

In *Philosophical Studies* (2003): 35–65.

RELEVANT ALTERNATIVES, CONTEXTUALISM INCLUDED¹

Ernest Sosa

I

Skeptics have traditionally argued as follows.

1. Knowledge of the external world requires perceptual knowledge.
2. If you're to know perceptually that *p* (e.g., that you see a fire), you must base your belief that *p* on sensory experience that discriminates situations in which *p* (in which you see a fire) from <*p*>-precluding alternative situations (in which you see no fire).
3. That being so, if you're to know perceptually that *p* you need tell-tale experience, sensory experience that ensures the truth of <*p*>, sensory experience that would be absent if any <*p*>-precluder were actual.
4. However, in familiar <*p*>-precluding skeptical scenarios (brain in a vat, evil demon, realistic dream, etc.) you might still undergo any and all of your present experience. (Indeed, the skeptic constructs his <*p*>-precluding

¹ Since this paper is for a conference on "Contextualism in Epistemology and Beyond," I have opted to sketch a retrospective of contextualism in epistemology, including highlights of the "relevant alternatives" approach, given how relevantism and contextualism have developed in tandem. We focus on externalist forms of contextualism, bypassing internalist forms such as Cohen 1988 and Lewis 1996, but much of our discussion will be applicable to contextualism generally. Internalist contextualism is helpfully discussed in papers by Stewart Cohen, Richard Feldman, and Jonathan Vogel, in Tomberlin 1999.

scenarios so that if they were so, then not only *might* you still undergo all of your present experience: you *definitely would* still do so.)

5. Therefore, you can attain no perceptual knowledge, and hence no knowledge of the external world.²

According to premise 2, no experience gives an adequate basis for perceptual knowledge unless it enables the knower to discriminate situations where the believed proposition $\langle p \rangle$ is true, from all incompatible alternative situations. The “relevant alternatives” response is to deny that situations in which $\langle p \rangle$ is true must be discriminated by the knower from all $\langle p \rangle$ -precluding alternatives. On the contrary, situations in which $\langle p \rangle$ is true must be discriminated only from those alternatives that are relevant.

And what determines whether an alternative is thus “relevant” (to one’s knowing that p)? Here people part ways, leading to three main positions. On any of the three, *either one defines “relevant alternative” (to $\langle p \rangle$), as “alternative that needs to be ruled out specifically (if one is to know that p)” or else one first gives some definition, or criterion, of relevance, or at least some sufficient conditions, thus preparing the way for a substantive claim that it is *such* alternatives, relevant ones, that need ruling out. Either way, the main questions will be these: (a) What sorts of alternatives are *relevant*? and (b) How does one properly *rule them out*?*

Each of our three positions concerns the requirements for knowing something, or at least for being correctly said to know it, and each will answer both of our two questions, (a) and (b). The first two positions require the ruling out of all relevant $\langle p \rangle$ -precluders. The *first* of these positions, held by Gail Stine, is also in line with writings of Charles Sanders Peirce. It requires, for the relevance of a $\langle p \rangle$ -precluder, that the subject have some reason in its favor (Stine) or at least some positive reason to doubt that it is false (Peirce). According to the *second* position, held by Fred

² I do not distinguish formally between situations or scenarios or states of affairs that are actual and the propositions that fully capture such situations, etc., and are true. The reasoning of interest to us could be cast equivalently either way.

Dretske and favored by Robert Nozick, a $\langle p \rangle$ -precluder is relevant only if it *might* be true even if $\langle p \rangle$ were false (i.e., only if it is *not* one that *would still* be false even if $\langle p \rangle$ were false). As for the *third* position, it asks, not what is required for someone S to know that p, but what is required for an attributor to *say with truth* that S “knows that p.” Because it holds that the truth conditions for knowledge attributions depend on their context, it is a *contextualist* position. Some of my discussion will concern a specific form of this position, one suggested by Keith DeRose’s work— “suggested” because, although he does not appeal explicitly to relevant alternatives, the effect is the same. As we shall see, in order to “know” that p it is still required only that one be in a position to rule out a proper subset of the $\langle p \rangle$ -precluders, a subset whose identity is determined by features of the attributor’s context. Such attributor contextualism is said to enable the best response to the skeptic, by enabling us to respond directly to his arguments, while also explaining the attraction of rival responses. Even where I discuss this specific form, what I say can be generalized to apply to contextualism more generally, both externalist and internalist.

II

One conception of relevance can be found in an early paper by Gail Stine, and a similar idea earlier yet in C.S. Peirce, who writes:

A person may, it is true, in the course of his studies, find reason to doubt what he began by believing; but in that case he doubts because he has a positive reason for it, and not on account of the Cartesian maxim. Let us not pretend to doubt in philosophy what we do not doubt in our hearts. (Peirce 5.2650)

Such passages from Peirce suggest the following condition of relevance, which, though similar to Stine’s, is still different.

CSP A $\langle p \rangle$ -precluding alternative A is not relevant to S’s knowing that

p simply in virtue of being possibly true; on the contrary, if, according to S's background beliefs, a given alternative A is false, then A will be relevant only once S has some positive reason to doubt that it is false.

Compare the Stine condition:

GS A <p>-precluding alternative is relevant to S's knowing that p only if there is some reason to think that it is true.

And suppose we wish, with Stine, to preserve the closure of knowledge over known entailment and deduction, by appeal to the following:

IA All irrelevant alternatives, such as skeptical scenarios, are known to be false without need of reasons or evidence; their sheer irrelevance ensures that their falsity is known to us.³

Can Peirce preserve closure as does Stine, by arguing that irrelevant alternatives are known to be false given their sheer irrelevance? Could he plausibly claim that the sheer irrelevance of skeptical scenarios ensures that we know them to be false? When spelled out, this would amount to claiming that he knows those scenarios to be false simply given his lack of any positive reason to doubt that they are false.

A problem for this derives from yes/no questions that we have no clue how to answer, such as: Are the stars even in number? Here we have no positive reason to say yes, and none to say no. Do we also lack any positive reason to doubt a yes answer, and any positive reason to doubt a no answer? If so, then, according to both CSP + IA and GS + IA, we wind up knowing either answer, whichever we accept, simply given that we accept it (perhaps from time immemorial). But such arbitrariness is unacceptable.⁴ So we must turn to the other possibility, that we *do* after all have a

³ Cf. Stine (1976; in DeRose and Warfield 1999), P. 153: "The point I want to make here is simply that if the negation of a proposition is not a relevant alternative, then I know it—obviously, without needing to provide evidence—and so obviously that it is odd, misleading even, to give utterance to my knowledge."

⁴ Similar points hold good for epistemic justification, it seems to me, though distinctions among sorts of such justification may be required.

positive reason to doubt each answer, the affirmative as well as the negative. What positive reason might that be? Why not the fact that the two answers are equiposed?

Though that seems the best response open to Peirce, it still falls short. Thus, consider the skeptic's alternative <I am a BIV>—or, for short, <BIV>—which precludes <I am sitting>. And compare now <BIV> and <Not-BIV>. Either these are equiposed or they are not. If they *are* equiposed, then I will have positive reason to doubt <Not-BIV>, to doubt the falsity of <BIV>. That being so, <BIV> will after all become a relevant alternative, in which case I do, on the Peircean approach, need some way to rule it out, some way to know that it is false. So I will get to know that <BIV> is false not through its sheer irrelevance, but in some other way. But then it is no longer clear how relevantism helps solve the problem of skepticism. For we then after all do require some way to downgrade the BIV hypothesis, and we will once again face the problem of how we could possibly do that without vicious circularity, and so on and so forth.

If on the other hand <BIV> and <Not-BIV> are *not* equiposed, that would have to be because something favors one over the other. So we do after all need what we were trying to do without, namely some factor that counters the skeptic's scenario, by raising the epistemic status of its falsity over its truth. And this will again force us to deal with skeptical problems that the appeal to relevance might have helped us to avoid, for example the problem of how we can escape vicious circularity while properly downgrading skeptical scenarios.⁵

According to our first sort of relevantism, then, an alternative is relevant only if there is some reason in its favor (Stine), or at least some reason to doubt that it is false (Peirce). However, we have found no defensible position of that sort.⁶

III

⁵ I prefer 'vicious circularity' to 'begging the question' because it better protects the relevant phenomenon from confusion with something involving not so much proper reasoning as proper dialogue. For issues of circularity that threaten once relevance gives no refuge, see Greco 2000, Pryor 2000, and Sosa 1994 and 1997.

⁶ Appendix I examines Stine/Peirce relevantism more closely.

The second position, in line with the writings of Fred Dretske and Robert Nozick, has two parts:

1. Any $\langle p \rangle$ -precluding alternative $\langle a \rangle$ is *relevant* so long as it *might* hold if $\langle p \rangle$ did not. (That is to say, so long as *it is NOT the case that: $\langle a \rangle$ would still be false even if $\langle p \rangle$ were false*).⁷

What now is required if one is to know a given fact? Must one rule out every relevant alternative to that fact, where to rule out an alternative is to know that it is false? This requirement is too strong, since it demands logical powers well beyond us. Any fact has infinitely many relevant alternatives, and the vast bulk of these are not believed false by anyone who believes that fact. Here then is a better way out, closely akin to the Dretske/Nozick line:

2. To rule out a $\langle p \rangle$ -precluding alternative in believing that p , your basis for believing that p must be one that you would not have for this belief if that alternative were true. (In effect, your basis enables you to discriminate the presence from the absence of that alternative.)

This two-part approach answers our two questions. Part 1 specifies which $\langle p \rangle$ -precluding alternatives are relevant, and part 2 tells us what it takes to rule them out. We are thus allowed the familiar claim that in order to know that p one need rule out only $\langle p \rangle$ -precluding alternatives that are relevant.

On this approach we get to keep our ordinary knowledge of hands, fires, zebras, and so on, but must give up our supposed knowledge that we are not radically deceived. Common sense thus splits the difference with skepticism, in a way

⁷ The “might” here is not just the weak “might conceivably,” nor “logically possibly might” (nor is it just a “might” of epistemic possibility). What in our sense “might” be so is *not* simply that whose negation is not necessarily so. Rather, the “might” here is something like “might in fact” or “might easily enough,” or the like. In any case, to say that $\langle a \rangle$ might thus hold if $\langle p \rangle$ did not, is not just to say that the falsity of $\langle p \rangle$ would not entail the falsity of $\langle a \rangle$ with alethic necessity. It is rather to say something a lot stronger, namely, that it is false that $\langle a \rangle$ would in fact still be false, even if $\langle p \rangle$ were false.

that (a) explains the attraction exerted by the skeptic, while it still (b) protects the bulk of our knowledge from his grasp.

Why does the Dretske/Nozick approach insist that in believing yourself illusion-free you must fall short of “knowledge”? That belief of yours is (i) true, (ii) safe (does not just happen accidentally to be true, but *would* be true if held on its present basis)⁸, and (iii) rational (rationally coherent within your overall view). Why exactly then is it not “knowledge”? Because, despite being safe, it is insensitive. That is to say, even if you *were* a victim of such illusion, you might still believe (and in fact *would* still believe), still on your present basis, that you were *not* a victim.

But why require sensitivity rather than safety? Why think that in order to qualify as knowledge a belief must be sensitive and not just safe? In earlier work I have ventured that it might be confusion that prompts this requirement, confusion deriving from how easy it is to assume, falsely, that strong conditionals contrapose, since anyone thus misled would naturally confuse sensitivity with safety. (Sosa 1999 and 2000.)⁹ In my view, we need to distinguish between truth-sensitive and truth-safe belief as follows (restricting ourselves to empirical beliefs):

Bp is *sensitive* iff it is based on an experiential reason $\langle r \rangle$ such that $\sim p \rightarrow \sim r$.

Bp is *safe* iff it is based on an experiential reason $\langle r \rangle$ such that $r \rightarrow p$.

(Where ‘ $x \rightarrow y$ ’ is short for ‘It would be so that x only if it were so that y’. This is to be distinguished from ‘If it had been so that x, then it would have been so that y’, which has unfortunate implicatures or worse.)

We now face the choice of what to require of a belief if it is to qualify as knowledge. Should we say that it must be sensitive, that it must be safe, or that it must be both safe and sensitive? Consider the intuitions that pertain to ‘R is a conclusive reason for P’, when this is defined, with Dretske, as ‘R would not be so unless P were so’, and when it is required that, in order to know P, S must have a conclusive reason for believing P. It is no longer clear whether these intuitions pertain to the

⁸ The idea of “safety” will gain prominence in what follows.

⁹ In an appendix I examine Dretske’s reasoning as a case study of how the confusion can lead us astray.

requirement that the belief must be sensitive, to the requirement that it must be safe, or to the requirement that it must be both sensitive and safe.

Of course those intuitions may remain strong even once we focus exclusively on the sensitivity requirement. So, it may be concluded, the Dretske/Nozick reasoning about knowledge and skepticism is largely unaffected by all these distinctions. How plausible is this?¹⁰

In fact the Nozick tracking account (in both its simple and its improved versions) runs against counterexamples that focus on its sensitivity conditional, counterexamples that hence apply also to Dretske's version of this approach.¹¹ And

¹⁰ Stephen Schiffer has argued that there is a criterion

... implicit in our concept of knowledge according to which I don't know that I'm not a BIV. It might seem that Nozick and, following him, [Keith] DeRose have put their finger on what the implicit criterion is: my belief that I'm not a BIV is insensitive—I'd have that belief even if I were a BIV. But, as DeRose is aware, sensitivity, either as glossed above or even as further precisified by Nozick, can't fully account for the plausibility of the sceptic's first premise, her claim that she doesn't know that she's not a BIV. Further refinement is needed, and it's not clear how it should go. The need for refinement is illustrated by the skeptical hypothesis that I'm a BIV', a BIV' being exactly like a BIV except that it lacks auditory sensations. Now I certainly believe that I'm not a BIV', and my belief that I'm not a BIV' is insensitive: if I were a BIV', I'd still believe that I wasn't. But it's obvious that I *know* that I'm not a BIV', for I know that I'm having auditory sensations and that I wouldn't if I were a BIV'. (Schiffer 1995/6, p. 331.)

It begs the question, however, to argue against those who deny closure, such as Dretske and Nozick, that one knows one is not a BIV' *through knowing that one is having sensations and knowing also that one wouldn't be having sensations if one were a BIV'*. One might as well argue directly that one knows one is not a BIV through knowing that one is sitting and knowing also that one would not be sitting if one were a BIV. For anti-closure thinkers, this is far from showing that one knows one is not a BIV. Nevertheless, even with no question-begging appeal to closure, it does seem plausible enough in any case that one does know one is not a BIV' and that this knowledge is not subject to whatever doubt may attach to the claim that one knows that one is not a BIV. So Schiffer does have a real problem for sensitivity-based solutions to the skeptical paradox, such as those of Dretske, Nozick, and to an extent, that of DeRose as well.

Schiffer takes the problem to remain even for the sensitivity-based approach as "further precisified by Nozick." This alludes to Nozick's modification of his simple tracking account, so as to include reference to the methods used by the subject. The revised version of the tracking account requires not just that if $\sim p$ then the subject S would not believe that p, but rather that if $\sim p$ and S used the same method M in forming a belief as to whether p, then S would not believe that p. This prompts the question of how methods are individuated, and Nozick tells us that "... any method experientially the same, the same "from the inside" will count as the same method." (Nozick 1981, in Luper-Foy 1987, p. 30). And now it is not so clear that the revised account will not deal with the BIV' case. For it is not clear what the truth value of the following conditional might be: If I were a brain in a vat while undergoing no auditory sensations, and I were to base my belief as to whether that was so *on my undergoing auditory sensations*, then I would not believe that I was a brain in a vat while undergoing no auditory sensations. The problem is in part how to evaluate a subjunctive conditional with a contradictory antecedent. Of course this is not to say that Nozick's account clearly gives the right result for that case; it is just to say that it does not clearly give the wrong result. In any case, Schiffer is at bottom right to pose the problem of the BIV', since it anyhow puts in doubt the sensitivity account as interpreted.

¹¹ For one such example, suppose (a) I know I see a hand, while (b) my belief is sensitive: if I saw no hand, and formed a belief about that on my present (perceptual) basis, then I would not believe on that basis that I did see a hand. And suppose I now deduce that I am not wrong in thinking I see a hand. This belief is then formed through a

there are moreover the problematic conjunctions that DeRose rejects as *abominable*, such as the conjunction “... that while you don’t know you’re not a bodiless (and handless!) BIV, still, you know you have hands.” (201) On the Dretske/Nozick account this is a true conjunction, despite which it does seem distinctly repugnant.

DeRose is himself drawn to the sensitivity condition. But he shrinks from its Nozickian consequences. So he takes a contextualist turn, an option that some “relevantists” had already adopted.¹² Contextualist accounts shift the focus from the question of what is required in order to know that *p*, to metalinguistic questions about the conditions required for correctly affirming “knows that *p*” of a subject *S*, conditions that must be satisfied by the attributor *A*.

IV

It would be preferable to avoid abominable conjunctions consequent on the failure of closure, while still somehow using sensitivity to explain how plausible the skeptic can seem. And this is the point of the “Rule of Sensitivity,” namely:

RS When it is asserted that some subject *S* knows (or does not know) some proposition *P*, the standards for knowledge (the standards on how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require *S*’s

certain method *M*, one that combines perceptual experience and inference. This gives rise to a problem for sensitivity-based accounts, because the belief that I reach through that method fails to be sensitive even when we take the method into account. Thus, if I *were* wrong in thinking I see a hand, while forming a belief on that matter through that method, then I would *still* believe that I was *not* wrong in thinking I see a hand. Not only *might* I still believe this under those conditions; in fact I definitely would still believe this under those conditions, if using that same method. Nor does it help to claim that the method here is objectionably one-sided, and yields always the same outcome, the positive belief that one does not wrongly believe that *p*, or the like. For even if we think of the method as that of forming a belief, either affirmative or negative, by basing it on one’s current or past experience, even so if I were wrong in thinking I see a hand, and still used the method of forming a belief about that matter by basing it on my then experiential basis, I would still form the belief that I was *not* wrong in thinking I saw a hand. So, even allowing the two-sided method, *appropriate basing on one’s current or past experience*, we still face a counterexample.

¹² Dretske 1970 (a paper containing both a contextualist and a noncontextualist account, without explicitly distinguishing the two) ; also Goldman 1976, Cohen 1986, and Cohen 1988.

belief in that particular P to be sensitive for it to count as knowledge.
(DeRose, p. 206)

Such contextualism focuses on how the assertion that you know that p ups the strength required in your belief, as well as that required in associated beliefs, if they are respectively to be knowledge. Such beliefs must correspond to the facts of the matter not only in the actual world but in all worlds up to the nearest worlds where it is false that p; and this happens because that knowledge–assertion highlights the question of the sensitivity of your belief that p. But the standards are *not* supposed to be raised similarly by the mere supposition, or even by the thought or the assertion, that your belief is insensitive. There is a reason why it is crucial to the position that the standards be raised only by attributions or denials of knowledge that one is illusion–free.¹³ The mere mention of that radical possibility must not be allowed to raise the standards. Otherwise we put at risk the desired combination of results: namely, (a) to give the skeptic his premise, while (b) still correctly alleging that *in ordinary contexts people know perfectly well that they sit, despite the fact that they are in no better position to know that they sit than to know that they are illusion–free, and despite the fact that were they to believe themselves illusion–free their belief would be insensitive.* In a philosophical context we make that whole italicized claim in adopting the contextualism under scrutiny. For what makes this contextualism distinctively attractive is that it can insist on such a claim even though it also grants the skeptic his skeptical premise, once he asserts it (which is allowed to change the context). It is this that makes it a compromise position able to make reasonable sense of all sides of the controversy.

Unfortunately, it seems arbitrary to so distinguish the conversational or other relevant effects of asserting that one knows oneself illusion–free from the effects of presenting that same possibility of illusion in other, non–assertive ways. But, again, if one allows any presentation of the illusion possibility to raise the standards, then one

¹³ Or by assertion or denial that one is enmeshed in some more or less specific skeptical scenario. DeRose 1995 appeals only to such contextual standards–raisers. I am suggesting that there is a reason to make that a policy, if you wish to use contextualism in dealing with skepticism.

would not be able to make the claim, as we reflect philosophically on the issues of skepticism, that in an ordinary context you can know about fires and hands, even though you're then no better positioned to know about any of that than you are to know yourself illusion-free, and even though believing yourself illusion-free would be insensitive. Unless we protect some such mention of the skeptical possibilities from automatically raising the standards, it is hard to see how the new contextualism could be correctly affirmed, either in speech or in thought.¹⁴

The issue I am raising is not peculiar to one particular form of contextualism. It is an issue for contextualism generally. How serious is it? Not very serious, it may be replied, since the contextualist could just claim in the metalanguage that people sitting in ordinary contexts who say "I know I am sitting" are often enough right. This may be thought to capture enough of what is right in the Moorean stance. But the Moorean stance concerns not what it would be true to say in an ordinary context using the verb 'knows'. It is rather a stance, adopted in a philosophical context, concerning what one then knows and, by extension, what people ordinarily know. At a minimum it is a position about whether people are right in their ordinary claims to know (which is not quite the same as whether they are right in their ordinary utterances of the form "I know that p"). Once we abandon the object language and ascend to the metalanguage, we abandon thereby the Moorean stance.

If the verb 'knows' allowed disquotations, then we could of course retain the Moorean stance even with ascent to the metalanguage. For our claim in the metalanguage that "S says 'I know that p' in an ordinary context and is right," would still entail something in the object language, namely that "S says that he knows that p, and is right." But that would abandon our attributor

¹⁴ By "the new contextualism" here I mean not just a position in linguistics or in philosophy of language about the context-dependent ways in which uses of certain expressions acquire semantic value. I mean rather a position that, while perhaps using such a linguistic theory, still takes a position also in epistemology as to who knows what when, despite the skeptic's arguments, or at least as to who might know what when; and as to what would be sufficient, or necessary, in somewhat general terms, for any such knowledge.

contextualism, which is incompatible with such semantic descent. At least, it is far from clear that the two are consistently combinable. Just compare the move from the premise that A has correctly *uttered* “S is tall” to the conclusion that A has *said* that S is tall. Is *that* move legitimate? An NBA basketball coach complains of Tom Recruit that he “is short.” Has he then *said* that Tom is short and has he thereby spoken truly? Plausibly he has, given Tom’s height of six feet, yet when a passerby in the street considers Tom “not short” he has equally plausibly *said* of Tom that he is not short, and seems also right, with equal plausibility. Contradiction. ¹⁵

In what follows I shall eventually offer another way to make sense of how we can report in the object language on people’s ordinary beliefs that they know they are sitting and the like. More generally, my proposed account tries to explain what is plausible in the various sides of the debate by attributing different respective standards to various attributors of knowledge. We can then report that the ordinary subject believes he knows by ordinary standards that he sits. The skeptic believes the ordinary subject does *not* know by *high* standards that he sits. Each claim is not only plausible but correct. Even if this account is combined with some variety of contextualism, it is not now the contextualism that enables us to explain what is plausible in the various sides of the debate. For, in explaining what is plausible in the Moorean stance we must attribute to ordinary subjects a *de re* belief concerning ordinary standards, and our own abstract, general

¹⁵ According to Stewart Cohen (in Cohen 1999, pp. 79-80), what is troubling and unacceptable about skepticism is the idea that our everyday claims to “know” are generally false. Contextualism attempts to show how the skeptical paradox can be resolved in a way that allows us to preserve the truth of our everyday utterances of the form “S knows that p,” and the like. But that seems to me incomplete. What is troubling about skepticism includes the idea that our ordinary claims to know are generally false, not just the idea that our ordinary utterances of “S knows that p” are generally false. And the semantic descent from the utterance report to the claim report is problematic for attributor-contextualism.

Moreover, it seems to me that the skeptic’s main premise exerts about as powerful an intuitive pull even if we think of it *explicitly* as the claim that *by ordinary standards* we do not know ourselves to be free of radical illusion. I do not myself *agree* with that premise, but I do find it very difficult to say why it is not right, given how plausible it is that we would have the very same experiential basis if we *were* at this moment undergoing such illusion.

endorsement of what ordinary subjects thus believe would then itself presumably amount to a claim concerning ordinary standards *as such*. So our word 'knows' as thus used in a philosophical context amounts to 'knows by ordinary standards'. And this expression is no longer relevantly contextual.

These last three paragraphs make it clear, or so I hope, how my deeper doubt about contextualism takes the form of a dilemma.¹⁶

¹⁶ Alvin Goldman's account of perceptual knowledge is also a major early source of contextualist relevantism, one that also makes a kind of appeal to sensitivity.¹⁶ Here are some key passages.

My theory requires no justification for external-world propositions that derives entirely from self-warranting propositions. It requires only, in effect, that beliefs in the external world be suitably caused, where 'suitably' comprehends a process or mechanism that not only produces true belief in the actual situation, but *would not produce false belief in relevant counterfactual situations*. (p. 63; my italics)

The qualifier 'relevant' plays an important role in my view. If knowledge required the elimination of all logically possible alternatives, there would be no knowledge (at least of contingent truths). If only *relevant* alternatives need to be precluded, however, the scope of knowledge could be substantial. This depends, of course, on which alternatives are relevant. (p. 47)

Goldman shies away from any account that requires an attributor of "knowledge" to have in mind a fully determinate set of alternatives believed relevant, and prefers to say rather this:

A person who *assents* to a knowledge sentence says that S discriminates the truth of *p* from relevant alternatives [where it is S who is said to know that *p*]; but he [the assenter, i.e., the attributor who thus attributes knowledge to subject S] may not have a distinct set of alternatives in mind. (p. 49)

Goldman shares DeRose's indirect appeal to sensitivity, but not his main problem. Earlier we formulated a DeRose-inspired account in terms of relevant alternatives. Unfortunately, its way of ruling on whether an alternative is relevant seemed arbitrary. For it seems arbitrary to suppose that an alternative *does* become relevant if it is *asserted*, in that context, that someone knows it does not obtain, while it does *not* become relevant if that alternative is only *mentioned* or in some other way made salient. By contrast, Goldman does not say what factors make us consider certain alternatives relevant, or make it the case that they are relevant. Indeed, in considering the sorts of factors that make for a judgment of relevance, Goldman fastens on what the attributor has saliently in mind, where such salience does *not* require that an actual knowledge assertion have been made. On the contrary, as seems clear from passages like the following, even just the mention of the problematic alternatives will induce a judgment of relevance in the mind of the attributor:

The speaker's own linguistic and psychological context are also important. If the speaker is in a class where Descartes's evil demon has just been discussed, or Russell's five-minute-old-world hypothesis, he may think of alternatives he would not otherwise think of and will perhaps treat them seriously [i.e., treat them as relevant]. (pp. 48-9)

While thus avoiding arbitrary restriction of the factors that induce relevance, or at least a *judgment* of relevance, this account still raises troublesome questions.

What, for example, does "relevance" mean? This question attaches to the following claim: "A person who *assents* to a knowledge sentence says that S discriminates the truth of *p* from relevant alternatives." (p. 49) For if this claim is right, then in attributing knowledge one is saying something about relevance. In order to have the concept of knowledge, one needs the concept of relevance. It may be replied that the possessor A of the concept of knowledge, he who says or thinks *that S knows that p*, need *not* have the concept of relevance, since it is required only that, in attributing knowledge to S, the attributor A must say *about* alternatives that are in fact relevant, that S discriminates the truth of *p* from *those* alternatives. A need not, in doing so, himself specify those relevant alternatives *as* relevant.

Nevertheless, any such explanation of the concept of knowledge, even if it is offered only as a partial explanation, will be adequate only given an adequately understood concept of relevance. The *explainer* of the concept

V

One important advantage has been claimed by contextualists to favor contextualist accounts: namely, that they enable a good explanation of what makes the skeptic's premise so plausible, of what makes it as plausible as it is that we cannot know ourselves to escape illusion. But this must be weighed against the problems faced by contextualist positions, both externalist sensitivity-based positions, and more internalist forms of contextualism. And we must also consider the prospects for a less costly option with that same advantage. In conclusion I would like to explore such an option.

This better option is enabled once we distinguish sensitivity from safety as a requirement for knowledge. DeRose shows how well the requirement of sensitivity explains our intuitive judgments about quite a few examples. And Goldman had also adduced important considerations favorable to this idea, as with the famous barns example. But the requirement of safety does equally well in explaining our intuitions about those very examples. What is more, if we opt for safety rather than sensitivity, we also apparently avoid abominable conjunctions. At least, we have as yet no sufficient reason to suppose that closure is in any danger once it is required that one's belief be, not sensitive, but only safe, in the sense that one would believe as one does (e.g., that one is sitting) on one's present basis only if one's belief were true.¹⁷

But why should there be any intuitive pull at all in favor of the skeptic's premise that one does not know oneself illusion-free? We now seek an explanation

of knowledge needs the concept of relevance. For the explainer explains that in attributing knowledge to someone, one says that the knower can rule out all *relevant* alternatives. So the explainer himself makes essential use of the concept of relevance. With no satisfactory account of that concept, we would be left wondering just which alternatives, or even which sorts of alternatives, the subject is said to be able to rule out, when the subject is said to know such and such. The present account remains incomplete, therefore, absent an adequate account of relevance. (Goldman 1989 contains relevant further discussion of these matters of relevance.)

¹⁷ A fuller comparison between sensitivity-based contextualism and a safety-based Moorean approach is found in the challenging DeRose forthcoming together with my response.

that comports with the requirement of safety rather than sensitivity. How might we explain such pull under a safety requirement? Believing oneself illusion-free is, after all, extremely safe, and is generally believed to be extremely safe. How then can the skeptic so plausibly deny that we know in so believing? What accounts for the plausibility of his claim if we require only safety and not sensitivity?

Part of the explanation here might be that it is so easy to confuse sensitivity with safety. Moreover, the safety theorist might do well to adopt a moderate contextualism of his own. Since safety comes in degrees, when is a belief *safe enough* to amount to knowledge?¹⁸ As so often elsewhere, here again threshold setting is, plausibly, context-driven. And we can now propose with familiar plausibility that the standards are set in part by what alternatives are salient in the context. Mention of radical $\langle p \rangle$ -precluding alternatives tends to raise the standards (and also to put in play a concept of “knowing for sure” that requires extreme safety). Fortunately, the pressure is resistible in obvious ways. Thus, in the context of a business transaction where money is needed, if someone says that he is “going to the bank for the money,” he will be interpreted as saying that he is going to a financial institution. But he can easily overcome that normal interpretational pressure by saying “I mean the river bank, where I have some buried cash.” Similarly, despite the mention of the skeptic’s alternatives, a Moorean can insist that he means ordinary knowledge, and indeed ordinary *certainty*, thus resisting the pressure to interpret him as meaning knowledge for *absolutely* sure, or *strict* certainty. The resolution of ambiguity, too, is often enough context-driven.

Difference in content seems in any case required once we distinguish among epistemic standards, at least between the ordinary and the philosophical, with

¹⁸ Actually, there is a family of notions of safety, and at least one member comes in degrees. We can thus distinguish

Safety 1: S would believe that p only if it was true that p.

Safety 2: Not easily would S believe that p without being right.

At least the second of these seems clearly to involve some dimension that admits degrees: *how* easily? (Or, perhaps, *how far from* easily?)

contexts to match. If someone in an ordinary context says “I know that p,” how do we in a philosophical context then describe him? We can’t just say “He thinks he knows that p,” without fear of misleading, if indeed our philosophical context imposes high standards for the semantic interpretation of our own speech. We could just say “He uttered the words ‘I know that p,’” but that would report less than we seem in a position to report. Better would be this: “He thinks that by ordinary standards he knows that p.” Of course this is not perfect, since the subject may not even have the notion of “standards,” and less the notion of “standards” that are “ordinary,” and in any case he did not speak of “standards” in saying what he did. So we may be taking some liberties in reporting on his attitudes by inference from his speech. But they seem reasonable liberties of a sort normally taken. Thus if someone has said: “I am here now” and we know who said it, who said it where and when, we can properly report that he thought *he was there then*. Although the subject did not use these words, we can still report his belief by using them ourselves.

Or take someone who says “It is 5 pm,” speaking from New Zealand in a telephone conversation with us in New England. We may then mislead if, without giving our interlocutor’s location, we cover the phone and report “He thinks it’s 5pm.” Better of course would be the report: “He thinks it’s 5 pm where he now is.” This even though the interlocutor did not mention, nor need he have had explicitly in mind, his current location. Indeed, he may not even have the concept of a time zone. Are we then wrong to report his belief as a belief that it is 5 pm in his time zone? The issues here are reminiscent of those concerning our beliefs about water and how that can depend on Earth/Twin Earth distinctions. If the similarity here is more than superficial, then such issues must go beyond cases of natural kinds. Time zones are hardly natural kinds. Nor, perhaps, are epistemic standards, since they may even depend on how vitally important it is to avoid a mistake in the context involved. Perhaps then, here again, in epistemic contexts, we are right to take broad liberties in reporting the thought of someone else by inference from his speech?

More plausibly, it seems to me, we are not really taking liberties. Rather, we are satisfied with reporting *de re* beliefs, leaving unspecified exactly how the subject latches on to the items (natural kinds, time zones, epistemic standards) by reference to which we report on his belief.

In any case, whether we retreat to *de re* reports or not, we arrive at a better epistemic contextualism. Whether a belief is properly considered “knowledge” depends on where it lies on a dimension of safety (among other dimensions). A threshold is required, and this is set by context-determined standards. There is at least a rough distinction between ordinary standards and high standards, or so I am willing to suppose as a simplification harmless for present purposes. If so, then from a philosophical context perhaps the best way to report someone’s relevant thought, by inference from his claim that he “knows he is sitting,” is to say: “He thinks that by ordinary standards he knows he is sitting.”

However, the full truth of the matter is likely more complex. A *lot* of philosophers find it intuitively plausible to say in philosophical contexts that we cannot know ourselves to be illusion-free. Here I include Dretske, Nozick, Cohen, DeRose, and Lewis, to name just a few.¹⁹ Yet, despite my respect for these colleagues, I find it extremely plausible to side with Moore, even in a philosophical context, in supposing that we do know ourselves to be free of radical illusion. Despite the plausibility of linguistic contextualism about epistemic vocabulary, I am led to think that in my relevant dialect, as in Moore’s, the move to philosophical contexts does not bring with it a raising of the standards to the heights alleged. One possibility that I see no way to rule out is that our dialects vary in that regard, that in my own philosophical dialect, as perhaps in Moore’s, there is no such automatic upping of the standards simply through mention of the skeptic’s alternatives. One way to make sense of the apparent disagreement in this area is just to postulate enough difference in dialect to explain the apparently quite robust difference of linguistic intuitions. Why not?

¹⁹ Among relevantists and contextualists, Goldman remains non-committal on this particular issue.

Even if there is such divergence in philosophical dialect, anyhow, we can still communicate in a common language if we are willing to make explicit reference to our relevant standards. We may just agree, for example, that even if by high standards we do not know ourselves to be clear of the skeptic's scenarios, we do know it well enough by ordinary standards. Traditional inquiry as to the conditions for knowledge, for epistemic rational justification, for warrant or aptness, and so on, can now proceed as before, now more explicitly as inquiry about what is required for knowing *by ordinary standards*. What is more, much of the traditional skeptical problematic must still be faced, even if the question we now address is that of how we manage to know *by ordinary standards*. For example, the question of how we can know ourselves to be clear of skeptical scenarios, and in particular of how we can know any such contingent fact without falling into vicious circularity, remains an undiminished threat, even despite our now explicit relativization to ordinary standards. This seems defensible, finally, even if we find considerable variety in ordinary epistemic standards, so long as they all fall under a determinable category of *ordinary epistemic standards*, one free of the heightening that enables skeptics to prove our lack of knowledge by their *very high* standards.²⁰

APPENDIX I: THE RELEVANCE OF PEIRCE

Recall the passage from Peirce according to which only real doubts are relevant, not just Cartesian doubts:

A person may, it is true, in the course of his studies, find reason to doubt what he began by believing; but in that case he doubts because he has a positive reason for it, and not on account of the Cartesian maxim. Let us not pretend to doubt in philosophy what we do not doubt in our hearts. (Peirce 5.2650)

²⁰ For more on our main issues see Klein 2002 and Sosa 2002.

This suggests that, for a $\langle p \rangle$ -precluding alternative A to be relevant to S's knowing that p, it is not enough that it be possibly true; if according to S's background beliefs, A is false, then A will be relevant only once S has some positive reason to doubt that it is false. According to a generalized Peircean condition of this sort, then, an alternative A is relevant for a subject S (at a given time) only if S has some reason in favor of doubting not-A, of doubting that A is false. But how should we understand the idea of doubting a proposition X? There are at least two possibilities. On the weaker of the two, in order to doubt X it is enough to consciously suspend any inclination to believe X, which is of course compatible with one's *also* suspending any inclination to disbelieve X, i.e., to believe Not-X. So that is *unbelief* doubt. On the stronger possibility, that of *disbelief* doubt, in order to doubt X one needs some inclination to disbelieve X, to believe Not-X. Any case of disbelief-doubt is a case of unbelief-doubt, but not every case of unbelief-doubt is a case of disbelief-doubt.

We can now distinguish two Peircean positions.

Strong Peirce A is relevant for S only if S has some reason to disbelief-doubt not-A

Weak Peirce A is relevant for S only if S has some reason to unbelief-doubt not-A

And compare now Gail Stine's position.

Stine A is relevant for S only if S has some reason in favor of inclining to believe A.

Strong Peirce is equivalent to *Stine*, which may be seen as follows. To disbelief-doubt that p one must incline to disbelieve that p. According to *Strong Peirce*, therefore, A is relevant only if one has some reason for inclining to disbelieve not-A; that is to say, only if one has some reason for inclining to believe A. And this is just the *Stine* condition.

Weak Peirce, on the other hand, imposes a weaker condition on relevance than does *Stine* (or *Strong Peirce*). More alternatives are relevant according to *Weak Peirce* than according to *Stine*. Every *Stine*-relevant alternative is *Weak-Peirce*-

relevant. But there are Weak–Peirce–relevant alternatives that are not Stine–relevant.

Suppose now we wish, with *Stine*, to preserve the closure of knowledge over known entailment and deduction, by appeal to the following:

IA All irrelevant alternatives, such as skeptical scenarios, are known to be false with no need of reasons or evidence; their sheer irrelevance ensures that their falsity is known to us.

Any yes/no question that we have no clue how to answer will suffice to refute the combination of *IA* with *Stine* (and therefore its combination with the equivalent *Strong Peirce* condition). For we will then know either answer given simply that we accept that answer, in conditions otherwise the same. But such arbitrariness is epistemically repugnant.

However, that leaves standing the combination of *IA* with *Weak Peirce*, to which we turn next. Now it seems we do have a reason to doubt the yes answer, and *also* a reason to doubt the *no* answer. For we have excellent reason to suspend both assent *and* dissent, and any inclination towards either. On the Weak Peircean stance, therefore, both answers become relevant. Hence *this* Peircean stance is not refuted as was the Strong Peircean stance (and that of *Stine*). It will not follow that we know either of the two equipoised answers given simply that we accept it, and given the sheer irrelevance of its negation. For the equipoised negation of either answer is now relevant and must be ruled out somehow.

But how does this Weak Peirceanism fare with regard to skeptical scenarios such as <I am a BIV> or, for short, <BIV>--which precludes <I am sitting>. Are <BIV> and <Not–BIV> equipoised? If they are, then I will after all need to rule out <BIV> in order to know <I am sitting>. And this will not serve the cause of relevantism. So we must suppose that <BIV> and <Not–BIV> are *not* equipoised. Something must hence favor <Not–BIV> over <BIV>, which means we do after all need what we were trying to do without: namely, some non–viciously–circular way of epistemically raising <Not–BIV>

over <BIV>, something beyond the sheer irrelevance of <BIV>. It may be said that there *is* some reason to incline one towards believing that one is undergoing no radical illusion, namely that one is having vivid experiences as of sitting and writing, etc. But it is hard to see how this could provide a reason for believing that one is not undergoing a perfectly realistic radical illusion as of sitting and writing, etc. The “relevant alternatives” approach was supposed to help us avoid reliance on any such supposed reasons that seem intuitively so implausible.

APPENDIX II: DRETSKE’S ANTI-CLOSURE RELEVANTISM AS A CASE STUDY

According to Dretske’s account, S knows that p only if there is a fact <e> of S’s being in such and such an experiential state, such that (i) S believes <p> on the basis of <e>, and (ii) <e> is a “conclusive reason” for <p>, i.e. <e> would not be so unless <p> were so.²¹ Here is how we are meant to understand the conclusive reasons condition: “Letting the symbol “<>” represent the appropriate modality (a yet to be clarified sense of *possibility* [which turns out to be a kind of relativized physical possibility]), ... R is a conclusive reason for P if and only if, ... $\sim\langle\rangle(R\&\sim P)$.” (Dretske 1971, in Dretske 2000, p. 3)

The basic locution of the account is ‘R would not be so unless P were so’. But this locution is ambiguous. It might mean ‘R would be so only if P were so’ or, alternatively, ‘If P were not so, then R would not be so’. Confusing these goes with symbolizing that locution as ‘ $\sim\langle\rangle(R\&\sim P)$ ’, which is a mistake, as is shown by the fact that

$\sim\langle\rangle(R\&\sim P)$ is equivalent to $\sim\langle\rangle(\sim P\&R)$.

For this in turn entails that ‘R would not be so unless P were so’ contraposes validly, which may be seen as follows:

(i) $\sim\langle\rangle(R\&\sim P)$ is equiv. to *R would not be so unless P were so*.

²¹ Dretske, “Conclusive Reasons,” in Dretske 2000. For simplicity, here we narrow our focus just to knowledge of empirical singular facts.

- (ii) $\sim\langle\rangle(\sim P \& \sim R)$ is equiv. to $\sim P$ would not be so unless $\sim R$ were so.
- (iii) $\sim\langle\rangle(\sim P \& R)$ is equiv. to $\sim\langle\rangle(\sim P \& \sim R)$. Therefore,
- (iv) 'R would not be so unless P were so' contraposes validly.

If one thinks this way, one will then suppose that the ambiguity attributed to 'R would not be so unless P were so' cannot matter, since ' $R \rightarrow P$ ' and ' $\sim P \rightarrow \sim R$ ' are then equivalent anyhow. So, 'R would be so only if P were so' and 'P would fail to be so only if R failed to be so' would be equivalent, and we could then read 'R would not be so unless P were so' either way equivalently.²²

On this basis, once we focus on the ' $\sim P \rightarrow \sim R$ ' reading, it will be easy to suppose that knowledge is not closed under entailment or deduction. Thus I might know that I see a fire based on my present sensory experience, but I could not know that I am not *just* under an illusion in believing (falsely) that I see a fire. For such victims of illusion do *not* discriminate between illusion and reality. If I were thus victimized in thinking I see a fire, I would still believe both that I see a fire and that I am not *just* under an illusion in so believing.

It is a mistake, however, to think of 'R would not be so unless P were so' as ' $\sim\langle\rangle(R \& \sim P)$ ', i.e., as equivalent to 'It is not physically possible (relative to certain conditions) that $R \& \sim P$ ', which would entail that the conclusive reasons conditional is contraposable. Even without going into the ambiguity between ' $R \rightarrow P$ ' and ' $\sim P \rightarrow \sim R$ ', moreover, it can be seen through examples that this conditional does not contrapose. Thus let $R = \langle$ Water flows from the kitchen faucet \rangle , $Q = \langle$ The house's main valve is open \rangle , and $P = \sim(R \& \sim Q)$. Given that, 'R would not be so unless P were so' is intuitively compelling, but compare its contrapositive: ' $\sim P$ would not be so unless $\sim R$ were so' or, spelled out more fully, ' $(R \& \sim Q)$ would not be so unless $\sim R$ were so'. This by contrast is clearly false.

²² I take all the sentences involved here to be "conditionals," since they all express something x that they condition on something y, either as a necessary or as a sufficient condition. Thus a sentence may express $\langle p \rangle$ as sufficient for $\langle q \rangle$, and hence $\langle q \rangle$ as necessary for $\langle p \rangle$. The contrapositive of that sentence would be exactly the same except only for negating each of 'p' and 'q' and inverting their positions. (Here I have omitted the corner quotes that strict formality would require.)

Since the conclusive reasons conditional fails to contrapose, it does after all matter just how we read it. We cannot just think of it as saying that $\sim \leftrightarrow (R \& \sim P)$. The direction matters, and we now have a choice of direction. We can no longer just assume that the intuitions that attach to 'R would not be so unless P were so' will favor either direction indifferently.

APPENDIX III: CAN EPISTEMOLOGY SURVIVE CONTEXTUALISM?

What you know at a given time depends of course on features of your context. You can't know you see a fire, for example, unless there is a fire there, which involves how it is in your context. Everyone would surely agree on that much. Interesting disagreement sets in on the ways in which your knowledge, and even your justification, can depend on features of your context. Some hold that, beyond the obvious dependence of knowledge on truth (often enough *external* truth), there is also the dependence of knowledge on external causal or counterfactual properties of the constitutive belief. Thus, whether S knows that p at t, K_p , has been said to depend on one or more of the following:

- whether B_p (S's belief that p) is accidentally true
- whether B_p is caused by the fact that p
- whether B_p tracks the truth
- whether B_p derives from a reliable process.

And these are all said to be (often enough) matters of the (external) context of the believer at the time of the belief.

Internalists have reached near-consensus in conceding that externalism is (or may be) right about *knowledge*, while retreating to internalism of *justification*. More specifically, it is conceded that the total epistemic condition that a belief must satisfy

in order to qualify as knowledge is (or may be) often enough an external condition that goes beyond the external truth of that belief. According to internalists it is whether one is epistemically *justified* that depends only on what is internal to one's mind. Externalists disagree. In their view even what one is justified in believing depends on features of one's (external) context. Such externalists thus accept a kind of "contextualism" that concerns the dependence of the epistemic status of beliefs on the (external) context of the believer. Externalism with regard to a given epistemic status (knowledge, for example, or justification) holds that it *does* depend on the external context of the believer, while internalism holds that it does *not*.

The view that a given epistemic status depends thus on context is generally known as "externalism," however, and not as "contextualism." The latter term is reserved mostly for a different view, one about the truth conditions of epistemic *attributions*. According to such attributor-contextualism, the *truth* of the attribution of a certain epistemic status by the use of a certain epistemic term (say, "knows") depends on features of the context of attribution that involve not *just* the truth, or the content, of the subject's belief; indeed the *content* (in one sense) is itself determined by such contextual features.

The main terms appealed to by various forms of externalism do in fact seem thus attributor-contextual. This is quite plausibly true of 'x is accidental', of 'x is caused by y', and of 'RP is a reliable process'. And it is even arguably true of 'x tracks y', i.e., of 'x would be so if y were so, and would not be so if y were not so'.²³ Nevertheless, we can still distinguish externalism from contextualism. Any form of externalism is still a view in epistemology proper, one that lays down externalist conditions, necessary or sufficient, for a given epistemic status, and does so in the object language. It may lay down the condition that the belief not be accidentally true, for example, or that it be caused by the fact that constitutes its content, or that it track the truth, or that it derive from a reliable process. Of course, in specifying its

²³ And it seems no less true of the most plausible forms of traditional internalism of knowledge, with their dependence on a "sufficient justification" requirement. See here the work of Stewart Cohen; for example, Cohen 1988.

favored condition—say “reliable derivation”—it depends on the context of use to provide an index relative to which only does the claim have a truth value. Even so, we can still distinguish the question expressed in a certain context by a contextual expression, one posed in the object language, from metalinguistic questions as to the contexts in which that expression can be used, to what semantical effects.

If epistemic terms such as ‘knows’ are contextual, *and if relevant contextual variation is high enough*, this can have important consequences for epistemology, some destructive, others more hopeful.

Context variability might be very low, in which case the consequences for epistemology seem harmless. For example, if it turns out that there is a simple set of indices relative to which the truth value of epistemic statements is to be evaluated, and especially if these are threshold-setting matters along some unified dimension or two, then we can gain understanding through a good account of the ways in which the truth in epistemic matters depends on such factors, perhaps even a convincing response to the skeptic.

If on the contrary it turns out that the factors that fix the truth values of our utterances are motley, resisting any unified theory, then there can be one of at least two outcomes, both harmful. Either the use of epistemic terminology in philosophical reflection or discourse turns out to be meaningless, or it is meaningful but of little or no relevance to its use in everyday thought and communication.²⁴

Either way epistemology is diminished: destroyed if the former possibility holds, but still diminished, in at least two ways, even if only the second holds. First, we will have less scope to draw on common human experience in delineating the contours of our epistemic concepts. For example, it has been argued that our concept of epistemic justification involves an indirect evaluation of the epistemic character of the subject, since this is something that we socially need to track, so that, not surprisingly, our (of course social) concept of justification, and indirectly our concept of knowledge, might plausibly be contoured with sensitivity to this need.

²⁴ See Putnam 1998a, 1998b, and 2001.

If radical contextualism is correct, then this sort of argument will be suspect, unless it so happens that despite the motley variability inherent in contextual epistemic terminology, there is some constant factor that involves the epistemic character of the subject.

In fact it is not obvious how to distinguish such motley contextualism from outright multiple ambiguity. Accordingly, it becomes doubtful that we possess firmly delineated concepts that span the marketplace, the sports arena, the dinner table, the playhouse, the confessional, the art museum, the analyst's couch, the newspaper, the seminar room, etc., etc. What do the inquiries of philosophers come to if that is how it is? Well, what questions do philosophers pose in the object language when theorizing about knowledge? Whatever they are, if motley contextualism (or the thesis of radical ambiguity) is right, then these questions would have little or no bearing on the issues expressed by the folk in syntactically similar terms. They may be as disparate as are questions about the sides of rivers and questions about financial institutions. And the converse is then of course equally so, in which case consider the fact pressed by the contextualist--that in ordinary situations we speak truly often enough in saying of someone that they "know that p." That fact is then of doubtful relevance to the question posed in our philosophical reflection as to whether people know the things that ordinarily they think they know. Even worse, we do not then clearly retain the right to suppose that people do ever ordinarily think that they *know* much of anything, not if our use of "knows" in so supposing is governed by the context of our supposition, namely the context of philosophical reflection, whose standards are now supposed to diverge impressively from those operative in ordinary situations.

But why exactly is it so bad that we lose our connection with folk thought and talk given the contextualism under scrutiny? Why is this any worse than the fact that 'work' and 'charm' mean in discussion of physics something different from what they mean in ordinary discourse? It is not as though the questions about physics lose any of their meaning or interest simply because of that fact. I would like to suggest *one*

reason why there is at least a difference of degree here between science and philosophy. Philosophy tends to study commodities that we want not just to understand but also to possess. Philosophy has a closer connection with individual and social practice than does much of science. Consider happiness, pleasure, justice, rightness, goodness, autonomy, free will, identity through time, aesthetic appreciation, friendship, love, consciousness, propositional attitudes, concepts, and, yes, good reasons, memory, and testimony, as well as coherence, rationality, and knowledge--philosophically interesting commodities all. And note how they are all matters of proper vital interest not only to the philosopher but to the human being. We want not only to understand them but also to have them in our lives. (Of course a similar point could be made about their opposites, with proper inversion, but I will not stop here to formulate my point in the more convoluted way required by this dual of my immediate concern.)

Note the caution: philosophy *tends* to study such commodities; and there is a *difference of degree* here between science and philosophy; and philosophy has a *closer* connection with individual and social practice. This is all compatible with the fact that biology studies life, which we do want in our lives. But it seems to me plausible that the interest of philosophy in its main areas of study--in epistemology, ethics, and metaphysics, for example--is largely driven by the importance of our *possessing* the studied commodities, whereas the agenda of physics, biology, and geology is set rather by the desire for understanding, prediction, and control, where *these* are the desired commodities (which, incidentally, are themselves desiderata whose study is of special interest to the philosopher), and it is of little concern to these sciences that the subject matter of their inquiries should *also* constitute human desiderata.

It is no surprise that philosophy should tend to study main human desiderata given the very strong historical association of our discipline with wisdom in how to live individually and collectively. At least one important tradition in perennial philosophy has focused on how to attain proper sensitivity to what matters and how

to attain it in a properly balanced and coherent life, or in a properly ordered society. In judging all of this, moreover, it would seem only sensible epistemically and practically to proceed in awareness of what is in fact valued, under what conditions. But if radical contextualism is correct this brings with it the seeds of doubt that what we study in philosophy is what is valued in life.

True, the commodities picked out in the ivory tower by philosophical terminology, including the examples above, and even the likes of 'existence', may still interest us philosophers not only as objects of understanding but as desiderata, regardless of their place in the lives of the folk. However, even if that should be so, and even if appropriate desire for them would survive their full study, which may be doubted about knowledge if some radical skeptics are right about what that (incoherently) requires, even so philosophy will be diminished, surely, if it becomes irrelevant to human thought and practice generally, philosophers being the only exceptions.²⁵

That would be bad enough, but not nearly as bad as the denial of meaning, period, to the questions raised in ivory tower contexts of philosophical reflection. Nevertheless, this more destructive outcome can be seriously entertained if one takes such words of our common language to have their life in everyday affairs, so that if torn from those contexts they turn useless for the conduct or expression of our thought.

It is hard to know how to react to such radical claims except by testing them, by considering whether we are or are not able to think or communicate in philosophical reflection or discussion by use of the same epistemic vocabulary used in everyday life. Of course, if the use of such vocabulary leads to irresolvable paradox and incoherence, then it may be argued that something has indeed gone terribly wrong, which might be plausibly attributed to our having torn our words from their natural contexts. But until we tried other ways out, that would be an overreaction.²⁶

²⁵ Which is not incompatible with the possibility that philosophy should discover new things to care about.

²⁶ How to argue in this area is not obvious. Take a radical contextualism, either of the motley sort or of the sort that requires ordinary contexts for so much as making meaningful sense through the use of epistemic

More sensible and promising is the more moderate contextualism according to which epistemic vocabulary remains subject to the same linguistic rules as we move from everyday contexts to the study or seminar room. However, shift in context brings with it difference in semantic values, difference in reference and/or truth conditions. And this is still compatible with the unfortunate consequence of the more radical contextualism rejected above, since the change in semantic value may still entail a radical change in subject matter. Thus, the subject matter of epistemological vocabulary in ordinary contexts may diverge so widely from its subject matter in philosophical contexts that “epistemological” discussion in one context is entirely irrelevant to “epistemological” discussion in ordinary contexts, even assuming that there is some unity on each side of this divide. Enough unity may remain across the divide, however, that philosophical discussion does not just lose all connection with shared human values.²⁷

Appendix IV

Response to Pryor²⁸

A. Concerning Safety and Its Epistemic Value

vocabulary. The fact that such a contextualism would destroy epistemology, is a point against it to the extent that one finds it incredible that epistemology is either meaningless or irrelevant. But it is hardly a decisive refutation, since contextualism is an empirical hypothesis, presumably, in the assessment of which one would need to take into account all its implications, and not only whatever implausible implications it may have for the status of epistemology.

²⁷ I am assuming that we commonly do desire the “philosophical desiderata” on my list, and properly so, and, moreover, that we desire that we have them to at least a degree deemed “sufficient” by ordinary standards.

²⁸ This is a response to James Pryor’s comments on the present paper (minus this appendix) at a conference on contextualism in epistemology held at the University of Massachusetts at Amherst in the fall of 2003.

1. Jim Pryor raises a problem for the idea that safety does as well as sensitivity in explaining certain pro-skeptic intuitions. It argues for a certain conception of the safety-defining conditional $r > p$, to be read “it would be so that r only if it were so that p ,” or (in an inequivalent but closely related alternative reading) as “Not easily would it be so that r without it being so that p .” (These are readings I explicitly lay down in my paper and in other places, where a safe belief that p is defined as one based on a reason $\langle r \rangle$ such that $r \rightarrow p$.)

2. Pryor reminds us that the standard Lewis/Stalnaker semantics rules true any strong conditional with truths in both its antecedent and consequent. This, however, would have the result that for any true propositions $\langle r \rangle$ and $\langle p \rangle$ whatever, S 's belief-that- p is safe given only that it is based on $\langle r \rangle$, which would be disastrous for the value of safety in an account of knowledge. So we need a different semantics for our safety conditionals.

3. Pryor considers two possible ways to understand the safety conditional and expresses a preference for one of them, one that spoils my safety-requirement for dealing with certain clear and familiar examples (such as the barns examples).

On one understanding of the truth conditions for the ' $r > p$ ' conditional, context will set a threshold determining the boundaries of the neighborhood of worlds throughout which the material conditional ' $r \rightarrow p$ ' must be true.

Pryor prefers an alternative account, suggested by Nozick, according to which

... we understand counterfactuals of the form “If A were the case, then C would be the case” as saying that C is true throughout the A -neighborhood of the actual world. In other words, you go out until you find the closest A -worlds, then C has to be true in those worlds, and as you keep going out, C has to keep being true, until you get

to the next band of not-A worlds. This proposal works the same regardless of whether A is true in the actual world. And it does not rely on context to set a threshold. Context just gives us a measure of closeness, and the domain of worlds we're considering.

Here's how he then expresses his preference for this second view:

Now, as I said, this second story sounds better to me because it gives the same treatment to counterfactuals with true antecedents and counterfactuals with false antecedents. It doesn't make some of these counterfactuals depend on context in a novel way, unexemplified by the other counterfactuals.

4. Applying this to my proposal on safety, Pryor has this to say (with the terminology modified slightly to accord with my own):

So when I read Sosa's criterion for safety:

(S1) Your belief that P is safe iff it's based on an experiential reason R such that:
... $R > P$.

The way I understand it is as follows:

(S2) In the closest band of worlds where ... R [is true], P is true in all of those worlds.

Or, in other words:

(S3) A world in which ...[R is false] is at least as close as any world in which ...[R is true] but P is false.

And he goes on to question the ability of safety so defined to deal with our intuitive responses to the barns examples.

5. However, S3 seems implausible. Consider, for example, two coins C1 and C2, and the four possible outcomes of a concurrent double-toss: C1-heads & C2-heads, C1-heads & C2-tails, C1-tails & C2-heads, and C1-tails & C2-tails. And suppose that the actual outcome is the first, and that the second and third possibilities are equally remote from the actual. In that case, the following conditional (*) would seem false: *that C1 would come up tails only if C2 came up tails*. Yet the S3-like condition would rule our conditional (*) true, not false. For a world in which C1 comes up heads is at least as close as any world in which C1 comes up tails while C2 comes up heads.

6. It may be thought that S2 gives the right result for that example, since it would appear that there are close enough worlds in which C1 comes up tails while C2 comes up heads. However, if we accept S2 for that reason, then we also get the right result about the barns, so long as there is a close enough world where my visual experience is still as of a barn before me even while there is no barn before me. And in the barns example there do seem to be about equally close worlds in some of which I have my barn experience while in others I do not.

7. So the Nozickian interpretation suggested by Pryor seems not to create a problem for the safety requirement, after all. Even if it did clearly create a problem, however, I would not find that a compelling reason against the safety requirement. For I have my doubts about that interpretation anyhow. The main reason offered in its favor is that it yields the following pleasing result: *that conditionals with false antecedents get the same treatment as conditionals with true antecedents*. One might agree with this, however, without accepting the second account for conditionals with true antecedents. One might instead accept the first account for conditionals with false antecedents. This is especially plausible for 'r → p' when read as "Not easily would it

be so that r without it being so that p.” For one might hold that in a case where $\langle r \rangle$ is false, what is required for the truth of ‘ $r \rightarrow p$ ’ interpreted thus is not just that it be true that p in the *closest* worlds in which it is true that r, but that it be true that p in those worlds and also in other more remote worlds up to a certain (context-determined) threshold.

For example, suppose an arrangement of light, bulb, fuse, wiring, and switch with the following features:

- a. The switch is off: Not-S.
- b. The light is off: Not-L.
- c. The bulb is not burnt out: Not-B.
- d. The fuse is not blown: Not-F.
- e. The wiring is not defective: Not-W.
- f. The following possibilities are arranged in the order of remoteness given, from least remote to most remote (where the possibilities under comparison are the numbered possibilities, each combining five specified sub-possibilities):

Possibility 1: S, Not-L, B, Not-F, Not-W

Possibility 2: S, Not-L, Not-B, F, Not-W

Possibility 3: S, Not-L, Not-B, Not-F, W

Given these assumptions, how do we assess the following claims:

- C1. The switch would be on while the light was off only if the bulb were burnt out.**
- C2. Not easily would the switch be on with the light off without the bulb being burnt out.**

These seem increasingly implausible as we narrow the margin of remoteness among our three possibilities, so that once the margin is very small, both C1 and C2 become quite implausible.

What is more, I can feel no difference in intuitions once we replace assumptions a and b above with:

a'. The switch is on, and

b'. The light is on,

leaving c - f unchanged.

Under these assumptions, claims C1 and C2 still seem increasingly implausible as we narrow the margin of remoteness among our three possibilities. And, again, C1 and C2 become extremely implausible once the margin is extremely small.

8. Is the Nozick-inspired account to be used in interpreting the likes of C1 and C2? If not, then it is not relevant to our definition of safety, which appeals to claims like C1 and C2. If, however, that account *is* so used, then claims C1 and C2 turn out true in the case hypothesized. This result being so implausible, that interpretation seems unacceptable and cannot then spell trouble for the idea that safety is both necessary for knowledge, and no less effective than sensitivity in explaining our intuitions about barns.

B. Concerning Contextualism and Its Epistemic Value

1. Pryor opens his paper with these remarks:

There are two central strands in Sosa's paper.

The first concerns how Contextualists are to give full endorsement to the knowledge-claims of ordinary subjects. Just saying, metalinguistically, that:

When an ordinary subject says "I know I am sitting," she expresses a true proposition.

doesn't seem to be a full enough endorsement. We also want to know what proposition it is she's expressing. And to be able to describe that proposition, Sosa argues we have to resort to expressions like "knows by ordinary standards," which are no longer contextual. Sosa thinks it is those expressions, rather than Contextualism, that enable us in philosophical contexts to explain what is plausible in the Moorean position.

That sounds right to me. I'm not really sure that the Contextualist needs to disagree.

2. Here is why I think the Contextualist needs to disagree. The Contextualist offers a Contextualist explanation of the distribution of intuitions concerning the skeptical paradox created by the following skeptical argument:

- P1. It is not the case that K(not-h)**
- P2. If P1 then C**
- C. It is not the case that K(o)**

Where 'o' is short for 'One sees a hand', 'h' for 'One's present experience is nothing but illusion', and 'K(p)' for 'One knows that p'. ('One' here refers to an arbitrary subject at an arbitrary time.)

Here are some alternative responses to that argument:

The skeptic: P1, P2; therefore C.

Dretske, Nozick, et. al.: P1, Not-C; therefore, Not-P2.

The Moorean: P2, Not-C; therefore, Not-P1.

A paradox such as this presents an inconsistent triad. Here the three inconsistent propositions are P1, P2, and Not-C. The most satisfying response to the paradox will not simply choose two of the members of the triad and reject the third just on that basis. A more satisfying resolution will do something to explain why one's rejected proposition holds the attraction that it holds for those who adopt other stances. This is what contextualists believe their account enables them to provide. But how, exactly, does it work?

Take the skeptic's stance. When we adopt it in a high-standards context, does contextualism help us to explain what is plausible in affirming Not-C, as the other two stances prefer to do? Here the contextualist appeals to the metalinguistic claim that in ordinary contexts it is true to say that one "knows that o." So the contextualist reasons that, in a philosophical context with high standards in place, we can still retain awareness of the fact that when one ordinarily takes oneself to know that o, one is often enough correct in that judgment. Through overlooking the fact that 'knows' is contextual, one then finds oneself intuitively pulled to think, even in a

philosophical context when the standards are raised, that one knows that *o* (when one does see a hand). This is how, despite the difference in contexts, the truth of ordinary utterances of the form “I know I see a hand” is relevant to us as we consider in a philosophical context whether we can ever know that we see a hand.

If I wonder whether I am in Alaska, it will not help me that someone in Alaska speaks truly in saying “I am in Alaska.” It would be an obvious fallacy to descend semantically from that metalinguistic fact to the conclusion that I am in Alaska. This is of course so obviously fallacious that the truth of the metalinguistic assertion just seems irrelevant to our question in the object language. The epistemic case is not at all like that, says the epistemic contextualist, because the contextuality of epistemic vocabulary is much less obvious. Because the contextuality of epistemic vocabulary is so easily overlooked, we are pulled intuitively to say and think incorrect things once the standards are raised.

How deeply satisfying is this explanation? Compare this. Where pro basketball is under discussion it can be true to say that Michael Jordan is not tall, unlike Patrick Ewing. Yet Jordan is taller than our six-footer friend Tom, whom we know to have no connection or even interest in basketball. So we would presumably be able to infer that Tom is not tall. However, there would remain some intuitive pull towards saying that six-footer Tom *is* tall. And the contextualist could then explain here by parity of reasoning that when we feel that pull we overlook the contextuality of tallness vocabulary.

Here is an alternative account: Such threshold-involving vocabularies come with standards. And it is usually clear enough what the operative standards are, which enables us to communicate well enough with such contextual vocabulary. But how do we report indirectly the speech or thought of others across contexts governed by different standards? Normally we just relativize to the relevant standards. We say that

the pygmy thinks that his friend is tall, where what we mean is clearly that the pygmy thinks his friend tall for a pygmy, or tall by pygmy standards.

So, similarly, when we say that our six-footer friend is not tall, since he is shorter than Jordan, who is not tall, what we mean is clearly that our friend is not tall by NBA standards. And now let's apply this to the skeptical paradox presented by the skeptic's argument (SA):

- P1. It is not the case that K(not-h).
- P2. If it is not the case that K(not-h) then it is not the case that K(o).
- C. It is not the case that K(o).

Presented in a high-standards context, this would be a sound argument, or so we are told. And let us grant this much for the sake of argument. But why do we still feel an intuitive pull to say that K(o)? Because we overlook the heightened standards and are tempted to import into the more demanding context what we say in an ordinary context with its lower standards.

By analogy with the case of tallness, then, once we see how matters stand, we should be able to relativize in ways familiar from such contextual vocabulary as 'tall', 'big', 'heavy', 'long', and many, many other terms with which we communicate smoothly and efficiently in ordinary discourse. We can just make explicit what is normally implicit in the use of such vocabulary, by replacing 'is F' with 'is F by such-and-such standards'. When we do that with our skeptic's argument SA, we get this:

SA+

- P1+. It is not the case that K+ (not-h).
- P2+. If it is not the case that K+ (not-h) then it is not the case that K+ (o).

C+. It is not the case that K+ (o).

And now the explanation of why we feel tempted to reject C+ is that we confuse 'K+ (o)' (read: "one knows by high standards that one sees a hand") with 'K- (o)' (read: "one knows by ordinary standards that one sees a hand"). So one is tempted to think the former true because one correctly thinks the latter true while failing to discern the important difference between the two.

Compare, however, the following argument:

SA-

P1-. It is not the case that K- (not-h).

P2-. If it is not the case that K- (not-h) then it is not the case that K- (o).

C-. It is not the case that K- (o).

Can't the skeptic still argue *this* way to similar effect? Won't the skeptic's reasoning in favor of P1+ be about as effective when it is used rather in support of P1-. If so, then the help we get from contextualism in dealing with the skeptic is limited at best. For one thing, now contextual terms have been relativized away. And we still face what is essentially the same skeptic, who is left standing even once we abandon our interest in K+, with its high standards, in favor of K-, with its ordinary standards.^{29 30}

²⁹ Here I have developed a line of reasoning sketched in my "Skepticism and Contextualism," *Philosophical Issues 10: Skepticism* (Supplement to *Nous* for 2000), p. 6: "Do the arguments of skeptics really concern only the attainment of some apex along a dimension of epistemic justification? Or do the most powerful and interesting skeptical arguments concern rather whether we can ever progress *to any distance whatever* from the nadir of justification? If the latter, then again, now in a different way, contextualist considerations may have limited relevance against the skeptic."

³⁰ Acknowledgments

REFERENCES

Austin, J.L. 1962. *Sense and Sensibilia*. Oxford: Oxford University Press.

Cohen, S. 1986. Knowledge and Context. *Journal of Philosophy* 83.

Cohen, S. 1988. How To Be a Fallibilist, in Tomberlin 1988, pp. 91–123.

Dancy, J., ed., 1988. *Perceptual Knowledge*. Oxford: Oxford University Press.

DeRose, K. 1995. Solving the Skeptical Problem. *Philosophical Review* 104: 1–52.

Reprinted in DeRose and Warfield 1999.

DeRose, K., and Warfield, T., eds. 1999. *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press.

DeRose, K. Forthcoming. Sosa, Safety, Sensitivity, and Skeptical Hypotheses. In Greco, forthcoming.

Dretske, F. 1971. Conclusive Reasons. *Australasian Journal of Philosophy* 49:1–22.

Reprinted in Dretske 2000.

Dretske, F. 1970. Epistemic Operators. *Journal of Philosophy* 67: 1007–1023. Reprinted in DeRose and Warfield 1999, and in Dretske 2000.

Dretske, F. 2000. *Perception, Knowledge and Belief*. Cambridge: Cambridge University Press.

Goldman, A. 1976. Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73: 771–91. In Dancy 1988.

Goldman, A. 1989. Psychology and Philosophical Analysis. *Proceedings of the Aristotelian Society* 89: 195–209. Reprinted in Goldman 1992.

Goldman, A. 1992. *Liasons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press.

Greco, J. 2000. *Putting Sceptics in Their Place*. Cambridge, UK: Cambridge University Press.

Greco, J. Forthcoming. *Sosa and His Critics*. Cambridge, MA, and Oxford, UK: Blackwell Publishers.

Klein, P. 2002. Skepticism. In Moser 2002, pp. 36–62.

Lewis, D. 1996. Elusive Knowledge. *Australasian Journal of Philosophy* 74: 549–67.

Luper-Foy, S., ed. 1987. *The Possibility of Knowledge: Nozick and His Critics*. Totowa, NJ: Rowman and Littlefield.

Moser, P., ed., 2002. *The Oxford Handbook of Epistemology*. Oxford, UK: Oxford University Press.

Nozick, R. 1981. Knowledge and Skepticism, *Philosophical Explanations*. Cambridge, MA: Harvard University Press. Reprinted in Luper-Foy 1987.

Peirce, C.S. 1931–58. *Collected Papers of Charles Sanders Peirce*. Cambridge, MA: Harvard University Press.

Pryor, J. 2000. The Skeptic and the Dogmatist. *Nous* 34: 517–49.

Putnam, H. 1998a. Strawson and Skepticism. In *The Philosophy of P.F. Strawson*. Chicago and La Salle, Ill.: Open Court.

Putnam, H. 2001. Skepticism, Stroud, and the Contextuality of Knowledge. *Philosophical Explorations* (January): 2–16.

Putnam, H. 1998b. Skepticism. In Marcelo Stamm (ed.), *Philosophie in Synthetischer Absicht*. Stuttgart: Klett-Cotta.

Schiffer, S. 1995/6. Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society*. London: The Aristotelian Society.

Sosa, E. 1994. Philosophical Scepticism and Epistemic Circularity. *Aristotelian Society Supplementary Volume* 68. In DeRose and Warfield 1999.

Sosa, E. 1997. Reflective Knowledge in the Best Circles. *Journal of Philosophy* 94: 410–30.

Sosa, E. 1999. How to Defeat Opposition to Moore, in Tomberlin 1999, pp. 141–53.

Sosa, E. 2000. Contextualism and Skepticism. *Philosophical Issues XI: Supplement to Nous*. (With critical comments by H. Kornblith, K. Lehrer, and J. Tomberlin, and replies.)

Sosa, E. 2002. Tracking, Competence, and Knowledge. In Moser 2002, pp. 264–87.

Stine, G.C. 1976. Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies* 29: 249–61. Reprinted in DeRose and Warfield 1999.

Tomberlin, J. 1988. *Philosophical Perspectives, 2, Epistemology: Supplement to Nous*.

Tomberlin, J. 1999. *Philosophical Perspectives, 13, Epistemology: Supplement to Nous*.