinted in K. DeRose, T. Warfield, ed., Skepticism (New York: Oxford University Press, 1999), and in E. Sosa, J. Kim, ed., Epistemology: An Anthology (Oxford: Blackwell Publishers, 2000).

4. I discuss a counter-example much like Sosa's first in my SSP, pp. 22-23. As Sosa points out, examples much like the one he uses were presented in a 1987 paper by Jonathan Vogel ("Tracking, Closure, and Inductive Knowledge," in S. Luper-Foy, ed., The Possibility of Knowledge (Rowman & Littlefield, 1987).

5. At least this is typically so. I assume Sosa would agree that there could be weird circumstances in which one doesn't know one is not a BIV. Of course, two such weird situations are: 1) where one actually is a BIV and 2) where one is convinced that one is a BIV. But it wouldn't be hard to construct weird situations in which someone isn't a BIV, believes she isn't a BIV, but about which even "Mooreans" would want to say she doesn't know she's not a BIV. Consider, for example, a subject who firmly believed she was a BIV until just a few minutes ago, when Lying Larry told her 20 different things, all of which the subject believed purely on Larry's say-so, where one of the 20 was the truth that the subject is not a BIV, and the other 19 things were all falsehoods.

6. True/False subjunctives — subjunctive conditionals with a true antecedent and a false consequent — also have the problem of having a true antecedent, and thus also seem somehow

wrong. Many people, though, do have a (fairly strong) intuition as to the truth value of true/false subjunctives, intuiting that they are false.

7. In his “Postscript” to “Proper Functionalism and Virtue Epistemology” (in J. Kvanvig ed., Warrant and Contemporary Epistemology, Rowman & Littlefield, 1996; pp. 271-280), Sosa proposes an account of knowledge that takes both $B(p) \rightarrow p$ (Sosa’s safety condition) and $p \rightarrow B(p)$ (Nozick’s fourth condition) to be requirements for knowledge.

8. See Nozick, Philosophical Explanations (Harvard UP, 1981), pp. 176, 680-681 (= Skepticism, pp. 162, 175-176 = Epistemology, pp. 81, 95-96).

9. Sosa opens “S&C” with these words:

Contextualism has gained center stage in epistemology mainly through its way with the skeptic, from the early days of “relevant alternatives” to important recent publications. While myself accepting elements of contextualism, I will detail reservations about its use in epistemology, and in particular about its use to dispose of skepticism. (p. 1)

Sosa writes above that he accepts “elements” of contextualism. What of it does he accept? A bit later in “S&C”, we find:

The main thesis of epistemic contextualism (EC) has considerable plausibility as a thesis in linguistics or in philosophy of language. In applying it to epistemology, however, it is possible to overreach, or so I am here arguing. (p. 3)

I think, then, that what Sosa accepts is what, in the second passage above, he says has “considerable plausibility”: the main contextualist thesis. In that case, it was a bit misleading to describe himself as accepting “elements” of contextualism. He accepts contextualism, not just elements of it, if I’m understanding him correctly. What he doesn’t accept is (at least many of) the applications of contextualism to epistemological problems — and particularly its alleged solution to skepticism.

In earlier work, Sosa had a brighter view of the prospects of applying contextualism to the problem of skepticism; see especially his 1988 paper, “Knowledge in Context, Skepticism in Doubt,” Philosophical Perspectives 2 (1988): 139-155.

10. See SSP, p. 1 and HDOM, p. 143.

11. These are the most attractive straightforward options, because they take the intuitively correct stand on two of the three matters at issue, and taking the intuitively correct stand on all three questions — 1, 2, $\sim C$ — suffers from the serious defect of being inconsistent. Of course, there are other possible, though less attractive, straightforward options on which one takes the intuitively correct position on only one of the three issues, and there’s even the very unattractive option of (perhaps incoherently) violating intuitions on all three issues: ~ 1 , ~ 2 , C .

12. As I note in SSP (p. 41, fn. 37), Moore himself responded in this way — by denying the skeptic’s first premise — to the dream argument — the version of AI that utilizing the hypothesis that one is dreaming as its H. I think it is very far from certain that Moore would

have responded in a similar way to versions of AI that utilize other (especially more radical) skeptical hypotheses. Still, needing a label, we call all deniers of AI's first premise "Mooreans".

13. For an explanation of such theories, see my "Contextualism and Knowledge Attributions," Philosophy and Phenomenological Research 52 (1992): 913-929; and my "Contextualism: An Explanation and Defense," in J. Greco and E. Sosa, ed., The Blackwell Guide to Epistemology (Blackwell Publishers, 1999): 187-205.

14. For more on "straightforward" vs. "contextualist" Mooreanism, see my "How Can We Know That We're Not Brains in Vats?," Southern Journal of Philosophy 38 (2000), Spindel Conference Supplement: 121-148; esp. sect. 6 ("A Nonheroic Alternative: A Moorean Contextualist Account"), pp. 133-136. Stewart Cohen, David Lewis, and Gail Stine, all seem, like me, to be contextualist Mooreans. Mark Heller advocates a different contextualist solution — a contextualist Nozickean solution (this category of solution, as compared with contextualist Mooreanism, is explained in "How Can We Know That We're Not Brains in Vats?," cited just above).

15. In a footnote attached to the above quotation, Sosa reports:

Informal polling of my classes has revealed (of course defeasibly) that those who find [the skeptic's first premise] false outnumber those who find it true, and quite a few prefer to suspend judgment. At every stage people spread out in some such pattern of three-way agreement-failure. (HDOM, p. 152, note 16)

I recently took a poll in a class of over 70 introductory philosophy students, with very different results. My question was whether or not each student, in her or his own opinion, knew that she or he was not a BIV. The preparation was explaining the BIV hypothesis in some detail, but while refraining from using terms of epistemic appraisal in my description of the hypothesis. (So, for instance, I was careful not to say anything like, "So, if you are such a BIV, you can't tell that you are," since I take "can't tell" to be a question-begging negative assessment of any belief one might have to the effect that one is not a BIV.) The question was put to the students before AI was even presented, so I had not indicated that a negative answer might jeopardize their knowledge of various Os — though, no doubt, many students were worrying about that on their own. I asked for a show of hands, first asking those who thought that did not know they weren't BIVs to raise their hands. A clear majority — about 2/3 of the class, it seemed, though I didn't count — of the hands went up. So most of these students agreed with the skeptic on this issue. When I asked who thought they did know they weren't BIVs, only three hands went up. So, as Sosa reports often happens, a sizable portion of the class didn't vote. Still, these are very different results from what Sosa reports. And from what I've found in the past. As best I can tell, the main difference between how I presented the question at this class and how I usually present it is that I explained the BIV hypothesis in more detail than I usually do.

16. On my account, we're not only well enough positioned with respect to I'm not a BIV to meet ordinary standards for knowledge, but also to meet even most of the extraordinarily high standards that are sometimes in play -- though not well-enough positioned to meet the skeptic's

absolute standards, of course. In SSP, I'm quite explicit about our being in as strong an epistemic position with respect to I'm not a BIV as we're in with respect to I have hands. Of course, this surprising comparative fact can be due either to our being in a surprisingly weak position with respect to the latter (as the skeptic would have it), or to our being in a surprisingly strong position with respect to the former. The surprise, on my account, is how well-positioned we are with respect to our not being BIVs. This verdict is defended by means of an account of why it can seem that our knowledge that we're not BIVs is shaky at best. For more on this, and on the importance of defending I'm not a BIV as being something we're in a very strong epistemic position with respect to, see again my "How Can We Know That We're Not Brains in Vats?"