



External World Skepticism

John Greco*

Saint Louis University

Abstract

Recent literature in epistemology has focused on the following argument for skepticism (SA): I know that I have two hands only if I know that I am not a handless brain in a vat. But I don't know I am not a handless brain in a vat. Therefore, I don't know that I have two hands. Part I of this article reviews two responses to skepticism that emerged in the 1980s and 1990s: sensitivity theories and attributor contextualism. Part II considers the more recent 'neo-Moorean' response to skepticism and its development in 'safety' theories of knowledge. Part III argues that the skeptical argument set out in SA is not of central importance. Specifically, SA is parasitic on skeptical reasoning that is more powerful and more fundamental than that displayed by SA itself. Finally, Part IV reviews a Pyrrhonian argument for skepticism that is not well captured by SA, and considers a promising strategy for responding to it.

Much of the recent literature in epistemology has focused on the following skeptical argument.¹

SA

Let o be some ordinary proposition about the external world, such as that I have two hands, and let h be a proposition describing some skeptical hypothesis, such as that I am a handless brain in a vat.

1. I know that o only if I know that $\neg h$.
2. But I don't know that $\neg h$.

Therefore,

3. I don't know that o .

The argument generalizes: we can take nearly any proposition about the external world, and we can choose a suitable skeptical hypothesis so as to generate an argument with a similar form. A robust skepticism about the external world threatens.

Part I of this article reviews the two major responses to **SA** that emerged in the 1980s and 1990s: sensitivity theories and attributor contextualism. Some objections that have been raised against these anti-skeptical responses are also reviewed. Part II considers a third anti-skeptical response to **SA** that has emerged more recently. This 'neo-Moorean' response can be taken in a number of directions, but here we focus on its development in 'safety'

theories. Part III considers reasons for thinking that the skeptical argument set out above is not of central importance. Specifically, there is good reason to think that **SA** is parasitic on skeptical reasoning that is more powerful and more fundamental than that displayed by **SA** itself. Finally, Part IV reviews a Pyrrhonian line of skeptical reasoning that is not well captured by **SA**, and considers a promising strategy for responding to it.

I. Sensitivity Theories and Contextualist Theories

The skeptical argument that we are considering proceeds as follows.

SA

1. $Ko \Rightarrow K\neg h$

2. $\neg K\neg h$

Therefore,

3. $\neg Ko$

Sensitivity theories respond to the skeptical argument by denying premise 1: that one knows that one has two hands only if one knows that one is not a handless brain in a vat. More generally, sensitivity theories deny the skeptical thought that one can know ordinary truths about the world only if one also knows that various skeptical scenarios are false.²

Contextualist theories grant that the skeptical argument is sound, and its conclusion true, relative to the philosophical contexts in which these are typically expressed. But contextualists limit the skeptical damage by adding a further claim: that the skeptical argument is unsound, and its conclusion false, relative to ordinary contexts in which ordinary knowledge claims are typically expressed.³

The burden of sensitivity theories is to explain why the relevant skeptical premise and the associated skeptical thought are false. To this end, sensitivity theorists develop accounts of knowledge, or at least partial accounts of knowledge, that are intended to do that job. The burden of contextualist theories is to explain how the skeptic's claims and ordinary knowledge claims can all be true. To this end, contextualists develop accounts of knowledge language intended to do that job. The next sections consider these strategies in more detail.

1. SENSITIVITY THEORIES AND THE DENIAL OF CLOSURE

Sensitivity theories deny premise 1 of **SA**, claiming that ordinary knowledge about the world does not require knowing that various skeptical hypotheses are false, or that various skeptical possibilities are not actual. The burden of this approach is to make that move plausible, and this is no easy task. For one, the skeptical premise is intuitively plausible as it stands. How could one know that one has two hands while *not* knowing that one is a handless brain in a vat? But second, the premise gains support from various 'closure principles' in the neighborhood.

Let's define a closure principle as a proposition that knowledge is closed under some operation. In other words, a closure principle says that if you start from knowledge, and perform some operation on that knowledge, you get more knowledge as a result. For example, it is plausible that knowledge is closed under known deduction: If you know that p , and you know that p entails q , then you know that q . One might quibble with this principle, insisting that knowledge requires believing that q on the *basis* of deducing it from p . But presumably the principle could be fixed up accordingly – for example, we can say that knowledge is closed under 'competent' deduction. The important point here is that there seem to be plausible closure principles in the neighborhood, and it seems that such principles will be satisfied by o and h in **SA** above. For example, it seems that if you know that you have two hands, and if you know that having two hands entails that you are not a handless brain in a vat, then you know (or can know by deducing it) that you are not a handless brain in a vat. Put another way, it seems that premise 1 of **SA** is well supported by plausible closure principles.⁴

So sensitivity theorists incur a heavy burden. They try to discharge it by providing an account of knowledge that explains why, contrary to first appearances, premise 1 and nearby closure principles are false. More exactly, sensitivity theorists propose a necessary condition on knowledge that is supposed to do this job. This approach has been championed by, among others, Fred Dretske and Robert Nozick.⁵ Here I will ignore the details of their respective views and focus only on the work that is supposed to be achieved by making 'sensitivity' a necessary condition for knowledge.

The most straightforward way to understand the sensitivity condition is in terms of a subjunctive conditional:

Sensitivity. S knows that p only if: If p were false, S would not believe that p .⁶

Intuitively, if one knows that p , then one would not believe p anyway if p were false. On the standard semantics for subjunctive conditionals, we can understand *Sensitivity* this way:

S knows that p only if: In the closest possible world where p is false, S does not believe that p .

The idea that knowledge requires sensitivity is intuitively plausible. For example, suppose that you now know that you have two hands. Intuitively, if you did not have two hands you would not believe that you did. If you did not have two hands, you would see that you didn't. On the other hand, suppose your belief that you have two hands is not 'sensitive'. Suppose, for example, that your brain is being manipulated so that you would believe that you have two hands even if you did not. Intuitively, in that case you do *not* know that you have two hands, and precisely because your belief that you do is not sensitive.

Placing a sensitivity condition on knowledge yields a distinctive strategy for responding to the skeptical argument above. First, the sensitivity theorist can accept premise 2 of **SA** and can explain why it is true. In short, one

does not know that skeptical hypotheses are false because one's beliefs to that effect are not sensitive. For example, you now believe that you are not a handless brain in a vat, fed ordinary experiences by a supercomputer stimulating your severed nerve endings so as to simulate an ordinary life. But if you were, you would still believe that you were not. Your belief that you are not a brain in a vat is not sensitive and therefore not knowledge.

More importantly for present purposes, the sensitivity theorist can reject premise 1 of **SA**, along with supporting closure principles in the neighborhood. Consider: We have just seen that your belief that you have two hands is sensitive. But we have also just seen that your belief that you are not a handless brain in a vat is not sensitive. Assuming that your belief that you have two hands satisfies other conditions on knowledge, it follows that you know that you have two hands even though you do not know that you are not a handless brain in a vat. Premise 1 of **SA** is false. By the same reasoning, relevant closure principles are false as well.

We will consider two problems for this anti-skeptical approach below. But first I want to consider a reaction to the approach that is natural but misguided. The reaction is to think that the approach begs the question against the skeptic in some unacceptable way. At most, one might think, the sensitivity theorist achieves only a 'stand-off' with the skeptic, since the approach assumes that ordinary beliefs about the world, such as that I have two hands, are sensitive. That is, the approach assumes that one is not a handless brain in a vat, and that is not an assumption that the skeptic is willing to concede.

This reaction is natural enough, but it misconceives the nature of the anti-skeptical project under consideration. The purpose of that project is not to persuade a non-believing skeptic, or to otherwise refute the skeptic in a way that is rhetorically satisfactory. Rather, the project is to reject something in the skeptical reasoning under consideration, whether or not a 'real' skeptic would be satisfied with that rejection. Put another way, the project is to critique the skeptical argument rather than to convince or persuade a skeptical person (ourselves or someone else). In the context of this project, we are looking for a response to skepticism that is theoretically adequate, as opposed to rhetorically or pragmatically adequate. And what would theoretical adequacy require? Just what the sensitivity theorist pretends to provide: an account of knowledge (in this case a partial account) that explains where the skeptical argument goes wrong, and thereby explains how knowledge is possible.⁷

So that is what sensitivity theorists pretend to do. Are they successful? Here I will review two objections that have been raised against the approach in this regard.

The first objection is that a rejection of premise 1 comes at too high a theoretical cost. One way to put the objection is this: Premise 1 of **SA** and relevant closure principles seem more plausible than the account of knowledge that rejects them. On this way of thinking, the rejection of

closure principles should be seen as a *reductio* of the sensitivity condition rather than a consequence of it. Putting the objection this way, however, amounts to little more than a conflict of intuitions. A more persuasive statement of the objection calls attention to what Keith DeRose calls ‘abominable conjunctions’.⁸ Specifically, it would seem that the sensitivity theorist is committed to embarrassing claims such as the following:

- a. I know that I have two hands, but not that I am a handless brain in a vat.
- b. I know that I am sitting in front of a computer, but not that I am merely deceived into thinking I am sitting in front of a computer.
- c. I know I am not in a vat, but I don’t know that I am not a handless brain in a vat.

DeRose’s point is evident. The conjunctions in a–c border on absurdity, and therefore count heavily against any theory that entails them.

A second objection against sensitivity theories is that they cannot accommodate clear cases of inductive knowledge. This sort of objection has been pressed by Jonathan Vogel, who offers the following two examples.

The Hole-In-One. Sixty golfers are entered in the Wealth and Privilege Invitational Tournament. The course has a short but difficult hole, known as the ‘Heartbreaker’. Before the round begins, you think to yourself that, surely not all sixty players will get a hole-in-one on the ‘Heartbreaker’. (‘New Relevant Alternatives Theory’ 165)

The Rookie Cop. Suppose two policemen confront a mugger, who is standing some distance away with a drawn gun. One of the officers, a rookie, attempts to disarm the mugger by shooting a bullet down the barrel of the mugger’s gun. . . . Imagine that the rookie’s veteran partner knows what the rookie is trying to do. The veteran sees him fire, but is screened from seeing the result. Aware that his partner is trying something that is all but impossible, the veteran thinks (correctly as it turns out) [that the] rookie missed. (‘Tracking’ 212)

Both of these cases have the following structure: there is a close world where a highly improbable possibility is actual. Put another way, in each case almost all close worlds are *p*-worlds, but there are some not-*p*-worlds close to the actual world. Notice that in this sort of case, sensitivity is violated. For example, in the nearest world where all sixty players will get a hole-in-one, you still believe that they won’t. In the nearest world where the rookie cop does not miss, the veteran cop still believes that he does. But it is absurd to think that there is no knowledge in the two cases.

Finally, the following case from John Hawthorne suggests that the problem is ubiquitous.

The Disappearing Desk. Suppose that there is a desk in front of me. Quantum mechanics tells us that there is a wave function that describes the space of nomically possible developments of the system that is that desk. On those interpretations of quantum mechanics according to which the wave function gives probability of location, there is some non-zero probability that, within a

short while, the particles belonging to the surface of the desk remain more or less unmoved but the material inside the desk unfolds in a bizarre enough way that the system no longer counts as a desk. Owing to its intact surface, the system would be reckoned a desk by normal observers. (4–5)

Sensitivity theorists therefore face two formidable objections. First, their account seems to entail counter-intuitive results, such as the denial of plausible closure principles and DeRose's abominable conjunctions. Second, their account seems inadequate to accommodate at least some cases of inductive knowledge. Perhaps the best way to answer these objections is to take a holistic approach. The idea here is that any account of knowledge will have costs and benefits in the face of **SA** and related problems. One might argue, then, that the theoretical benefits of sensitivity theories outweigh the costs, relative to competitors. Perhaps this is the best way to understand the case put forward by Dretske and Nozick, and more recently by Kelly Becker.⁹ The following sections put the competitors more clearly in view.

2. ATTRIBUTOR CONTEXTUALISM

Contextualist responses to **SA** grant that the skeptical reasoning is sound, and that the skeptical conclusion is true, in most cases where they are expressed. But contextualists deny that this has widespread skeptical consequences. The guiding idea is that the truth-conditions of knowledge attributions vary across conversational contexts, with the following anti-skeptical effect: Although the skeptic typically says something true when she says, 'You don't know', ordinary speakers (in ordinary conversational contexts) typically say something true when they say, 'I do know'.

Here is a simple illustration of how this might work. Suppose that in contexts where the skeptic is making her arguments the standards for 'knowledge' are very high. For example, in these contexts it requires a great deal of evidence, of very great quality, for sentences of the form 'S knows that p' to come out true. But suppose also that in ordinary contexts standards are much lower. In these ordinary contexts one requires much less evidence for sentences of the form 'S knows that p' to come out true. If these suppositions hold, then our knowledge language will work exactly as the contextualist suggests: the skeptic will be right when she claims 'S does not know that p', but ordinary speakers will often enough be right when they claim 'I do know that p'.¹⁰

In addition to this anti-skeptical effect, contextualists claim two further advantages of their view. First, it explains why skeptical arguments can seem so convincing. Skeptics and non-skeptics alike have long noted a puzzling dynamic: skeptical arguments can seem persuasive while we are engaging them, but then their power fades as soon as we cease from philosophizing.¹¹ The contextualist response to skepticism has a ready explanation for this phenomenon. Namely, in contexts where we are considering skeptical

scenarios and the like, the standards for knowledge get raised unusually high, and so *the skeptic is right* when she claims ‘We do not know’ in those contexts. On the other hand, the standards for knowledge operative in ordinary contexts are not so high, and so when we return to these in our non-philosophical lives *we are right* when we claim ‘We do know’ in those contexts. In each context, the claims that we make seem true because they *are* true!

Second, contextualism achieves its anti-skeptical effect without denying plausible closure principles. Relative to skeptical contexts, we ‘know’ neither that ordinary propositions about the world are true nor that skeptical scenarios are false. Relative to ordinary contexts, however, we ‘know’ both that ordinary propositions about the world are true and that skeptical scenarios are false. We may not be able to *say* that we know that skeptical scenarios are false, since mentioning skeptical scenarios tends to move us out of an ordinary conversational context and into one that is skeptical, thereby raising the standards for ‘knowledge’ and making the saying false. The point remains, however, that closure principles hold across all contexts: in no single context is the claim ‘I know that I have two hands’ true and ‘I know that I am not a handless brain in a vat’ false.

Some philosophers have argued that contextualist theories concede too much to the skeptic, however. As we have seen, the contextualist is happy to say that the skeptic is right relative to skeptical context – when the skeptic claims ‘You don’t know that you have two hands’, or ‘No one knows he is not a brain in a vat’, these claims are true in the contexts where they are made. But many philosophers would like to deny just that. That is, many philosophers want to say that the skeptic is wrong when she makes such claims. One way to press the point is to consider Moore’s statements when he says things like ‘Of course I know that here is a hand’, and ‘Of course I know that the world has existed for more than five minutes’. Many philosophers want to say that Moore is right when he opposes the skeptic in these ways, even though Moore’s claims are made in a context where he is engaging skeptical arguments.

DeRose has argued that the contextualist need not make the concessions at issue.¹² It is consistent with contextualism, DeRose points out, to hold that skeptical considerations fail to raise the standards for knowledge as high as skepticism requires, so that knowledge claims come out true, and skeptical claims come out false, even relative to standards operative in philosophical contexts. One way this might happen is if non-skeptics have ‘veto power’ over skeptical attempts to raise the standards for knowledge too high. The thought here is that, in general, raising and lowering standards requires more than just willing it to be so. That being the case, contextualists can posit conversational mechanisms that prevent skeptical standards from coming into play even in philosophical contexts. A second option open to contextualists is to posit ‘gaps’ in the truth-values of knowledge claims, so that competing standards in philosophical contexts result in knowledge claims being neither true nor false.¹³

DeRose is surely right to point out these options for contextualism. But many philosophers will not be satisfied. On the 'Gaps View', our knowledge claims continue to come out not true relative to philosophical contexts, although skeptical claims come out not true as well. Some will think that this still concedes too much. What about the 'Veto Power View'? Here our knowledge claims come out true and skeptical claims come out false, relative to both ordinary and philosophical contexts. However, it is not clear that this view still gives us a *contextualist* response to the skeptical argument. The real anti-skeptical work, it would seem, will require a theory of knowledge that explains why skeptical claims are false, and how non-skeptical claims can be true, across the board.

II. Neo-Moorean Responses

Recall premise 1 of **SA**: that I don't know that I have two hands unless I know that I am not a handless brain in a vat. G. E. Moore famously noted that this thought cuts both ways. One can just as well use it to argue that, since I *do* know that I have two hands, therefore I also know that I am not a brain in a vat.¹⁴ Neo-Mooreans follow Moore on this tack, but try to provide an account of knowledge to back it up. That is, they try to explain *how* one knows, in the typical case, that skeptical scenarios are false. In this sense, neo-Mooreans are involved in the same project as, and incur a burden analogous to, sensitivity theorists. They are in the business of providing an account of knowledge that can explain and support a theoretical response to the skeptical argument under consideration. Hence,

Neo-Moorean responses follow G. E. Moore by denying premise 2 of **SA**. More generally, neo-Moorean responses deny the skeptical thought that, in the typical case, one does not know that various skeptical scenarios are false.¹⁵

One example of this neo-Moorean approach is provided by James Pryor, who offers an account of perceptual justification (and perceptual knowledge) on which one can be justified in believing that one has two hands without being *antecedently* justified in believing that skeptical scenarios are false. Rather, one can be justified in believing that one has two hands, and even know that one does, on the basis of one's perceptual experience alone, without further evidence about one's perceptual conditions, the reliability of one's experience, the reliability of one's perceptual powers, or the like. Having gained this sort of justification via perceptual experience, one can then go on to reason that various skeptical scenarios are false, mimicking Moore's reasoning above.¹⁶

Much of Pryor's approach is consistent with both internalist and externalist accounts of justification and knowledge.¹⁷ Here I will focus on an externalist version of the neo-Moorean strategy – one that invokes a 'safety' condition on knowledge as opposed to the sensitivity condition discussed above.

1. SAFETY THEORIES

We said that S's belief that p is *sensitive* just in case it satisfies the following condition:

If p were false, S would not believe that p .

Alternatively,

In the closest possible world where p is false, S does not believe that p .

Ernest Sosa has argued that a belief is better safe than sensitive, where S's belief that p is safe just in case it satisfies the following condition:

S would believe that p only if p were true.

Here Sosa means to propose an alternative necessary condition on knowledge. Accordingly, we have:

Safety. S knows that p only if: S would believe that p only if p were true.¹⁸

The intuitive idea here is that, in cases of knowledge, one could not easily have been wrong. Alternatively, not easily would S believe that p when p is false. One's belief is therefore 'safe' in that sense. There is some controversy over how to best capture this intuitive idea, however. First, we can distinguish between a strong and a weak reading of the subjunctive conditional in *Safety*.

Strong Safety. S knows that p only if: In close possible worlds, always if S believes that p then p is true. (In close possible worlds, never does S believe that p and p is false.)

Weak Safety. S knows that p only if: In close possible worlds, usually if S believes that p then p is true. (In close possible worlds, almost never does S believe that p and p is false.)

Plausibly, *Strong Safety* does no better with the counterexamples raised against sensitivity theories in Part I. Those examples were constructed so that there are a small number of not- p -worlds very close to the actual world, insuring that the sensitivity condition is violated in cases that seem to be knowledge. Notice, however, that the condition expressed in *Strong Safety* is also violated in those examples. For example, there is a close world where S believes that the rookie misses, but it is false that the rookie misses. Since Sosa means to endorse a safety condition as an alternative to a sensitivity condition, it makes sense to interpret him as endorsing *Weak Safety*.¹⁹

Pritchard has recently argued for a position between *Strong Safety* and *Weak Safety*. The guiding idea is that knowledge is most threatened by error in the closest nearby worlds. Pritchard ties this idea to more general considerations about luck. In general, Pritchard argues, judgments about luck place more weight on those counterfactual events that are modally closest. For example, suppose that a sniper fires two shots, the first of which misses your head by inches and the second of which misses by yards.

Intuitively, you are luckier to be missed by the first shot than to be missed by the second. Put differently, you more easily could have been hit by the first shot than by the second. Pritchard's idea is that knowledge is intolerant of luck in a similar way, and that this is what the safety condition should capture. Accordingly, Pritchard suggests the following:

P-Safety. S knows that *p* only if 1) in all close possible worlds, usually if S believes that *p* then *p* is true, and 2) in the closest possible worlds, always if S believes that *p* then *p* is true.²⁰

We may now consider the safety theorist's approach to **SA**. As with other versions of the neo-Moorean approach, the strategy is to deny premise 2 of **SA**. That is, the strategy is to deny that I don't know that skeptical possibilities are false. First, consider my ordinary beliefs about the world, such as my belief that I have two hands. In normal environments, where no brains in vats or deceiving demons exist, many such beliefs will count as safe. For example, in normal environments where I believe that I have two hands, I would believe this only if it were true. But the same is true of my belief that I am not a handless brain in a vat. Since there are no close worlds where I am a handless brain in a vat, there are no close worlds where I believe that I am not but I am. The safety condition is satisfied. Assuming that remaining conditions on knowledge are satisfied as well, a safety theory allows that I know that I am not a handless brain in a vat. Similar considerations will apply to other skeptical scenarios.

Is it fair to assume that remaining conditions on knowledge are satisfied? An obvious worry is that, taken by itself, the safety condition is quite weak, and so it is no surprise that it is easily satisfied. To adequately determine the anti-skeptical force of the position, we need to know what conditions must be added to safety to get sufficient conditions for knowledge. One kind of case is especially relevant in this context; namely, those where S believes a proposition that is true in all close worlds, and therefore satisfies the safety condition by default.

For example, suppose that S is severely color-blind, so that he is unable to discriminate green from non-green objects. Suppose also that S forms a perceptual belief that the frog he sees is green (and S has no other reason for believing that the frog is green). Finally, suppose that frogs are by nature green, due to some feature of frog DNA. Accordingly, frogs are green in all nearby possible worlds. Given that S is color-blind, S could easily be wrong about the colors of other objects in the environment – he could easily mistake a non-green object for a green object. But S could not be easily wrong that the frog is green, since (we are supposing) this is a stable fact about a natural kind.

S's belief that the frog is green is safe, but clearly S does not know that the frog is green. What more is needed? Sosa argues that we must add a broader cognitive ability, one that gives rise to the safety of the particular belief in question. According to Sosa, a belief's safety 'must be fundamentally

through the exercise of an intellectual virtue', where an intellectual virtue is a reliable or trustworthy source of truth.²¹ In the frog case above, *S* lacks a broader perceptual ability (for discriminating green objects from non-green objects) to ground the safety of his belief, and this explains why *S* does not know.

Sosa's suggestion, then, is to add a virtue-theoretic condition to a safety condition. In general, knowledge is true safe belief grounded in a broader intellectual virtue or ability. Greco has argued that a safety condition falls out of the virtue-theoretic condition.²² Following Sosa, we may think of an intellectual virtue as a kind of ability – an ability to form true beliefs and avoid false beliefs within a relevant range and under relevant conditions. Visual perception, for example, is (very roughly) an ability to form true beliefs about the locations and orientations of mid-size physical objects, under conditions of good light, etc. But abilities in general are to be understood in modal terms.

To see the point, consider that one might have success in the actual world without ability. In short, one might be lucky. For example, I might successfully hit a baseball in the actual world, but only because, by good fortune, my bat is in the right place at the right time. Ability requires counterfactual success – one has ability only if one continues to hit the ball in worlds that are relevantly close. For example, in worlds where the ball comes in a little higher or a little faster, the player with ability adjusts her swing accordingly.

But now the same applies to intellectual abilities. The person with excellent perception forms true beliefs and avoids false beliefs in the actual world, but continues to do so in relevantly close worlds. And that entails that a safety condition will be satisfied by her perceptual beliefs: In relevantly close worlds, if *S* (perceptually) believes that *p* then *p* is true.²³

Recall that one of the advantages claimed for contextualism is that it explains the pull of skeptical considerations and the appeal of the skeptical arguments such as **SA**. In particular, contextualism explains the appeal of premise 2 of **SA**, the claim that I do not know that I am not a handless brain in a vat. Neo-Moorean accounts insist that premise 2 is false; that in the typical case one *does* know that one is not a handless brain in a vat. But then how do neo-Moorean accounts, and safety accounts in particular, explain the appeal that skeptical arguments such as **SA** have in the first place? Sosa's explanation is that it is easy to confuse safety with sensitivity. Beliefs about skeptical possibilities are not sensitive. Knowledge requires safety. But since sensitivity and safety are easily confused, one might confusedly think that one's (in fact safe) belief fails to satisfy a necessary condition on knowledge.²⁴

Duncan Pritchard offers a different explanation of pro-skeptical intuitions.²⁵ According to Pritchard, in typical cases we do know that skeptical possibilities are false, but claiming that we know violates pragmatic rules governing what is assertable in a conversational context. In particular, Pritchard invokes Grice's 'conversational maxim of evidence', which states that one's assertions

should be supported by adequate evidence. Crucially, however, what counts as ‘adequate evidence’ changes with the ‘purpose or direction’ of the conversational context. This implies that what counts as assertable changes as well, with the result that knowledge is sometimes unassertable in a context. The application to skeptical considerations and skeptical arguments should be evident: in philosophical contexts where skeptical arguments and considerations are in play, the direction and purpose of those conversations make knowledge claims about the external world unassertable, since knowledge of the external world is exactly what is at issue. For example, in such contexts it is inappropriate to assert that I am not a handless brain in a vat, or that I know that I am not, even if both claims are literally true.

2. THE OBJECTION TO SAFETY THEORIES: THAT’S TOO EASY!

A number of objections have been raised against safety theories of knowledge, but here I will focus on a family of objections directed specifically at the safety theorist’s neo-Moorean response to skepticism. The unifying theme of this family of objections is that the safety approach makes responding to skepticism too easy.

One way to understand this charge is that a safety approach ‘begs the question’ against skepticism in an inappropriate way. Specifically, the approach assumes that there is no close world where one is a brain in a vat or the victim of a deceiving demon, and so it assumes that one is not so victimized in the actual world. But this objection recalls the ‘natural but misguided’ objection that we saw raised against sensitivity theories in Part I. Our response there was to clarify the nature and purpose of a sensitivity theory (and now a safety theory) of knowledge. Such theories do *not* try to give accounts that would be persuasive in a debate with a committed skeptic. Rather, the idea is to give an account of knowledge that challenges something in the skeptical argument – that explains where the skeptical argument goes wrong, and thereby explains how knowledge is possible. The goal is not to offer something that is dialectically appropriate in a debate. It is to offer something theoretically adequate in an explanation.

Another version of the objection charges that safety theories beg the question in a different sense: they deny some essential component of the skeptical problematic. For example, it is sometimes claimed that the skeptic is working with an internalist conception of epistemic justification (where ‘epistemic justification’ names the sort of justification that knowledge requires). Insofar as safety theories adopt an externalist approach to justification, they deny an essential assumption of the skeptic’s reasoning. This sort of objection is surely misguided, however, in that any anti-skeptical approach must deny *something* in the skeptical argument. The argument of **SA**, for example, is formally valid. Any approach that means to avoid its conclusion must deny at least one of the argument’s premises.²⁶ Moreover, safety theories do not *merely* deny some assumption of the skeptic’s reasoning

– they motivate that move with a theory of knowledge that explains why the premise in question is false. That sort of theoretical work is not ‘too easy’.

A third version of the objection does not claim that safety theories make the response to skepticism too easy. Rather, the charge is that safety theories make *knowledge* too easy. For example, safety theories make it possible to know the world through safe perception. There is typically no requirement that the perceiver herself can explain how she knows, or that she can otherwise reconstruct the knowledge-producing process or circumstances. Here we should heed an insight from James Van Cleve, however – that knowledge of the world is either ‘easy or impossible’.²⁷ We can gloss Van Cleve’s point this way: either knowledge of the world is impossible or near impossible, as skepticism claims, or it is widespread, as common sense claims. If the latter, then a sufficiently anti-skeptical account must explain not only how knowledge is possible, but how it is widespread. In other words, a sufficiently anti-skeptical account must explain how knowledge of the world could be easy. Understood in this light, it is a virtue rather than a vice of safety theories that they do just that.

Nevertheless, I want to argue that there is something right about the ‘that’s too easy’ objection. The way to articulate that thought, however, is to find fault with **SA** itself, and with the way that argument articulates the skeptical problem. If safety-based responses to **SA** fail to adequately address the problem of skepticism, it is because **SA** does not capture the problem adequately in the first place. I will explore arguments to that effect in Parts III and IV.

III. Doubts about SA

Recall the skeptical argument that we have been discussing.

SA

1. I know that *o* only if I know that $\neg h$.
 2. But I don’t know that $\neg h$.
- Therefore,
3. I don’t know that *o*.

As we have seen, premise 1 of the argument is plausible in its own right, and it is further supported by plausible closure principles, for example that knowledge is closed under known entailments. But what about premise 2? Is that premise initially (or pre-theoretically) plausible? Some philosophers have insisted that it is not. On the contrary, these philosophers want to insist, it is initially obvious that I *do* know that I am not a handless brain in a vat. That is not to deny that there might be good reasons in favor of premise 2. The skeptic might very well have good *arguments* that would make 2 plausible. The present point is only that 2 lacks plausibility in the absence of such arguments. Pre-theoretically, i.e., before such arguments are brought to bear, premise 2 seems not plausible at all.²⁸

If this is right, then **SA** is parasitic on other skeptical arguments. In particular, premise 2 of **SA** is plausible only insofar as other skeptical arguments are plausible. Accordingly, we should focus our attention on those other arguments. That's where the real action is.

1. HUME'S ARGUMENT

One such argument is inspired by David Hume.²⁹ The argument begins with the assumption that our beliefs about the external world are at least partly based on how things appear. For example, I believe that I am presently seated at my desk at least partly because that is the way things visually appear to me. But that can't be the whole story, the argument continues. I must also be assuming, at least implicitly, that the way things appear is a good indication of the way things really are. If I were not relying on that assumption, Hume argues, then the fact that things appear to me a certain way would not be a reason to think that they are that way. But now how am I to justify this assumption about the reliability of appearances? How can I know that the way things appear is a good indication of the way things really are? According to Hume, there is no way to justify that assumption. For example, suppose I were to rely on appearances, reasoning that, as far as I can tell, the way things appear to me appear to be a reliable indication of the way things really are. This, of course, would be to argue in a circle, taking for granted the very thing at issue. Here is Hume's argument put more formally.

H

1. All my beliefs about the external world depend for their evidence on both a) the way things appear to me, and b) an assumption that the way things appear to me is a reliable indication of the way things really are.
 2. But the assumption in question can't be justified.
- Therefore,
3. All my beliefs about the external world depend for their evidence on an unjustifiable assumption. (1, 2)
 4. Beliefs that depend for their evidence on an unjustifiable assumption do not count as knowledge.
- Therefore,
5. None of my beliefs about the external world count as knowledge. I don't know anything about the external world. (3, 4)

Clearly, a linchpin of Hume's argument is premise 2: that an assumption regarding the reliability of appearances cannot be justified. In support of premise 2, Hume considers various possibilities for justifying the assumption in question. One consideration that Hume emphasizes is that the assumption is itself a contingent claim about the external world. That is, the assumption claims that sensory appearances are, as a matter of contingent fact, related to the way things are in a particular way. This suggests that the assumption can be justified, if at all, only in the way that contingent claims about the

external world are justified in general – i.e., by relying on the way things appear! But this, of course, would be to argue in a circle, taking for granted the very thing at issue.

Here again is the reasoning in support of 2.

H2

1. All my beliefs about the external world depend for their evidence on both a) the way things appear to me, and b) an assumption that the way things appear to me is a reliable indication of the way things really are.
2. The assumption in question is itself a belief about the external world.
Therefore,
3. The assumption depends on itself for its evidence. (1, 2)
4. Beliefs that that depend on themselves for their evidence can't be justified.
Therefore,
5. The assumption in question can't be justified. (3, 4)

A natural thought is that the assumption that appearances are a reliable guide to reality can be justified in some other way, perhaps by some sort of *a priori* reflection that proceeds independently of appearances. But Hume thinks that this line of reasoning is a dead end. This is because the assumption in question makes a contingent claim about the way things are – it is a matter of contingent fact, and not a matter of necessity, that appearances do or do not reflect the way things really are. But that sort of fact cannot be known through *a priori* reflection. In short, *a priori* reflection gives us knowledge of necessary truths rather than contingent truths.

Finally, we may note that Hume's argument provides independent support for premise 2 of SA. Insofar as my belief that I am not a handless brain in a vat involves a claim about the external world, Hume's argument applies. In sum, I can't know that I am not a handless brain in a vat because my evidence for that belief essentially involves an assumption that appearances to this effect are a reliable indication of the way things are. But there is no way to justify that assumption without going in a circle, and so my belief that I am not a handless brain in a vat depends on inadequate evidence.

2. DESCARTES'S ARGUMENT

A different skeptical argument is inspired by Descartes's *Meditation One*. Here I follow Barry Stroud's influential reconstruction of that argument.³⁰

To understand the argument, consider the claim that one sees a goldfinch in the garden, based on one's observation that the bird is of a particular size and color, and with a tail of a particular shape. Suppose now that a friend challenges one's claim to know, pointing out that woodpeckers also are of that size and color, and also have tails with that shape. As Stroud points out, this seems to be a legitimate challenge to one's claim to know that the bird is a goldfinch. More generally, if one's evidence for one's belief that the bird is a goldfinch is consistent with the possibility that it is in fact a woodpecker, then one does not know on the basis of that evidence that it

is a goldfinch. Based on this sort of reasoning, the skeptic proposes the following plausible principle:

1. A person knows that p on the basis of evidence E , only if E rules out alternative possibilities to p .

Further support for this sort of principle comes from reflection on scientific enquiry. Suppose that there are several competing hypotheses for explaining some phenomenon, and suppose that these various hypotheses are 'live' in the sense that current evidence does not rule them out as possibilities. It would seem that one can not know that one of the hypotheses is true until further evidence rules out the remaining ones. Again, principle 1 above looks plausible.

The second step in the skeptical argument is to point out that there are various possibilities that are inconsistent with what we claim to know about the external world. For example, it is possible that things appear to me visually just as they do now, but that I am actually lying in my bed asleep rather than sitting at my desk awake. It is possible that things appear to Descartes's just as they do, but that he is actually the victim of an evil demon, or that he is a handless brain in a vat. (Here and in step three, we seem to have assumptions in common with **SA**. I will emphasize the difference between the two arguments shortly.)

The third step in the skeptical argument is to claim that our evidence does not in fact rule these possibilities out. The gist of the present claim is something like this: These possibilities are consistent with all the evidence that we have or could have at our disposal. Even if, practically speaking, we don't usually give such possibilities any thought, upon reflection we have no evidence available to us that counts against them, and in favour of our preferred beliefs.

If we put these three claims together we have the materials for a powerful skeptical argument. Here is the argument stated more formally.

D

1. A person knows that p on the basis of evidence E , only if E rules out alternative possibilities to p . (Principle 1 from above.)
2. It is a possibility that I am not sitting at my desk awake, but merely dreaming that I am.
Therefore,
3. I know that I am sitting at my desk only if my evidence rules out the possibility that I am merely dreaming. (1, 2)
4. But my evidence does not rule out this possibility.
Therefore,
5. I do not know that I am sitting at my desk. (3, 4)

And of course, the skeptical argument stated in **D** is supposed to generalize. That is, it is supposed to apply to beliefs about the external world in general. We therefore have:

6. The same line of reasoning can be brought to bear against any belief about the external world.

Therefore,

7. No one knows anything about the external world. (5, 6)

I said that this is a powerful argument. The argument is not powerful in the sense that it is convincing – we shouldn't start to worry that we really don't know anything about the external world. Rather, the argument is powerful in the sense that it is not easy to see where it goes wrong.

One place to look for a weakness in the argument is premise (4). What does it mean, exactly, to say that my evidence does not 'rule out' the possibility that I am dreaming? One way to interpret the premise is along the lines of **SA**: to say that I can't rule out the possibility that I am dreaming is to say that I cannot *know* that I am not dreaming. But this way of reading premise 4, I have suggested, robs the skeptical argument of its intuitive force – it is pre-theoretically implausible that one can't know that one is not dreaming. There are at least two other ways to interpret premise 4 that make the argument more interesting, however. I turn to those now.

One way to understand 'ruling out' a possibility is as follows: A body of evidence *E* rules out a possibility *q* if and only if *E* supports not-*q* in a non-circular way. Here we understand 'support' to express a semantic notion: Evidence *E* supports proposition *p*, in the relevant sense, just in case *E* entails *p* or *p* is probable in relation to *E*. Putting these ideas together, we get the following reading of premise 4 of **D**:

4a. My evidence for my belief that I am sitting at my desk neither entails nor makes probable (in a non-circular way) the proposition that I am not dreaming.

One reason for accepting 4a is the considerations put forward by Hume's argument above. That is, one might think that my evidence for believing that I am sitting at my desk is the way things appear to me, together with my assumption that the way things appear to me is a reliable indication of the way things are. But as Hume's reasoning shows, there is no non-circular way to justify the assumption in question, and therefore no good evidence for either that assumption or further beliefs that are based on it. In particular, my evidence cannot entail or even make probable (in a non-circular way) the proposition that I am not dreaming. Insofar as this is the reasoning behind 4a, argument **D** is parasitic on argument **H**.

There is, however, another way to understand the notion of evidence ruling out alternative possibilities. On this understanding, a body of evidence *E* rules out alternative possibilities to *p* just in case *E* discriminates the state of affairs represented by *p* from alternative states of affairs. For example, hearing my wife coming in the door from work, my auditory experience rules out the possibility that it is my children coming home from school or a burglar coming in through a window. In effect, I have the capacity to 'tell the difference', so to speak, and this is what allows me to know that it is my

wife who has just come in the house. On this understanding of ‘ruling out’, it does seem plausible that my evidence must rule out alternative possibilities in order to ground knowledge. For example, how could I know that my wife has just come home, on the basis of hearing her come through the door, if I could not discriminate that state of affairs from my daughter’s coming through the door? Moreover, premise 4 of argument **D** is plausible on this understanding of ‘ruling out’. We now have,

- 4b. My evidence does not discriminate my sitting at my desk from my merely dreaming that I am sitting at my desk.

One might think, in fact, that this claim is clearly right.

Finally, we may note that either reading of ‘ruling out’ yields an argument that lends support for premise 2 of argument **SA**. That premise claims that I cannot know that I am not a handless brain in a vat. One might think that this is so because my evidence does not support the negation of alternative hypotheses. In particular, it does not support the negation of the hypothesis that I *am* a handless brain in a vat. Alternatively, one might think that 2 is true because my evidence does not discriminate the case where I am not a handless brain in a vat from the case where I am.

Plausibly, neo-Moorean theories have resources for responding to the skeptic’s reasoning, even when reconstructed in these ways. For example, a virtue-theoretic approach can distinguish between reasoning powers, which involve an inference from prior evidence, and perceptual powers, which give rise to beliefs about the world by means of non-inferential processing. This approach would help ground a challenge to premise 1 of **H**, which seems to presuppose an inferential origin for all our beliefs about the world, perceptual beliefs included. Likewise, a virtue-theoretic approach to perception can ground a challenge to argument **D**. Plausibly, perception involves an ability to discriminate certain states of affairs from others. For example, seeing that I am sitting at my desk involves discriminating that state of affairs from my standing at my desk or my sitting in the next room. But as with abilities in general, perceptual abilities are relative to an environment and to a range of appropriate conditions. And now the relevant point is this: I might very well (here and now, under friendly conditions) have the sort of perceptual abilities required for knowledge, even if I would lack those abilities in a very different environment, under unfriendly conditions. That is, plausibly I do have the ability to discriminate my sitting at my desk from alternative possibilities, even if I would lack that ability were I a brain in vat or the victim of a Cartesian demon.

Neo-Mooreans have resources, then, for responding to the skeptical reasoning reconstructed from Descartes and Hume. Whether these resources prove adequate will depend on the details of their development.³¹ Suppose these resources do prove adequate, so that we have adequate answers to the arguments reconstructed in **D** and **H**. Even so, one might think that a different skeptical concern has not yet been addressed.

IV. *The Pyrrhonian Problematic*

Since the 1980s, externalist responses to skepticism have been in ascendancy. Sensitivity theories and safety theories are cases in point. For present purposes, we can understand the distinction between internalism and externalism as follows. First, all theories require something akin to reliability for knowledge: a belief counts as knowledge only if it is reliably formed, or formed by a reliable power, or some such thing. Internalists add a further condition on knowledge: that the knower justifiably believes that her belief is reliably formed. Externalist theories omit any such further condition. Put differently, externalist theories require only *de facto* reliability, whereas internalist theories require a perspective on one's reliability.

Externalist theories, we have seen, have ample resources for addressing skeptical arguments. And yet externalist responses to skepticism have left many philosophers dissatisfied. The recurrent protest is that externalist theories do not get to the bottom of skeptical concerns. In recent years this protest has crystallized as follows: Externalist theories (sensitivity theories and safety theories included) do not adequately address Pyrrhonian concerns.³² In this final part of the article I want to explore this persistent theme. In section 1 I will reconstruct the Pyrrhonian problematic. In Section 2 I will consider a strategy for resolving it.

1. THE PROBLEMATIC

The Pyrrhonian problematic begins with a familiar tri-lemma. All knowledge, says the skeptic, must be grounded in good reasons. But this simple requirement ensures that all grounds for knowledge will be inadequate. For either a) one's reasons will go on in an infinite regress; b) they will come back in a circle; or c) they will end arbitrarily. But none of these outcomes is satisfactory – none provides knowledge with grounding in *good* reasons. And therefore, the skeptic concludes, knowledge is impossible.

Externalists, we have seen, have an answer to the problem. Namely, they deny the first assumption of the skeptic's reasoning – that all knowledge must be grounded in good reasons. *Some* reliable cognition involves grounding in good reasons. But not *all* reliable cognition does. Perception, for example, might be highly reliable, but involve nothing by way of inference from good reasons. Introspection is plausibly like that as well, as is logical intuition and memory. And perhaps there are other kinds of reliable, non-inferential processes as well. The picture of knowledge that results is foundationalist in structure: A foundation of non-inferential knowledge, produced by non-inferential but reliable processes, provides the basis for further knowledge, produced by reliable inferences from the foundations. On this account, the skeptic is just wrong to think that all knowledge must be grounded in good reasons. Put another way, the skeptic is just wrong to think that all knowledge-producing processes are reasoning processes.

As noted above, however, many philosophers have not been satisfied with this straightforward externalist response. They give voice to a different concern that, they believe, lies deeper in the Pyrrhonian problematic. Here is Laurence Bonjour:

[A]lthough the foregoing dialectical motive for externalism is abundantly clear, it is nevertheless far from obvious that what results is a plausible account of epistemic justification. . . . And indeed such views may well be suspected of being merely *ad hoc* in relation to the difficulties arising from the epistemic regress problem. . . . How can the fact that a belief is reliably produced (or indeed any sort of fact that makes a belief likely to be true) make my acceptance of that belief rational and responsible when that fact itself is entirely unavailable to me? (26–7)

Barry Stroud resists the idea that knowledge and justification require more than *de facto* reliability. Nevertheless, Stroud argues, the externalist fails to give us a satisfying *understanding* of our knowledge. Stroud writes,

If I ask of my own knowledge of the world around me how it is possible, I can explain it along ‘externalist’ lines by showing that it is a set of beliefs I have acquired through perception by means of belief-forming mechanisms which are reliable . . . As a good externalist . . . I think I do know or have good reasons to believe my theory; I believe that I fulfil the conditions which that very theory says are sufficient for knowing or having good reason to believe it. Do I now have a satisfactory understanding of my knowledge of the world? Have I answered to my own satisfaction the philosophical question of how my knowledge of the world is possible? I want to say ‘No’. (301–2)

Sosa is sympathetic, and accordingly endorses the following passage from Sextus.

Let us imagine that some people are looking for gold in a dark room full of treasures . . . [N]one of them will be persuaded that he has hit upon gold even if he *has* in fact hit upon it. In the same way, the crowd of philosophers has come into the world, as into a vast house, in search of truth. But it is reasonable that the man who grasps the truth should doubt whether he has been successful. (7:259)

Sosa comments:

Most would not disdain the good fortune of those who strike it rich in the dark, but it is no doubt a lesser state than that of finding gold guided by good eyesight in clear light. Enlightened discovery is more admirable than is any comparable luck that may reward groping in the dark. For one thing, enlightened discovery is success attributable to the agent; luck in the dark is not. (‘How to Resolve the Pyrrhonian Problematic’ 231)

According to Sosa, enlightened discovery requires a perspective on the reliability of one’s cognitive faculties. Real knowledge requires that one’s belief be reliably produced, but also that one sees that one’s belief is reliably produced.

Here it might seem that we are stuck with a clash of intuitions. On the one hand, we have the externalist's intuition that knowledge of the world is possible, and indeed that it is widespread. Hence the externalist rejects the skeptical idea that knowledge requires evidence of one's reliability, seeing that any such requirement leads soon enough to skepticism via infinite regress or vicious circle. On the other hand, we have the internalist intuition that *de facto* reliability is not enough – that knowledge requires exactly the perspective that the externalist rejects. At the very least, an adequate understanding of our knowledge requires this. Not all of these intuitions can be correct. Something has to give.

One strategy for breaking the impasse is to reject the internalist requirement on independent grounds. For example, one might note that the requirement issues in a completely general skepticism, and not merely skepticism about the world. If all knowledge requires a perspective on one's reliability, and if such perspective is impossible without infinite regress or vicious circle, then any knowledge whatsoever is impossible, even the knowledge that I think or that I exist. Even an external world skeptic should balk at this result. A related strategy is to argue that the internalist requirement is incoherent. Boiled down to its bare bones, the requirement is to vindicate the reliability of our cognitive resources, while not allowing (on pain of circularity) that any resources be brought to the task. In effect, it asks for evidence of reliability, while at the same time disallowing any evidence that one could possibly have. Once the internalist requirement is seen in this light, anyone in their right mind should balk.

Finally, one might try to accommodate the internalist's intuition without accepting it outright. A version of this third strategy is developed in detail by Sosa.

2. SOSA'S ACCOMODATIONIST STRATEGY³³

Central to Sosa's strategy is to distinguish between two kinds of knowledge. 'Animal knowledge', he tells us, can be defined by minimal conditions regarding *de facto* reliability or something in the neighborhood. 'Reflective knowledge' is of a higher grade, requiring the sort of perspective on reliability that the internalist holds dear. Endorsing this distinction, Sosa argues, does not commit one to the position that our ordinary knowledge language is ambiguous. Rather, the point is to acknowledge a class of knowledge that is of especially high quality and value, whether ordinary language recognizes the kind or not. With that distinction in hand, we can then interpret the internalist's requirement as a demand on high-quality, reflective knowledge.

That cannot be the whole story, however, since the Pyrrhonian argument will now go through for reflective knowledge. So long as the demand for a perspective is fully general, the Pyrrhonian's reasoning will apply, and no matter what label we put on the desirable epistemic state. This leads us to a second important element of Sosa's strategy: to conceive higher-quality

knowledge as a matter of degree. The Pyrrhonian reasoning takes hold precisely because the requirement for a perspective is conceived as fully general. It is that fully general requirement that issues in a regress or circle. But we can acknowledge the value of an epistemic perspective, Sosa argues, even when that perspective is not fully general. That value lies, in part, in the coherence that such a perspective confers, and on attendant understanding. But coherence and understanding come in degrees. More is better than less, and some is better than none at all.

These last considerations apply to Stroud's complaint as well. A fully general (yet non-circular) understanding of our knowledge is indeed impossible. That goal, in fact, might very well be incoherent. But one can nevertheless admit the value of more understanding over less. And even full generality is in principle possible, so long as one does not also eschew circularity.

Sosa's strategy for accommodating the internalist's intuition is now in place: acknowledge the value of an epistemic perspective, i.e., a perspective on the sources of one's beliefs, and on the reliability of those sources. But acknowledge that value in the context of a more general value: that of explanatory coherence and the understanding that such coherence brings with it. In this way we avoid any requirement on knowledge, or even on reflective knowledge, that issues in a general skepticism. At the same time, we give the skeptic his due: we acknowledge the superiority of knowledge with understanding over knowledge without understanding. Or what is the same, we acknowledge the inferiority of the latter in comparison with the former.

Acknowledgement

Thanks to Duncan Pritchard and Ernest Sosa for comments on an earlier draft and other relevant materials.

Short Biography

John Greco is the Leonard and Elizabeth Eslick Chair in Philosophy at Saint Louis University. He received his Ph.D. from Brown University in 1989. Areas of research include: epistemology, skepticism, Hume, and Reid. Recent publications include: *The Oxford Handbook of Skepticism* (Oxford 2008); *Sosa and his Critics* (Blackwell, 2004); *Putting Sceptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry* (Cambridge University Press, 2000); 'Virtue, Luck and the Pyrrhonian Problematic', *Philosophical Studies* (2006); 'Virtue Epistemology', *The Stanford Encyclopedia of Philosophy* (2004); 'Reid's Reply to the Skeptic', in *The Cambridge Companion to Reid* (2004); 'Knowledge as Credit for True Belief', in *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford University Press, 2003); and 'How to Reid Moore', *Philosophical Quarterly* (2002).

Notes

* Correspondence address: Philosophy Department, Saint Louis University, 221 N. Grand Blvd, St. Louis, Missouri 63103, USA. Email: jgreco2@slu.edu.

¹ For a useful collection of that literature see DeRose and Warfield. For excellent discussions see DeRose, 'Solving the Skeptical Problem'; Pritchard, 'Recent Work'; 'Sensitivity'; Cohen, 'Ascriber Contextualism'.

² For example, see Nozick; Dretske.

³ For example, see Cohen, 'Knowledge and Context'; 'How to be a Fallibilist'; DeRose, 'Solving the Skeptical Problem'; Lewis.

⁴ But see David and Warfield.

⁵ Nozick; Dretske. See also Becker.

⁶ This serves as a first approximation. Various qualifications have been proposed, but since they are not important for present purposes I will ignore them here.

⁷ For more on the proper methodology for responding to skepticism, see Greco, *Putting Skeptics in their Place*, especially chapter 1.

⁸ DeRose, 'Solving the Skeptical Problem' 27–9.

⁹ See Becker.

¹⁰ See Cohen, 'Knowledge and Context'; 'How to be a Fallibilist'. For developments of the view, see DeRose, 'Solving the Skeptical Problem'; Lewis.

¹¹ Thus Hume writes, 'The *intense* view of these manifold contradictions and imperfections in human reason has so wrought upon me, and heated my brain, that I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another . . . Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices for that purpose, and cures me of this philosophical melancholy and dilerium . . . I dine, I play a game of back-gammon, I converse, and am merry with my friends; and when after three or four hour's amusement, I would return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter into them any farther' (*Treatise of Human Nature* Book I, Part IV, Section VII).

¹² See DeRose, 'Single Scoreboard Semantics'.

¹³ This is the view that DeRose prefers.

¹⁴ Cf. Sosa, 'How to Defeat Opposition to Moore'.

¹⁵ For example, see Sosa, 'Skepticism and Contextualism'; 'How to Defeat Opposition to Moore'; Pryor, 'Skeptic and the Dogmatist'; 'What's Wrong'; Williamson, *Knowledge and Its Limits*; 'Scepticism and Evidence'; Greco, 'How to Reid Moore'; Pritchard, 'Resurrecting the Moorean Response'; *Epistemic Luck*.

¹⁶ See Pryor 'Skeptic and the Dogmatist'; 'What's Wrong'.

¹⁷ See Pryor 'Skeptic and the Dogmatist'.

¹⁸ See Sosa 'Skepticism and Contextualism'; 'How Must Knowledge be Modally Related'; 'How to Defeat Opposition to Moore'. Sosa's account is developed most recently in *Virtue Epistemology*. See also Williamson, *Knowledge and Its Limits*; 'Scepticism and Evidence'; Greco, 'How to Reid Moore'; Pritchard 'Resurrecting the Moorean Response'; *Epistemic Luck*. Again, this serves as a first approximation. I will continue to ignore qualifications that are unimportant for present purposes.

¹⁹ *Weak Safety* is not without its problems however. See Greco, 'Knowledge as Credit'.

²⁰ Pritchard, 'Sensitivity, Safety and Anti-Luck Epistemology'; 'Knowledge, Luck and Lotteries'.

²¹ See Sosa, 'How Must Knowledge be Modally Related'.

²² See Greco, 'Knowledge as Credit'.

²³ Duncan Pritchard has argued against wedding a safety condition with a virtue-theoretic or agent reliabilist condition. See Pritchard, *Epistemic Luck*; Greco, 'Worries about Pritchard's Safety'.

²⁴ See Sosa, 'Skepticism and Contextualism'; 'How to Defeat Opposition'; *Virtue Epistemology* ch. 2.

²⁵ See Pritchard, *Epistemic Luck* 79–86.

²⁶ This is true even of contextualist responses to skepticism. In effect, contextualist theories claim that premise 2 of SA is false relative to ordinary contexts.

²⁷ Van Cleve.

- ²⁸ See Greco, *Putting Sceptics in their Place* especially ch. 2. See also Sosa, 'How Must Knowledge be Modally Related'; 'How to Defeat Opposition'; 'Skepticism and Contextualism'.
- ²⁹ See Greco, *Putting Sceptics in their Place* ch. 2.
- ³⁰ See Stroud. See also Bruekner; Cohen, 'Two Kinds'; Vogel, 'Varieties of Skepticism'; Pritchard, 'Structure of Sceptical Arguments'.
- ³¹ See Greco, *Putting Sceptics in their Place*; Sosa, *Virtue Epistemology*.
- ³² For example see Bonjour; Lammenranta.
- ³³ This section is reconstructed from Sosa, 'How to Resolve the Pyrrhonian Problematic'; 'Replies' especially 290–304, 309–15; *Virtue Epistemology* ch. 6. For a less accommodating approach, see Greco, 'Virtue, Luck and the Pyrrhonian Problematic'. For a more accommodating approach, see Lammenranta.

Works Cited

- Becker, K. *Epistemology Modalized*. Oxford: Routledge, 2007.
- Bonjour, L. 'A Version of Internalist Foundationalism'. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. L. Bonjour and E. Sosa. Oxford: Blackwell, 2003.
- Bruekner, A. 'The Structure of the Skeptical Argument'. *Philosophy and Phenomenological Research* 54 (1994): 827–35.
- Cohen, S. 'Ascriber Contextualism'. *Oxford Handbook of Skepticism*. Ed. J. Greco. New York, NY: Oxford UP, 2008.
- . 'How to be a Fallibilist'. *Philosophical Perspectives* 2 (1988): 91–123.
- . 'Knowledge and Context'. *Journal of Philosophy* 83 (1986): 574–83.
- . 'Two Kinds of Skeptical Argument'. *Philosophy and Phenomenological Research* 58 (1998): 143–59.
- David, M. and T. Warfield. 'Knowledge-Closure and Skepticism'. *Epistemology: New Philosophical Essays*. Ed. Q. Smith. Oxford: Oxford University Press, forthcoming.
- DeRose, K. 'Single Scoreboard Semantics'. *Philosophical Studies* 119.1–2 (2004): 1–21.
- . 'Solving the Skeptical Problem'. *Philosophical Review* 104 (1995) 1–52. Reprinted in *Skepticism: A Contemporary Reader*. Eds. K. DeRose and T. Warfield. New York, NY: Oxford UP, 1999.
- and T. Warfield, eds. *Skepticism: A Contemporary Reader*. New York, NY: Oxford UP, 1999.
- Dretske, F. 'Conclusive Reasons'. *Australasian Journal of Philosophy* 49 (1971): 1–22.
- Greco, J. 'How to Reid Moore'. *Philosophical Quarterly* 52.209 (2002): 544–63.
- . 'Knowledge as Credit for True Belief'. *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Eds M. DePaul and L. Zagzebski. Oxford: Clarendon Press, 2003.
- , ed. *Oxford Handbook of Skepticism*. New York, NY: Oxford UP, 2008.
- . *Putting Sceptics in their Place: The Nature of Skeptical Arguments and their Role in Philosophical Inquiry*. New York, NY: Cambridge UP, 2000.
- , ed. *Sosa and His Critics*. Malden, MA: Blackwell Publishing, 2004.
- . 'Virtue, Luck and the Pyrrhonian Problematic'. *Philosophical Studies* 130 (2006): 9–34.
- . 'Worries about Pritchard's Safety'. *Synthese* (forthcoming).
- Hawthorne, J. *Knowledge and Lotteries*. Oxford: Oxford UP, 2004.
- Lammenranta, M. 'The Pyrrhonian Problematic'. *Oxford Handbook of Skepticism*. Ed. J. Greco. New York, NY: Oxford UP, 2008.
- Lewis, D. 'Elusive Knowledge'. *Australian Journal of Philosophy* 74 (1996): 549–67.
- Nozick, R. *Philosophical Explanations*. Cambridge, MA: Harvard UP, 1981.
- Pritchard, D. *Epistemic Luck*. New York, NY: Oxford UP, 2005.
- . 'Knowledge, Luck and Lotteries'. *New Waves in Epistemology*. Eds. V. Hendricks and D. Pritchard. Aldershot: Ashgate, 2006.
- . 'Recent Work on Radical Skepticism'. *American Philosophical Quarterly* 39.3 (2002): 215–57.
- . 'Resurrecting the Moorean Response to Scepticism'. *International Journal of Philosophical Studies* 10 (2002): 283–307.
- . 'Sensitivity, Safety and Anti-Luck Epistemology'. *Oxford Handbook of Skepticism*. Ed. J. Greco. New York, NY: Oxford UP, 2008.
- . 'The Structure of Sceptical Arguments'. *Philosophical Quarterly* 55 (2005): 37–52.

- Pryor, J. 'The Skeptic and the Dogmatist'. *Noûs* 34.4 (2000): 517–49.
- . 'What's Wrong with Moore's Argument?' *Philosophical Issues* 14 (2004): 349–78.
- Sosa, E. 'How Must Knowledge be Modally Related to What is Known'. *Philosophical Topics* 26 (1999): 373–84.
- . 'How to Defeat Opposition to Moore'. *Philosophical Perspectives* 13 (1999): 141–53.
- . 'How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes'. *Philosophical Studies* 85.2–3 (1997): 229–49.
- . 'Replies'. *Sosa and his Critics*. Ed. J. Greco. Oxford: Blackwell Publishers, 2004.
- . 'Skepticism and Contextualism'. *Philosophical Issues* 10 (2000): 1–18.
- . *A Virtue Epistemology*. Oxford: Oxford UP, 2007.
- Stroud, B. 'Scepticism, "Externalism", and the Goal of Epistemology'. *Proceedings of the Aristotelian Society* Suppl. vol. 68 (1994): 290–307.
- Van Cleve, J. 'Is Knowledge Easy – Or Impossible? Externalism as the Only Alternative to Skepticism'. *The Sceptics*. Ed. S. Luper. Aldershot: Ashgate Publishing, 2003.
- Vogel, J. 'The New Relevant Alternatives Theory'. *Philosophical Perspectives* 13 (1999): 155–80.
- . 'Tracking, Closure, and Inductive Knowledge'. *The Possibility of Knowledge*. Ed. S. Luper-Foy. Totowa, NJ: Rowman and Littlefield, 1987.
- . 'Varieties of Skepticism'. *Philosophy and Phenomenology Research* 68 (2004): 1–37.
- Williamson, T. *Knowledge and Its Limits*. New York, NY: Oxford UP, 2000.
- . 'Scepticism and Evidence'. *Philosophy and Phenomenological Research* 60 (2000): 613–28.