

RECENT WORK ON RADICAL SKEPTICISM

Duncan Pritchard

0. INTRODUCTION

This discussion surveys recent developments in the treatment of the epistemological problem of skepticism. These are arguments which attack our *knowledge* of certain truths rather than, say, our belief in the existence of certain entities. In particular, this article focuses on the *radical* versions of these skeptical arguments, arguments which purport to show that knowledge is, for the most part, impossible, rather than just that we lack knowledge in a particular discourse.

Although most of the key recent developments in this area have taken place since the late 1980s and early 1990s, it is necessary to also discuss some of the movements that have developed since 1970 in order to give these recent developments the necessary setting. The date of 1970 is dictated by the publication in that year of Fred Dretske's seminal article "Epistemic Operators," which both pushed a "relevant alternatives" theory of knowledge to the fore of discussion and also brought into focus one possible line of argument against the so-called "closure" principle for knowledge. In so doing, it provided one of the main sources of response to the emergent

interest in the "infallibilist" motivation for radical skepticism, as expressed by, for example, Peter Unger (1971; 1975). As argued below, this relevant alternatives model was developed along two extremely influential lines in subsequent discussion. On the one hand, we find the first wave of relevant alternatives theorists that includes Dretske (1971) himself and Robert Nozick (1981), who both advocate a modal theory of knowledge that results in the rejection of the principle that knowledge is "closed" under known entailment (the so-called "closure" principle for knowledge). On the other hand, we have the later wave of relevant alternatives theorists, and in particular Stewart Cohen (1986; 1987; 1988; 1990; 1991; 1998a; 1999; 2000a), who opt for an explicitly contextualist thesis that retains closure. This type of theory in turn evolved into the kind of sophisticated contextualism (henceforth, "semantic" contextualism) advocated by such figures as Keith DeRose (1995), David Lewis (1996) and Cohen himself that is currently in vogue in the literature.

Due to its dominance in the recent debate, this relevant alternatives movement in contemporary epistemology will be the focus for this discussion. Nevertheless, this

discussion will also cover, to a lesser extent, certain other responses to skepticism that have cropped up in the recent literature. For example, it will examine the work of the so-called “new skeptics.”¹ These are philosophers, such as Unger (1971; 1975), Barry Stroud (1984; 1989; 1994; 1996), Thomas Nagel (1986), and Richard Fumerton (1990; 1995) who have contributed to a renewed sympathy toward skeptical arguments, in contrast to earlier post-war treatments of skepticism which tended to regard it as simply a by-product of a particularly dubious philosophical language game.

We will also be examining a different brand of anti-skeptical theorizing—which takes its cue from Wittgenstein’s remarks on “hinge” propositions in *On Certainty*—that has more in common with the new skeptics than those who endorse a theory along relevant alternatives lines. The kind of philosophers who fall into this (less clearly defined) camp include Peter Strawson (1985), Crispin Wright (1985; 1991; 2000a), Hilary Putnam (1992), and Avrum Stroll (1994). Relatedly in this respect, we will also be discussing the rather different contextualist thesis (henceforth, “inferential” contextualism) that is advanced by Michael Williams (1991) and which takes its cue from both relevant alternatives theories and from the Wittgensteinian stimulus.

Finally, we will be discussing a new camp of anti-skeptical theorists who have recently emerged, the “neo-Mooreans.” Proponents of this kind of view, including Mark Sainsbury (1997), Ernest Sosa (1999a; 2000a), Timothy Williamson (2000a; 2000c, chapter 8), and Duncan Pritchard (2001b; 2001e; 2002a; 2002b; 2002c), advocate a non-contextualist anti-skeptical theory that parallels, in a key respect, G. E. Moore’s notorious “common-sense” reply to the skeptic. It is argued that the plausibility of

such a thesis rests almost entirely upon the correctness of a certain construal of epistemological externalism.

In the interests of brevity, however, there are a number of important recent contributions to the debate on skepticism that will not be discussed here. In particular, we will be taking for granted a certain form of the skeptical argument that makes essential use of the closure principle. In so doing, we will be setting aside concerns discussed by Anthony Brueckner (1994b) and Stewart Cohen (1998c) that skeptical arguments might instead rest upon a different epistemic principle entirely, what they term an “underdetermination” principle.² Relatedly, we will also be taking a very modern Cartesian understanding of the skeptical argument for granted, and therefore ignoring the emergence of renewed interest in classical Pyrrhonian skeptical arguments by such figures as Williams (1988b) and Robert Fogelin (1994).³ Finally, attention will be confined to those anti-skeptical theories of knowledge which are genuinely capable (given, of course, that they are correct) of meeting the skeptical challenge. Accordingly, we will not be discussing the so-called “semantic” responses to skepticism that have been drawn from work by, for example, Putnam (1981) and Donald Davidson (1983). Though influential, such arguments only show that, at best, *some* radical skeptical hypotheses must be false.⁴ Unfortunately, however, the skeptic needs only one of her hypotheses to go through in order to achieve her skeptical result, and thus we cannot acquiesce in this kind of nominal anti-skeptical conclusion (although, obviously, the nominal anti-skeptical conclusion, if it can be adequately supported, is extremely interesting in other respects).⁵

In §1, the template radical skeptical argument that will be considered is put forward and the kinds of claims that it

depends upon are highlighted, along with the challenge that it poses. In §2, the basic relevant alternatives approach to skeptical arguments, and the attack on the closure principle for knowledge that is associated with it, is discussed. The contrast between the relevant alternatives approach and infallibilism is also considered. Sections 3 and 4 look at the first main development of the relevant alternatives view along modal lines and consider the beginnings of the contextualist challenge to this approach, along with some of the other main lines of critique. Sections 5 and 6 scrutinize the semantic contextualist thesis and consider some objections that might be raised against it. Section 7 examines the various hinge proposition theses that have been put forward as regards skepticism, and then contrasts one refined version of the hinge proposition thesis with an inferential rendering of the basic contextualist approach that runs along superficially similar lines. Section 8 looks at the neo-Moorean response to skepticism and considers some problems that it faces. Section 9 considers the impact of the “new skeptics” to current epistemological debate, and examines how this position relates to the epistemological externalism/internalism distinction. Finally, §10 offers some brief concluding remarks on the prospects for future discussion in this area.

1. THE SKEPTICAL PARADOX

Following DeRose (1995), Sosa (1999a), Jonathan Vogel (1999), and others, we will understand the radical skeptical paradox to consist in the joint incompatibility of three claims, each of which appears, on the surface of things and taken individually, to be perfectly in order. Take ‘SH’ to refer to some suitable skeptical hypothesis, such

as the hypothesis that one might be a “brain in a vat” being “fed” one’s experiences by computers (Putnam 1981, chapter 1; Pollock 1986, 1–3), and take ‘O’ to be some “ordinary” proposition that one would typically take oneself to know and which entails the falsity of the skeptical hypothesis in question (such as that one has two hands). We can then formulate these three incompatible claims as follows:

- (1) I know O.
- (2) I do not know not-SH.
- (3) If I do not know not-SH, then I do not know O.⁶

Clearly, these three propositions cannot all be true, and yet they do seem to possess at least an initial degree of plausibility. If one knows anything, then one ought to be able to know the ordinary proposition at issue in line (1). Furthermore, as line (2) makes explicit, the very sort of propositions which one seems unable, in principle, to know are the denials of radical skeptical hypotheses, since these hypotheses concern scenarios which are phenomenologically indistinguishable from everyday life. Finally, it seems relatively uncontroversial to argue with line (3), that in order to know these ordinary propositions one must be able to know that the relevant skeptical hypotheses are false, since they seem to constitute defeaters for our everyday knowledge. One cannot consistently endorse all three of these claims, however. We are thus in a bind.

Of course, the skeptic offers a very simple way out of this puzzle, which is to deny, on the basis of lines (2) and (3), that we ever have knowledge of the ordinary propositions at issue in line (1). That is, the skeptic argues as follows:

- (S1) I do not know not-SH.
- (S2) If I do not know not-SH, then I do not know O.

Hence:

(SC) I do not know O.

For example, a skeptical argument which employed the brain-in-a-vat hypothesis might well run as follows:

(S1*) I do not know that I am not a brain-in-a-vat.

(S2*) If I do not know that I am not a brain-in-a-vat, then I do not know O.

Hence:

(SC*) I do not know O.

Clearly, however, this radical skeptical suggestion regarding how we should respond to these three incompatible claims is less of a proposal than a *reductio* of epistemological theorizing. Since one can repeat such an argument with any agent and just about any ordinary proposition (though one might have to vary the skeptical hypothesis at issue to suit), it follows that we are unable to know anything much at all, and this seems like an intolerable conclusion to draw.

If we are to evade skepticism, then, we are going to have to motivate one (or more) of the following three claims. First, despite appearances, we do (or at least *can*) know the denials of radical skeptical hypotheses after all. Second, despite appearances, it does not follow from the fact that we lack knowledge of the denials of radical skeptical hypotheses that we thereby lack knowledge of ordinary propositions as well. Third, despite appearances (and granted the truth of *modus ponens*⁷), these three claims *are* consistent.

2. RELEVANT ALTERNATIVES, INFALLIBILISM, AND CLOSURE

Of the three anti-skeptical strategies listed, the second looks, *prima facie*, to be the most promising. The reason why this seems to be the weakest element of the

skeptical argument is that, although it is at first pass intuitive, on reflection it is far from immediately obvious that our knowledge of everyday propositions should be dependent upon anti-skeptical knowledge in this fashion. One response to the problem of skepticism has thus been to deny this premise in the skeptical argument by arguing that one can perfectly well know everyday propositions while failing to know the denials of anti-skeptical hypotheses such as the brain-in-a-vat hypothesis.

One motivation for this line of argument has been to argue that skeptical error-possibilities are just not *relevant* to everyday knowledge in the way that everyday error-possibilities are. After all, we do not ordinarily demand that agents should rule out skeptical error-possibilities before we ascribe them knowledge. This *relevant alternatives* line of argument, which has its roots in remarks by J. L. Austin (1961), has been given an explicit expression in the work of Dretske (1970).⁸ As Dretske is aware, however, simply denying (S2) of the skeptical argument on these grounds is not enough, rather one needs to also engage with the epistemological theses that underlie this premise and offer a fully fledged account of what this notion of epistemological relevance involves.

One epistemological thesis that is often thought to provide support for (S2) is that of *infallibilism*. This is the thesis that, roughly, for an agent to know a proposition that agent must be able to eliminate *all* error-possibilities associated with that proposition. Provided that one is willing to make the plausible move of construing “eliminate” here in terms of the ability to know the negation of, then one straightforwardly gets the requisite link between infallibilism and (S2) since the skeptical hypothesis in question (whichever skeptical hypothesis it is) will clearly be an error possibility which must be known to be false

if the agent is to have knowledge of the ordinary proposition at issue. Accordingly, an inability to know the denial of the skeptical hypothesis will suffice to ensure that the agent lacks knowledge of the ordinary proposition, just as (S2) says. In effect, infallibilism is the opposing thesis to the relevant alternatives line because it counts *every* alternative as being relevant.

Although infallibilism may seem to be an obviously false epistemological thesis, a persuasive case can be made in its defense. In particular, Unger (1971; 1975) has been a prominent defender of a version of infallibilism (although in more recent work, such as Unger (1984; 1986) he has moved toward a thesis which is more in line with semantic contextualism, as discussed below). In these early works Unger argued that “knowledge” is an “absolute term” like “flat” or “empty.” According to Unger, what is interesting about absolute terms is that they are never really satisfied, although we often talk as if they are. So, for example, nothing is ever *really* flat or *really* empty because, respectively, no surface is ever completely free of friction and no container could ever be a vacuum. Accordingly, even though we might loosely talk of Holland’s “flat” roads or John’s “empty” fridge, reflection indicates to us that such assertions are, in fact, false (Holland’s roads have some bumps on them, however small, and John’s fridge, while empty of food, is “full” of air, not to mention refrigerator parts). Similarly, Unger’s point is that what the skeptic is responding to in her arguments is the fact that, strictly speaking, nothing is *every really* known because to be *really* known the agent would have to rule out every possibility of error and this is an impossible hurdle to clear (at least for a non-omniscient being). So although we might talk of knowing lots of things, reflection indicates to us, as it does with our use of “flat”

and “empty,” that our claims to know are all, in fact, false.⁹

We will return to consider infallibilism again below. In the meantime, however, we can set this thesis to one side because there is a (logically) weaker thesis that would also suffice to support (S2). Accordingly, so long as we are able to deny the weaker thesis then we can get a rejection of infallibilism by default. This weaker thesis is the principle that knowledge is “closed” under known entailment, or the “closure” principle for short. Roughly, this principle states that if an agent knows a proposition (such as that she is currently seated), and knows that this proposition entails a second proposition (such as that she is not a brain-in-a-vat), then she also knows the second proposition. More formally:

The “Closure” Principle

For all S, ϕ , ψ , if S knows that ϕ , and S knows that ϕ entails ψ , then S knows that ψ .¹⁰

Whereas infallibilism supports (S2) by demanding that an agent should be able to know the denials of *all* error-possibilities, closure merely demands that the agent know the denials of those error-possibilities that are known to be logical consequences of what one knows. For example, if one knows the ordinary proposition that one is currently seated, and one further knows that if one is seated then one is not a brain-in-a-vat, then one must also know that one is not a brain-in-a-vat. Conversely, if one does not know that one is not a brain-in-a-vat, then, given that one knows the entailment in question (which ought to be uncontroversial), one thereby lacks knowledge of the ordinary proposition in question, just as (S2) says. And note that, unlike (S2), the plausibility of closure is not merely *prima facie*. We reason in conformity with closure all the time in cases where we gain knowledge of previously unknown propositions via knowledge of

other propositions and the relevant entailment. Indeed, closure is in this respect far more compelling than infallibilism, since what credibility the latter thesis has is gained by philosophical argument rather than by *prima facie* reflection on our actual epistemic practice. The theoretical burden imposed upon anyone who advocates the denial of (S2) is thus very strong, since it requires a principled rejection of the intuitive principle of closure.¹¹

The standard proposal put forward to support the denial of closure has been some variation of the original relevant alternatives model advanced by Dretske (1970; 1971; 1981). Essentially, the idea is to claim that knowledge only transfers across known entailments where the entailments in question are “relevant.”¹² Thus, knowledge that one is sitting down will transfer across a known entailment to the relevant proposition that one is not standing up, but it won’t transfer across a known entailment to the irrelevant proposition that one is not a brain-in-a-vat. Accordingly, the link between ordinary knowledge and anti-skeptical knowledge required by the skeptic is severed and ordinary knowledge is secured. Dretske himself puts the point as follows:

The general point may be put this way: there are certain presuppositions associated with a statement. These presuppositions, although their truth is entailed by the truth of the statement, are not part of what is *operated on* when we operate on the statement with one of our epistemic operators. The epistemic operators do not *penetrate to* these presuppositions (Dretske 1970, p. 1014).

In effect, what Dretske is arguing here is that in everyday contexts an agent’s acquisition of knowledge of the propositions at issue in that context *presupposes* the falsity of certain irrelevant error-possibilities. That they are taken for granted is, for Dretske, entirely legitimate (i.e., he rejects

infallibilism). Nevertheless, the negations of these error-possibilities are often entailed by what is known in that context and thus, if closure held, it would follow that an agent could come to have knowledge of what is presupposed in her knowledge simply by knowing the relevant entailment. It is this that Dretske objects to, arguing that one’s epistemic position regarding the antecedent proposition will not transfer to the consequent proposition where the consequent proposition has performed this “presuppositional” role.

An example will help clarify matters. Consider the following two propositions (adapted from ones adduced by Dretske [1970]):

- (P) The animals in the pen are zebras.
- (Q) The animals in the pen are not mules cleverly disguised to look like zebras.

Dretske argues that in normal circumstances one can come to know (P) without making any special checks to ensure that the irrelevant error-possibility at issue in (Q) is false. Instead, all the agent needs to do is have evidence that eliminates relevant error-possibilities (such as, for example, evidence to support her belief that it is the zebra enclosure and not the ape enclosure that she is looking at). This is fortunate, because if we demand that the agent must rule out the kind of error-possibility at issue in (Q) (and thus, one might reasonably assume, know [Q]) before she can know (P), then we will end up setting the requirement for knowledge at a very high level. Indeed, it will be highly unlikely that your average agent would be able to know a proposition like (P) if this demand is made, because the average agent would not be able to tell a zebra apart from a cleverly disguised mule. Nevertheless, Dretske acknowledges that the agent’s knowledge of (P) presupposes that the error-possibility at issue in (Q) is false. Here is the crux,

however. If we allow closure to stand, then it will follow from the agent's knowledge of (P), and her knowledge of the entailment from (P) to (Q), that she thereby knows (Q) also, even though we have already granted that the agent in question is not in a position to be able to know such a thing. Dretske puts the point as follows:

If you are tempted to say [*that the agent does know (Q)*], think for a moment about the reasons that you have, what evidence you can produce in favor of this claim. The evidence you *had* for thinking them zebras has been effectively neutralised, since it does not count toward their *not* being mules cleverly disguised. Have you checked with the zoo authorities? Did you examine the animals closely enough to detect such a fraud (Dretske 1970, p. 1016)?

Dretske thus concludes that we should instead allow that an agent might be able to know (P) while failing to know (Q), and thus, given that the entailment is known, that closure fails.¹³

This is certainly a very compelling argument, and it does at the very least offer a *prima facie* case against closure. The job is not quite done, however, because we also need to be given an account of knowledge which will flesh out this notion of relevance. After all, we have strong intuitions that our epistemic concepts do license closure. It is to this end that Dretske (1971) went on to develop his "modal" account of knowledge, an account which was adapted and supplemented by Nozick (1981).

3. MODAL EPISTEMOLOGY AND THE DENIAL OF CLOSURE

So Dretske needs a theory of knowledge which, while being plausible, can also explain how we can know everyday propositions while failing to know the denials of skeptical hypotheses, even in cases where

we know that the everyday proposition in question entails the denial of the relevant skeptical hypothesis. Given these arduous demands, his proposal is, to say the least, ingenious. In essence, what Dretske does is to adduce the following modal condition on knowledge, what we will call "Dretskean Sensitivity":

Dretskean Sensitivity

A necessary condition of an agent's knowledge of a contingent proposition, ϕ , is that she has a true belief in ϕ in the actual world and that, in the nearest possible world or worlds in which ϕ is false, she does not believe ϕ .¹⁴

The basic idea behind Dretskean sensitivity is that for a belief to count as knowledge it must at least "track" the truth in the sense that, not only is it true, but, had what is believed been false, the agent would not have believed it. With this condition in play, Dretske can get the result he wants.

Suppose the actual world is pretty much as one takes it to be. It follows that one's belief that one is currently in one's office and one's belief that one is not a brain-in-a-vat are both true. Now consider whether the former belief counts as an instance of knowledge. On this account it does (at least pending any further conditions that one wants to add to Dretskean sensitivity), because in the nearest possible world in which one is not in one's office—the world in which, for example, one is in the corridor outside one's office—one no longer believes that one is. One's belief thus "tracks" the truth adequately to be a candidate for knowledge. In contrast, consider one's belief that one is not a brain-in-a-vat. The problem with this belief is that in the nearest possible world in which this belief is false (i.e., the brain-in-a-vat-world), one continues to have a belief that one is not a brain-in-a-vat because in this world one is the victim of a widespread deception. Accordingly, this belief fails to

meet the necessary condition for knowledge set out in Dretskean sensitivity and thus is not even in the running to be an instance of knowledge.

Dretske can thus use this notion of doxastic sensitivity to explain why closure fails by showing how knowledge needs to be understood relative to a certain relevant range of possible worlds which is variable depending upon the proposition at issue. As a result, one can know one proposition relative to one set of possible worlds, know the entailment to a second proposition, and yet fail to know the second proposition relative to a *different* set of possible worlds. More specifically, one can have knowledge of everyday propositions relative to one set of possible worlds that is “near to” the actual world, know that it entails the denial of a skeptical hypothesis, and yet lack knowledge of the denial of the skeptical hypothesis because knowledge is here relative to the “far off” possible worlds that are quite unlike the actual world. The notion of “relevance” at issue in the basic relevant alternatives account is thus cashed out in explicitly modal terms. Moreover, since a number of logical principles fail in modal contexts because of this sort of variability, it should not come as much of a surprise to find that closure meets a similar fate. Dretske is thus in a position to offer a plausible account of knowledge that can accommodate all of the claims that we saw him wanting to make earlier.¹⁵

4. PROBLEMS WITH THE DRETSKEAN LINE

This conception of knowledge, along with the later more elaborate version advanced by Nozick (1981),¹⁶ has been extremely influential. It is worth noting, however, that it is an epistemologically *revisionist* anti-skeptical proposal because it results in the denial of the highly plausible principle of closure. With this in mind,

there is a *prima facie* tension involved in adopting such a proposal, despite the compelling defense of this position that Dretske, Nozick, and others have offered. After all, the idea that we must know the known consequences of what we know is extremely strong and the rejection of this claim cannot be taken lightly. As DeRose (1995, §5), has pointed out, dropping closure means allowing what he calls “abominable conjunctions,” such as that one knows that one has two hands but one does not know that one is not a brain-in-a-vat. That closure should hold has been one of the main motivations for alternative interpretations of the core relevant alternatives anti-skeptical thesis that do not result in the rejection of closure, as we will see in a moment.

Besides this line of criticism against the Dretske-Nozick account, a number of other claims have been made. We will focus here on those avenues of critique that are directed at the view presented as an anti-skeptical proposal, rather than as an analysis (albeit perhaps only a partial one) of knowledge. For example, Edward Craig (1989; 1990b) has argued that the Dretske-Nozick proposal is either impotent at meeting skeptical arguments or unnecessary. That is, the strategy only works on the assumption that skeptical possible worlds are indeed far off worlds, and Craig argues that if we are entitled to *that* supposition then we have no need of an anti-skeptical strategy. Conversely, if we are not entitled to that supposition, then the modal analysis of knowledge offered by Dretske and Nozick leaves us in an *impasse* with the skeptic that is no better than we were in before. That is, all that the Dretskean approach has achieved is to show us that, *provided the world is in fact pretty much as we take it to be*, skepticism is false and this still leaves the issue of whether the

world is in fact pretty much as we take it to be unresolved.

Craig's objection is surely wrong, however, because the skeptical argument purported to show that knowledge was *impossible*, and the Dretske-Nozick account at least refutes this claim by showing that knowledge is possible; that we can have knowledge provided that certain conditions obtain.¹⁷ What is true, however, is that there is nothing in the Dretske-Nozick line which demands that we should be able to become reflectively aware that we have met these conditions if we are to have knowledge, and this element of the view may well be problematic. For if this is the case then the existential force of skepticism—that we could indeed, for all we can tell, be a victim of a skeptical hypothesis—is just as powerful as ever. The worry here is that there is always going to be something intellectually unsatisfactory about an anti-skeptical proposal that is run along epistemologically externalist lines.¹⁸ We will consider the role of epistemological externalism and internalism in the skeptical debate in more detail below, since it raises issues which affect all anti-skeptical proposals, regardless of whether they retain closure.

The large body of critical appraisal of the Dretskean proposal falls, however, on the rejection of the closure principle.¹⁹ There are two ways in which this critique is often run. Either critics argue directly for the retention of closure and thereby against the Dretskean line by default, or else they try to offer an alternative construal of the motivation for the Dretskean line that retains closure. Peter Klein (1981; 1995) is a good representative of the former position, arguing that, *contra* Dretske, we can indeed come to know (/be justified in believing) that the animal in the pen is not a cleverly disguised mule on the basis of our knowledge (/justified true belief) that it is a zebra

and our knowledge of the entailment. Central to Klein's view is a commitment to epistemological internalism, whose role in the skeptical debate will be discussed below, and a certain view about the structure of reasons that we do not have the space to go into here.²⁰ The most interesting attacks on the Dretskean approach to closure do not come from this quarter, however, but from those who claim to be motivated by similar epistemological concerns as those which motivate Dretske himself. That someone who endorses a radically different conception of the epistemological landscape should not find the Dretskean proposal plausible is not nearly so intriguing as dissent from those who sign up to many of the key Dretskean claims. In particular, it is interesting to note that two of the main rival views to the Dretskean line are wrought out of the same basic relevant alternatives claims.

Consider again the core thought expressed by proponents of a relevant alternatives model. This is that certain error-possibilities are irrelevant to the determination of knowledge, and thus that one can have knowledge merely by eliminating the salient error-possibilities. With this characterization of the core relevant alternatives thesis in mind, the natural question to ask is why the core relevant alternatives thought should be spelled out along Dretskean lines. After all, the Dretskean line *does* take far off skeptical possible worlds to be relevant to the determination of knowledge (albeit only knowledge of denials of skeptical hypotheses), whereas the basic relevant alternatives idea was surely that such far off worlds were manifestly *irrelevant* to the determination of knowledge (*any* knowledge). There thus seems to be an ambiguity in the relevant alternatives thesis. Either we take it as meaning that relevance is determined by the nearest not-P world (no matter how far out that might be), and thus end up with the thesis that Dretske and Nozick

propose, or else we construe it as simply demanding that only nearby possible worlds are relevant worlds. This is no mere technical dispute either, since a great deal hangs upon which alternative we adopt.

If we adopt the latter reading of the core relevant alternatives thesis then we are left with the thought that knowledge possession only requires tracking the truth in nearby possible worlds and on this construal the motivation for denying closure fades.²¹ For not only will an agent's belief in an everyday proposition typically track the truth in nearby possible worlds, so will her belief that she is not a brain-in-a-vat. Admittedly, this belief will not track the truth in the *nearest* possible world in which she is a brain-in-a-vat, but since this possible world is far off, this fact alone should not suffice on this construal of the relevant alternatives thesis to undermine her knowledge. An entirely different reading of the core relevant alternatives thesis thus seems to license the denial of the first premise of the skeptical argument, (S1)—that we are unable to know the denials of skeptical hypotheses—rather than the denial of closure and thus the rejection of the second premise, (S2). Moreover, that this reading of the relevant alternatives thesis does not result in the failure of closure means that it does possess some considerable dialectical advantage over the Dretske-Nozick thesis.

We will consider how such an approach to skepticism might function in more detail below. First, however, we will look at a different reading of the core relevant alternatives thesis that also does not result in the denial of closure. Where this line differs from the one just canvassed is that it does not straightforwardly allow that one can know the denials of skeptical hypotheses either. Instead, it argues that the standards for relevance are variable such that, although one knows everyday

propositions and thus the denials of skeptical hypotheses relative to a low standard of relevance, one lacks knowledge of these everyday propositions and thus of the denials of skeptical hypotheses at high standards of relevance. In this way, the hope is that this thesis can explain both our skeptical and anti-skeptical intuitions (and our attachment to closure), while nevertheless denying the universal correctness of the skeptical argument. This anti-skeptical thesis is *contextualism*, and since the version that we will be considering regards the mechanisms that alter these standards of relevance to be conversational mechanisms, we will call it *semantic contextualism*.

5. SEMANTIC CONTEXTUALISM

Semantic contextualism—as put forward by such figures as Cohen (1986; 1987; 1988; 1990; 1991; 1998a; 1999; 2000a), DeRose (1995) and Lewis (1996)—is without doubt the most dominant form of contextualist anti-skeptical theory in the current literature.²² In what follows, the discussion of the semantic contextualist thesis will primarily focus on DeRose's version since this is the most developed (and, arguably, the most influential) characterization of the thesis which incorporates most of the main features of the other two accounts. Later on, however, some of the key differences between these three versions of the semantic contextualist view will be discussed. Moreover, in §8, we will examine a different type of contextualist theory that is advanced by Williams (1991)—what we will call *inferential contextualism*—which does not conform to the basic semantic contextualist template.

Before we look at the DeRose version of the semantic contextualist thesis, however, it is worthwhile first being clear about the following supposed features of

the “phenomenology” of our engagement with skepticism since semantic contextualism can most naturally be viewed as a response to—and to some extent as an accommodation of—these features:

- I. Ascriptions of knowledge to subjects in conversational contexts in which skeptical error-possibilities have been raised seems improper.
- II. In conversational contexts in which no skeptical error-possibilities are in play it seems perfectly appropriate to ascribe knowledge to subjects.
- III. All that may change when one moves from a non-skeptical conversational context to a skeptical context are mere conversational factors.

The first two features represent what Williams (1991, chapter one) refers to as our “biperspectivalism,” our intuition that skepticism is compelling under the conditions of philosophical reflection but never able to affect our everyday life where it is all but ignored. The third feature creates the tension because it highlights our sense that one of these judgements must be wrong. That is, since conversational topic has no obvious bearing on the epistemic status of a subject’s beliefs, it ought to be universally true (i.e., *whatever* the conversational context) that the subject either does or does not know the propositions in question. So either our knowledge ascriptions in everyday contexts are right (and thus the skeptic is wrong), in which case we shouldn’t take skepticism so seriously in conversational contexts in which skepticism is at issue; or else the skeptic is right, and thus our everyday practice of ascribing knowledge is under threat. Semantic contextualism opposes this thought with the suggestion that what is actually occurring is not a contradiction but a responsiveness, on the part of the attributor of knowledge, to a fluctuation in the epistemic standards (and with them the

subject’s possession of knowledge) caused by a change in the conversational context. In effect, and *contra* the third claim listed above, semantic contextualism holds that mere changes in the conversational context can have an effect on the epistemic status of one’s beliefs so that it can be true both that one has knowledge in everyday contexts and that one lacks it in skeptical conversational contexts.

Accordingly, we find DeRose (1995) arguing that the basic contextualist strategy pivots upon the acceptability, and appropriate use, of the following contextualist thesis:

Suppose a speaker A (for “attributor”) says, “S knows that P,” of a subject S’s true belief that P. According to contextualist theories of knowledge attributions, how strong an epistemic position S must be in with respect to P for A’s assertion to be true can vary according to features of A’s conversational context (DeRose 1995, p. 4).

Here we get the essentials of the semantic contextualist view. In particular, we get (i), the claim that the strength of epistemic position that an agent needs to be in if she is to have knowledge can fluctuate from context to context (which makes the thesis contextualist); and (ii), the claim that what is at issue in the determination of contexts in this respect is the conversational context (which makes the contextualist thesis semantic). The interesting question now is how the details of how this account is to function are to be spelled out.

The first thing that DeRose tries to capture is the intuition that as one moves from one conversational context to another one’s epistemic situation (one’s total informational state for instance) could remain exactly the same. DeRose accommodates this intuition in conjunction with the contextualist picture by arguing, as the above quotation indicates, that although one’s “epistemic position” is constant at any one

time, the epistemic position that one *needs to be in* so as to count as possessing knowledge can be variable. Strength of “epistemic position” is characterized by DeRose as follows:

being in a strong epistemic position with respect to P is to have a belief as to whether P is true match the fact of the matter as to whether P is true, not only in the actual world, but also at the worlds sufficiently close to the actual world. That is, one’s belief should not only be true, but also should be non-accidentally true, where this requires one’s belief as to whether P is true to match the fact of the matter at nearby worlds. The further away one gets from the actual world, while still having it be the case that one’s belief matches the fact at worlds that far away and closer, the stronger a position one is in with respect to P (DeRose 1995, p. 34).

In order to see this, imagine that Lars believes that his car is outside on the basis of a certain fixed informational state (which involves, perhaps, his memory of the car being there a few hours ago, his grounds for believing that no one would steal it, and so forth). Now imagine an (almost) exact counterpart of Lars—Lars*—who is in exactly the same cognitive state except that he has the extra piece of information that the car was there a minute ago (perhaps he looked). Clearly, Lars* will be in a slightly better epistemic position with respect to his belief that his car is outside than Lars. Although they will, in general, track the truth across the same set of possible worlds, Lars* will track the truth in a few extra possible worlds, such as the possible worlds in which his car was stolen ten minutes ago.

DeRose then goes on to describe the mechanism which changes the conversational context (and thereby alters the epistemic standards at play) in terms of the thesis of “Dretskean sensitivity” that we saw above. Recall that for an agent to have a

belief in P that is sensitive in this way, the agent must not only have a true belief in P in the actual world, but also *not* believe P in the nearest possible world (or worlds) in which P is false. DeRose’s thought is that in any particular conversational context there is a certain set of propositions that are explicitly at issue and that the agent must, at the very least, be sensitive to all these “explicit” propositions if she is to know them. Moreover, the most demanding of these propositions—the proposition which has a negation that occupies the furthest-out possible world—will set the standard for that conversational context since this not-P world will determine the extent of possible worlds that one’s beliefs must be able to track if one is to be truly said to know a proposition in that context. Knowing a proposition thus involves being in an epistemic position sufficient to track the truth across the range of possible worlds determined by the most demanding proposition explicit to that context. Crucially, however (and this detail will be expanded upon in a moment), this point also applies to propositions which are *implicit* to a conversational context (i.e., propositions which one believes but which are not explicit to that conversational context). In order to know such a proposition—even if one’s belief in that proposition is not sensitive—one need only be in a sufficient epistemic position to meet the standards of that context. (The importance of this point will soon become apparent).

DeRose then characterizes the mechanism that brings about an upward shift in epistemic standards as follows:

When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards for 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

belief in that particular P to be sensitive for it to count as knowledge (DeRose 1995, p. 36).

That is, what changes a conversational context is when a new proposition is made explicit to that context which is more demanding than any of the propositions currently explicit in that context. This will thus increase the range of possible worlds at issue in the determination of knowledge, and thereby increase the strength of epistemic position required in order to be truly said to know.

What motivates this claim is the fact that, as Lewis (1979) famously argued, when it comes to “context-sensitive” terms like “flat” or “knowledge,” the conversational “score” tends to change depending upon the assertions of that context. We may all agree that the table in front of us is “flat” in an everyday context, but, *ceteris paribus*, if someone enters the room and denies that it is flat we do not thereby disagree with her. Instead, we take it that she means “flat” in some more demanding sense and so raise the standards for “flatness” so as to make her assertion true (this is what Lewis [1996, 559] calls a “rule of accommodation”). That is, we take it that the new participant of our conversational context means flat in some more restricted sense so that the barely perceptible bumps on the table before us are sufficient to make the claim “This table is flat” false. DeRose considers the Lewis line to have captured something intuitive about the pragmatics of how we use our “context-sensitive” terms and, moreover, believes that epistemic terms such as “knowledge” behave in a similar way.²³

An example will help clarify matters here. Imagine an agent in a quotidian context in which only everyday propositions, such as whether or not one is currently having dinner with one’s brother (P), or whether or not the garden gate has been

closed (Q), are at issue. Sensitivity to these everyday propositions will only require the consideration of nearby possible worlds and thus the strength of epistemic position demanded will be very weak. Let us say, plausibly, that the possible world in which one is not having dinner with one’s brother is “further out” than the possible world in which the garden gate is not closed. This proposition will thus determine the range of possible worlds at issue in the determination of knowledge in that conversational context. Let us further suppose that the agent in question does have a sensitive belief in this proposition. The issue of what other propositions the subject knows will now be decided by whether the agent’s belief in those propositions will track the truth across the range of possible worlds determined by not-P. If, for instance, the agent’s belief that the garden gate is closed matches the truth as to whether Q in all of the possible worlds within that range, then she will know Q. Equally, however, if the subject’s belief in a proposition which is implicit to that conversational context tracks the truth across this range of possible worlds then that proposition will be known also, even if the agent’s belief in that proposition is not sensitive.

Consider, for instance, the agent’s belief that she is not a brain-in-a-vat, a proposition which, as we saw above, one cannot be sensitive to because in the nearest brain-in-a-vat-world one still believes that one is not a brain-in-a-vat. On the contextualist model, however, if one is in a conversational context in which such a proposition is not explicit, then one can know this proposition just so long as one has a belief as to whether this proposition is true which matches the facts as to whether it is true within the range of possible worlds at issue. And, clearly, this demand will be trivially satisfied in the above scenario where the subject has a sensitive belief in the

ordinary proposition, P. Given that one has such a sensitive belief in P, then it must be the case that the brain-in-a-vat skeptical world is, modally speaking, far out, for if it weren't, then this would affect the sensitivity of the subject's beliefs in ordinary propositions like P. Accordingly, on this view, all the subject needs in this context is a stubborn belief that she is not a brain-in-a-vat in order to be truly said to know this proposition in this conversational context. The contextualist can thus capture the second element of the "phenomenology" of our engagement with skepticism that we noted above—that, in quotidian conversational contexts, we are perfectly willing to ascribe knowledge of everyday propositions and also feel that we must know the denials of skeptical hypotheses as well.

One might wonder why the word "feel" is used here. The reason is that, on the contextualist account, if one were to explicitly mention these anti-skeptical propositions (as one would if one were to verbally ascribe knowledge of them to oneself), then one would thereby make that proposition explicit to the conversational context and so change the epistemic standards needed for knowledge accordingly.²⁴ In order to have knowledge that one is not a brain-in-a-vat within the new conversational context, one's belief that one is not a brain-in-a-vat must now exhibit sensitivity (which, as we saw above, is impossible), and the possible worlds relevant to the determination of that sensitivity will be relevant to one's knowledge of even everyday propositions. Accordingly, one will now lack knowledge both of the denial of the skeptical hypothesis (because one's belief in this respect is not sensitive), *and* of the everyday propositions (since even though one's beliefs in these propositions are sensitive, one can never be in an epistemic position that would support knowledge of them which would be strong

enough to track the truth in far off brain-in-a-vat-worlds). The contextualist thus claims to have captured the other two aspects of the "phenomenology" of our engagement with skepticism—that we are completely unwilling to ascribe knowledge in skeptical conversational contexts, and that this is so even when the only thing that may have changed from the non-skeptical conversational context in which we were willing to ascribe knowledge is the course of the conversation. Moreover, the contextualist has done this without either conceding the universal truth of skepticism (since skepticism is false in everyday contexts), or denying closure (since there is no single context in which one both knows an everyday proposition while lacking knowledge of the denial of a skeptical hypothesis).²⁵ Accordingly, DeRose claims to have "solved" the skeptical paradox in an entirely intuitive manner.

6. PROBLEMS WITH SEMANTIC CONTEXTUALISM

The semantic contextualist proposals made by Lewis, Cohen, and others run along similar lines to the DeRose thesis. The key difference between DeRose and Lewis is that Lewis cashes out his thesis in terms of a series of rules which determine when we may and may not properly ignore a certain error-possibility rather than in terms of a general modal account of knowledge. The key difference between Cohen's position and that advanced by Lewis and DeRose is that it is centered upon the concept of justification rather than knowledge and incorporates a certain view about the structure of reasons (Cohen 1999; 2000a). That the DeRose view differs in these ways from the views presented by Lewis and Cohen may work in its favor. As Williamson (2001) has pointed out, DeRose's more straightforward position

may be insulated from the kind of *ad hoc* charges regarding Lewis's employment of numerous rules governing when an error-possibility is properly ignored (as put forward, for example, by Williams [2001]). Moreover, the stress on justification in Cohen's account can make it unappealing to those who find such a notion problematic, at least to the extent that it is understood as playing an essential role in knowledge acquisition.²⁶ There are thus *prima facie* grounds for thinking that if any semantic contextualist theory will work, then it will be one run along the lines presented by DeRose.

In any case, the most interesting objections against semantic externalism do not rest upon specific elements of the particular versions of this position, but rather strike against the general approach. Perhaps the most obvious concern facing semantic contextualism is that its claim that the epistemic status of an agent's beliefs can be variable in response to mere conversational factors is highly contentious. For instance, *contra* Cohen's (1991) contextualist characterization of his relevant alternatives position, Dretske writes:

Knowledge is relative, yes, but relative to the extra-evidential circumstances *of the knower* and those who, like the knower, have the same stake in what is true in the matter in question. Knowledge is context sensitive, according to this view, but it is not indexical. If two people disagree about what is known, they have a genuine disagreement. They can't both be right (Dretske 1991, p. 191).

That is, Dretske rejects outright the thought that the truth value of a knowledge claim can be dependent upon anything other than what the concrete features of the situation are. A detailed discussion of this line of attack against semantic contextualism will not be undertaken here, however, since the semantic contextualist proposal is clearly meant to be understood as a revisionistic

proposal in this respect. Indeed, this is to be expected given that, as we have already seen, any plausible anti-skeptical proposal will have to deny some claim that is otherwise thought to be intuitive (Dretske himself denies closure, for example). To query this element of the thesis alone is thus not to give it a run for its money (although, if one can couple this line of critique with other concerns then it can carry some dialectical weight).

The other more general lines of criticism against semantic contextualism fall into three main camps. First, there is the allegation that semantic contextualism leaves us no better off than we were before. Second, there is the worry about how the contextualist is to explain our knowledge (albeit only in quotidian contexts) of the denials of skeptical hypotheses. As we noted above, we do have a strong intuition that we can never know such propositions in any context (recall that this intuition was one of the motivating factors behind the Dretskean theory). Third, there is the claim that semantic contextualism is guilty of overkill in its approach to skepticism in that it essentially incorporates a claim which would suffice to undermine radical skepticism by itself (i. e., regardless of whether it is combined with a general contextualist thesis). Accordingly, on this view, contextualism is both unnecessary and ill-motivated.

The worry that motivates the first concern is that what semantic contextualism does concede to the skeptic is the "hierarchical" structure of her doubts. That is, on the semantic contextualist view the skeptic is indeed working at a high epistemic standard, one that is more demanding than our everyday standards. The problem with conceding this much to the skeptic is that it appears to legitimate the concern that the skeptic's standards are the *right* standards, and thus that, although we are content to

ascribe everyday knowledge in quotidian contexts, we ought not to ascribe such knowledge because, strictly speaking, we lack knowledge relative to the proper skeptical standards that should be employed. This line of thought should be familiar as the general infallibilist line that we saw Unger arguing for in §2 and one might think that one could dismiss it on the same grounds that were canvassed there. Matters are not quite so simple, however. Indeed, it is interesting that in more recent works Unger (1984; 1986) has toned down his commitment to infallibilism in favor of a different conception of the skeptical debate which argues that there is nothing to tell between the contextualist reading of our epistemic concepts and what he calls “invariantism,” which would license infallibilist conclusions. Unger’s point here is that all that contextualism can achieve is, at best, a kind of *impasse* with the skeptic such that we know that one of these views must be right but where we are unable to definitively determine which. Is it that the standards are high and only seem variable because we are content to talk loosely in everyday contexts (in which case the skeptic is right), or is it that our epistemic concepts are genuinely variable and thus that we do have knowledge relative to low everyday epistemic standards (in which case the skeptic is wrong)? The skeptical thought that arises here is that if we are unable to offer sufficient reasons for preferring the latter scenario over the former then contextualism does not put us in a better situation than we were in before.²⁷

The second worry—that contextualism is unable to explain how we can have knowledge of the denials of skeptical hypotheses—is particularly striking because if this claim holds then semantic contextualism does not even supply us with an *impasse* with the skeptic. Instead, the position would simply be incoherent. The

thought here is that since our presumably *empirical* knowledge in this respect cannot be coherently thought of as being the result of an empirical investigation, we cannot make sense of it at all. Perhaps the bravest response to this line of attack in the recent literature can be found in the suggestion made by both Cohen (1999; 2000a) and DeRose (2000) that it might be possible to understand an agent’s knowledge of the denials of skeptical hypotheses along *a priori* lines. Cohen motivates this thought relative to what he terms the “*a priori* rationality” of denying skeptical error-possibilities (e.g., Cohen 2000a, p. 104), while DeRose employs Putnamian reflections on semantic externalism as a means of showing how we might have *a priori* knowledge of empirical truths.²⁸

Even if these lines of argument could be made palatable, however (and, it must be admitted, neither is particularly appealing), then a further line of attack (the third worry listed above) would instantly emerge. For if we can indeed make sense of our putative knowledge of the denials of skeptical hypotheses then what could the motivation for an epistemologically revisionist thesis like contextualism possibly be? Recall that the epistemological revisionism incorporated in contextualism only seemed necessary because our knowledge of these propositions seemed so insecure. If we can have knowledge of these propositions, however, then why not simply motivate one’s anti-skepticism by straightforwardly denying the first premise of the skeptical argument, (S1), rather than by going contextualist? We will return to consider this kind of anti-skeptical proposal—known as the “Neo-Moorean” view—in more detail in §8.

7. “HINGE” PROPOSITIONS AND INFERENCEAL CONTEXTUALISM

Before we look at this neo-Moorean approach to radical skepticism, however, it is worthwhile to first consider two lines of argument which, at least superficially, might appear to be analogue approaches to that of the Dretskean and semantic contextualist line. Where these anti-skeptical approaches differ, in the first instance at least, from those just canvassed, is that they primarily take their stimulus not from the relevant alternatives debate but from Wittgenstein’s last notebooks (published as *On Certainty*) which were on the topic of knowledge and skepticism.

Both of these accounts focus upon the Wittgensteinian notion of a “hinge proposition,” though they each use the notion in a different way. The first camp, which includes such figures as Strawson (1985), Wright (1985; 1991; 2000a), Putnam (1992), and Stroll (1994), might be considered to be offering the purest hinge proposition thesis because for them the notion is the primary unit upon which their anti-skepticism rests. In contrast, the conception of this notion employed by Williams (1991) is developed along explicitly contextualist lines. Interestingly, however, the type of contextualism that emerges is not of a semantic variety. These views will be considered in turn. First, however, we need to look a little at what Wittgenstein had in mind in his employment of this notion.

Wittgenstein describes hinge propositions as follows:

the *questions* that we raise and our *doubts* depend upon the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.

That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted.

But it isn’t that the situation is like this: We just *can’t* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put (Wittgenstein 1969, §§341–343).²⁹

As this quotation indicates, what is odd about these propositions is that, unlike other seemingly empirical propositions, our belief in them does not seem to either stand in need of evidential buttress or, for that matter, be legitimately prone to coherent doubt. And this property is not explained merely by the fact that these propositions are “*in deed* not doubted,” since the situation is rather that we do not doubt them because, in some sense, we *ought* not to doubt them. Even despite their lack of sufficient evidential support, their immunity to coherent doubt is part of “the logic of our scientific investigations.”

In proposing this notion Wittgenstein was explicitly challenging the conventional epistemological wisdom that a belief is only legitimately held if it is sufficiently evidentially grounded (otherwise it is open to legitimate doubt), and that *no* belief in an empirical proposition is beyond coherent doubt should the grounds for that belief be found wanting. In particular, Wittgenstein’s remarks here were primarily targeted at G. E. Moore’s (1925; 1939) famous “common-sense” response to the skeptic. In effect, what Moore did was reverse the skeptical train of reasoning by arguing, on the basis of his conviction that the skeptical conclusion must be false, that he did know the denial of the relevant skeptical hypothesis after all. In general, the Moorean response to the skeptic is to respond to the skeptic’s *modus ponens* argument with the corresponding *modus tollens*.

Wittgenstein claimed that where Moore went wrong was in treating a hinge proposition as if it were just a normal empirical

proposition. In particular, he focussed upon Moore's use of the everyday proposition "I have two hands" in this respect, arguing that this is, in normal circumstances, a hinge proposition and thus that the certainty that we attach to it is not due to any evidence that we might have for its favor but rather reflects the fact that in those circumstances it is what "stands fast" in our assessment of other propositions. Nothing is more certain than this proposition in ordinary circumstances, and thus no other belief could be coherently regarded as providing support for it or be coherently thought to undermine it. As Wittgenstein puts the matter at one point:

My having two hands is, in normal circumstances, as certain as anything that I could produce in evidence for it.

That is why I am not in a position to take the sight of my hand as evidence for it (Wittgenstein 1969, §250).

Given that this is so, however, Moore cannot use the certainty he has for this proposition in order to derive support for his belief in the denials of skeptical hypotheses since, strictly speaking, his belief in this proposition is not grounded at all. Indeed, Wittgenstein goes further to argue that, to the extent that Moore regards this proposition as being grounded by the evidence he has for it, then the circumstances are no longer "normal" in the relevant respects and thus the proposition is no longer a hinge proposition. And in these conditions, it would be ridiculous to try to buttress one's belief in the denials of skeptical hypotheses with the support that one has for this proposition (Wittgenstein 1969, §20) compares it to basing one's belief in the existence of the external world on the grounds one has for thinking that there are other planets).³⁰

Wittgenstein is not endorsing radical skepticism by attacking Moore in this way,

for the general anti-skeptical line that emerges is that where the skeptic, like Moore, goes wrong is in considering certain basic propositions as being coherently open to doubt. In doing so, or so the argument runs, she fails to pay due attention to the pivotal role that certain, apparently ordinary, propositions can play in our systems of beliefs and thus to the manner in which even everyday propositions can be held with conviction without thereby standing in need of a corresponding degree of evidential buttress.

So construed, the hinge proposition line is largely a *diagnosis* of why we are so tempted by skeptical arguments rather than a refutation of the skeptical argument. Indeed, in the hands of Strawson (1985), Putnam (1992), and Stroll (1994) this is what the hinge propositions line is largely designed to do. So construed, however, it is difficult to see what comfort it offers since it is far from clear that it results in the denial of any of the intuitions that we saw at issue in our formulation of the skeptical argument in §1. In particular, this diagnostic thesis neither results in the denial of closure nor in a contextualist thesis. Moreover, the standard hinge proposition line explicitly allows that hinge propositions are not known, strictly speaking, since they fall outside of the ambit of epistemic evaluation (being instead the nodes around which epistemic evaluation takes place). The anti-skeptical import of the basic hinge proposition line is thus moot.³¹

In the hands of Wright (1985; 1991; 2000a) and Williams (1991), however, one finds a much stronger version of this approach. We will focus on each in turn. In effect, Wright argues that doubting hinge propositions does not merely reflect a misunderstanding of the epistemological landscape, rather it also leads to intellectual self-subversion. The most sophisticated of these arguments can be found in Wright

(1991), where he attempts to show that dreaming skepticism must be unsound because it leads to a more radical form of doubt which undermines the very presuppositions that the skeptic needs to employ in order to get her dreaming skepticism off the ground. Thus, dreaming skepticism is necessarily intellectually self-subverting and thus it can be disregarded with impunity. What this line adds to the previous approaches is a principled epistemic ground for discounting the skeptic's doubt. In particular, since it is epistemically irrational to doubt the hinge propositions in question, the implication of the argument that Wright offers is that we can indeed know the denials of skeptical hypotheses on the basis of our knowledge of everyday propositions even though our belief in the former is not evidentially grounded. As Wright (1991, pp. 107–108) puts the point, “the impossibility of *earning* a warrant that one is not now dreaming does not imply that no such warrant is ever *possessed*.”³²

Although Wright does succeed in at least showing how the diagnosis of radical skepticism that the hinge proposition line offers might lead to a refutation of the skeptical argument, a number of problems with his approach remain. One worry concerns the examples of hinge propositions that he uses. Whereas Wittgenstein seemed to have everyday propositions in mind in his use of this notion (and “I have two hands” in particular), Wright's argument works only if one takes the denials of skeptical hypotheses as hinge propositions, a move that receives only ambiguous support in the text.³³ A more serious worry concerns how, exactly, we are to cash out this idea that the warrant which underpins our putative knowledge of the denials of hinge propositions can be “unearned” in the fashion that Wright envisages. The idea that we could have knowledge of these propositions *purely because* we have

knowledge of everyday propositions does seem to be question begging. Again, then, the worry about closure in the context of the skeptical debate is brought to the fore. The intuition that closure captures is that an anti-skeptical theory must incorporate the claim that we can know the denials of skeptical hypotheses, even though such knowledge bears few, if any, of the usual hallmarks of empirical knowledge. This tension can only be resolved by either denying closure or allowing that we can know the denials of skeptical hypotheses, but each of these moves raises tensions of its own.

Wright responds to this latter worry in more recent work. Drawing upon remarks made by Martin Davies (1998), Wright (2000a) argues that we need to distinguish between the principle of closure and what he terms the principle of “transmission”. In essence, his way around this concern about closure is to maintain that while knowledge does indeed *transfer* across known entailments as closure demands, it does not *transmit*, where transmission requires something more demanding. In particular, transmission demands that the “cogency” of the argument be preserved, which is its aptitude to produce *rational conviction*. Here is Wright:

A *cogent* argument is one whereby someone could be moved to rational conviction of the truth of its conclusion (Wright 2000a, p. 140).

Wright thus concedes that if one does know the everyday propositions then, trivially, one must know the denials of skeptical hypotheses that are known to be entailed by those propositions. Nevertheless, he also argues that the fact that closure holds in these cases does not mean that any argument that attempts to establish anti-skeptical knowledge on this basis will be apt to convince a third party. As Wright notes, the argument is question

begging in the relevant respect since it is only by taking for granted the denials of radical skeptical hypotheses (i.e., by taking the relevant hinge propositions for granted) that the agent is able to have knowledge of the everyday propositions in the first place. This is a compelling move to make since it diagnoses the arguments against closure by characterising them in terms of how what is lacking from the conclusion is not knowledge but something just as valuable, a conclusion that is apt to convince. It also explains the failure of the Moorean approach, for where Moore goes wrong is not in lacking the knowledge that he claims to have, but rather in attempting to secure any conviction on the part of his audience by offering his argument. Indeed, this is a very Wittgensteinian claim to make since Wittgenstein himself argues that:

Moore's mistake lies in this—countering the assertion that *one cannot know* that, by saying “I do know it” (Wittgenstein 1969, §521, my italics).

As this quotation implies, the problem with Moore's argument is precisely not the falsity of the conclusion (which would be to validate either radical skepticism or the denial of closure), but rather concerns the manner in which he proposes it and the ends that it is designed to serve.

Still, compelling though this approach is, we are still in need of an account of knowledge which can explain how it is that we can know such propositions as the denials of skeptical hypotheses, something which Wright himself does not offer.³⁴ For this contribution to the anti-skeptical debate we must look elsewhere. Such an account will be considered in a moment, for one could view the neo-Moorean strategy as presenting the required analysis. First, however, we will briefly look at another influential reading of the hinge proposition thesis, due to Williams (1991).

As might be expected, what identifies Williams's inferential thesis as a contextualist account is that it incorporates the claim that one should understand knowledge relative to a context of some description. Like the semantic contextualist view put forward by DeRose and others, Williams argues that certain error-possibilities are only epistemically relevant, and thus potentially knowledge defeating, in certain contexts. Moreover, Williams is also keen to retain the closure principle. He does so on contextualist grounds, arguing that provided one keeps to the one context then closure will hold. As with the semantic contextualist position, then, apparent failures of closure are simply due to equivocations between different contexts.

The line Williams takes thus shares a central core of claims with the semantic contextualist view. Nevertheless, their disagreements are significant. The two main areas of disparity between the two theories are as follows. First, Williams does not individuate contexts along a “conversational” axis, but rather in terms of the inferential structure of that context (hence the name, *inferential* contextualism). Second (and as we will see this point is closely related to the inferential thesis), Williams does not allow a context-independent hierarchy of contexts. That is, unlike the semantic contextualists, Williams does not, for example, regard the skeptical context as being a more epistemically demanding context. Rather, it is just a context which employs a different epistemic structure.

Indeed, on Williams's view, contextualism just *is* the thesis that there is *no* such hierarchy of epistemic contexts—instead, each context is, epistemically speaking, autonomous. To think otherwise is, he thinks, to fall victim to the doctrine of “epistemological realism”—the view that the objects of epistemological inquiry have an inherent, and thus context-independent,

structure. In contrast, the inferential contextualism that he advances is defined as the denial of this thesis. It holds that:

the epistemic status of a given proposition is liable to shift with situational, disciplinary and other contextually variable factors: it is to hold that, independently of such influences, a proposition has no epistemic status whatsoever (Williams 1991, p. 119).

And Williams is quite clear that this last phrase “has no epistemic status whatsoever” is meant to indicate that there is no context-independent means by which we can evaluate the standards wrought in different contexts. He describes his view as a “deflationary” theory of knowledge, in that it holds that there need be nothing that ties all instances of knowledge together other than the fact that they are instances of knowledge. He writes:

A deflationary account of “know” may show how the word is embedded in a teachable and useful linguistic practice, without supposing that “being known to be true” denotes a property that groups propositions into a theoretically significant kind. We can have an account of the use and utility of “know” without supposing that there is such a thing as human knowledge (Williams 1991, p. 113).

It is as a consequence of such a view that the very validity of the epistemological enterprise, at least as it is commonly understood, is called into question:

If we give up the idea of pervasive, underlying epistemological constraints; if we start to see the plurality of constraints that inform the various special disciplines, never mind ordinary, unsystematic factual discourse, as genuinely irreducible; if we become suspicious of the idea that “our powers and faculties” can be evaluated independently of everything having to do with the world and our place in it: then we lose our grip on the idea of “human knowledge” as an object of theory (Williams 1991, p. 106).

And to say that “human knowledge” is not a suitable object of theory is itself to say that the epistemological project cannot be *systematically* conceived. On Williams’s view there is no epistemological analysis to be conducted outside of contextual parameters and, accordingly, there are no context-independent standards either as the semantic contextualist model would suggest.

It is as a consequence of this stance that Williams is forced to concede that the skeptic is perfectly correct in her assessment of our knowledge, but only because she is operating within an epistemic context that utilises a standard which defeats our everyday knowledge. That is, the skeptic’s conclusions are correct, but, by being confined to a specific epistemic context, they lack the hegemony that she requires in order to cause the intended epistemic harm. It does not follow from the truth of skepticism that we lack the everyday knowledge that we attribute to ourselves, or even that such knowledge is inferior to the knowledge that the skeptic has in mind (which would, in line with the semantic contextualist view, presuppose a hierarchy). That is:

The skeptic takes himself to have discovered, under the conditions of philosophical reflection, that knowledge of the world is impossible. But in fact, the most he has discovered is that knowledge of the world is *impossible under the conditions of philosophical reflection* (Williams 1991, p. 130).

So whereas the “hierarchical” camp of semantic contextualists tend to individuate contexts in terms of a context-transcendent criterion of rigor, Williams’s schema allows no such ordering of contexts. For him, a context is individuated purely in terms of the epistemic structure it endorses—in terms of the inferential relations that obtain between the types of beliefs that that context is interested in. And since no context employs universal standards, this contextual

epistemic structure is also identified in terms of what it takes for granted—which propositions it regards as being immune from doubt in terms of that context. Williams calls the defining assumptions of a context of inquiry its “methodological necessities,” and this notion is explicitly meant to capture the chief insights behind the Wittgensteinian notion of a hinge proposition. When we do history, for example, we take the general veracity of historical documentation for granted, as well as the denials of certain skeptical scenarios such as that the world came into existence five minutes ago replete with the traces of a distant ancestry (the so-called “Russellian Hypothesis”). To doubt such methodological necessities (/hinge propositions) is not, he argues, to conduct our historical investigations in a more exacting fashion, but rather to engage in a different sort of investigation altogether, one that is guided by traditional epistemological concerns.

The methodological necessities of the traditional epistemological project which, Williams claims, spawns the skeptical threat, are meant to involve a commitment to this false doctrine of epistemological realism. This leads, he argues, to an antiquated foundationalism,³⁵ one manifestation of which is the traditional epistemologist’s concern with the problem of the external world. This problem is meant to reflect the inadequacy of beliefs of one type—concerning immediate experience—at serving the purpose of epistemically supporting beliefs of another type—concerning material objects in the external world. With the problem so characterized, Williams maintains that it should come as no surprise to find that it is without a solution (there are no beliefs concerning immediate experience that are able to act as epistemic guarantors for beliefs concerning objects in the material world). But, he contends, this need not result in a general external world skepticism because

this conception of the inferential ordering needed for warranted beliefs about “external” objects is far from obligatory. In terms of another type of inquiry, such as psychological investigations of perception for example, we may legitimately begin with beliefs about material objects and draw inferences about immediate experience. And since no context has any epistemological ascendancy over any other, the project of justifying psychological beliefs with reference to external world beliefs is just as valid as any epistemological theory which demanded that the inferential relations should point in the opposite direction.

Williams’s inferentialist contextualist view is thus both more radical and more demanding than its semantic counterpart. On the one hand, it is more radical because Williams does not concede to the skeptic that her skepticism functions at a higher epistemic standard. Instead, he argues that such skepticism merely reflects a different, and faulty, conception of the epistemological landscape. Williams’s view therefore evades one strand of criticism that we saw levelled at the semantic contextualist account above. On the other hand, it is more demanding because, for related reasons, Williams does not believe that mere changes in the conversational context can suffice to bring about a different epistemic context. Instead, there must be an actual difference in the inferential structure that is employed, and thus, given his contextualism, a difference in what is being taken for granted relative to what. Nevertheless, inferential contextualism may well carry with it even more troubling problems of its own, not least the worry that this approach is allied to a general quietistic philosophical approach. Critical appraisal of this theory has been limited, however, since this variant of the contextualist thesis has tended to be obscured by its semantic counterpart.³⁶ As the dust

settles on the current wave of discussion of the semantic contextualist proposal, one would expect this distinctive thesis to gain a greater degree of attention.

8. NEO-MOOREAN RESPONSES TO SKEPTICISM

Back, then, to Wright, or, more specifically, to what was absent in Wright's account, *viz.*, an analysis of knowledge that could account for the conclusions that he presents. In essence, what Wright is offering is a neo-Moorean response to skepticism in that he allows, with Moore, that if we do know everyday propositions then we must know the denials of radical skeptical hypotheses that are known to be entailed by them. Where Wright differs from Moore is in not allowing that one can coherently argue to this conclusion in any way that could secure rational conviction. Like Moore, however, Wright fails to supplement this account with an analysis of knowledge that would support it. In so doing, Wright is failing to properly engage with other proponents of the debate regarding radical skepticism—such as Dretske, Nozick, and DeRose—who do supplement their anti-skeptical account with an analysis of knowledge (albeit, perhaps, only a partial one) that backs up their theory. Although Wright does not offer such an account, however, there are analyses of knowledge in the literature that might provide support for his view, and it is to these analyses that we now turn.

Recall that the Dretskean line made “Dretskean sensitivity” an essential component of knowledge, where this demanded that an agent should have a belief which is not only true, but which “tracks” the truth of the proposition in question in the nearest possible world in which that proposition is false, no matter how “far off,” modally speaking, that world is. It was this

element of the thesis that secured the skeptic's first premise, (S1), because no one's belief in the denial of a radical skeptical hypothesis could track the truth in this sense. Similarly, DeRose also made Dretskean sensitivity relevant to knowledge, albeit only as regards those propositions that were at issue in that context. Accordingly, in skeptical conversational contexts where skeptical hypotheses were at issue agents lacked knowledge of these propositions. What both of these anti-skeptical approaches have in common is thus that they allow that, at least in demanding contexts, we lack knowledge of the denials of “far off” skeptical error-possibilities (and thus that the first premise of the skeptical argument, [S1], is true).

In contrast to these lines of thought, Sosa (1999a; 2000a) has argued that we should instead regard what he terms “safety” as being central to knowledge rather than sensitivity. In essence, he characterizes this notion as follows (Sosa 1999a, 142):

Safety

A necessary condition of an agent's knowledge of a contingent proposition, ϕ , is that the agent should have a true belief in ϕ in the actual world and, across a wide range of nearby possible worlds, if she believes ϕ then ϕ is true.³⁷

Similar proposals have also been put forward by Sainsbury (1997), Williamson (2000a; 2000c, chapter 8), and Pritchard (2002c; cf. Pritchard 2001e; 2002a). The anti-skeptical advantage that this kind of proposal offers over both the Dretskean line and either of the contextualist views that we have discussed is that it allows one to endorse a version of the Moorean proposal which neither issues in the denial of closure nor results in contextualism.

Take the former point first. What prompted the denial of closure was the fact that, if we take sensitivity as a necessary condition on knowledge, then we

must allow that a subject can know everyday propositions while being unable to know all the known consequences of those everyday propositions—i.e., the denials of skeptical hypotheses. In contrast, on this view we can allow both that agents have knowledge of everyday propositions *and* that they can know the denials of radical skeptical hypotheses. Suppose, for example, that the agent does have a safe belief in an everyday proposition. Not only is this belief true in the actual world, but, across the range of nearby possible worlds where she believes this proposition, it is true there as well. Given that this belief really is safe, however, then there will not be any skeptical possible worlds in the realm of nearby possible worlds which determine that safety (henceforth, the “realm of safety”). For if there were such worlds present in the realm of safety, then this would suffice to undermine the agent’s knowledge of the everyday proposition since there would then be a nearby possible world in which the agent still believes the everyday proposition but where this proposition is false (because the skeptical hypothesis is true). But since skeptical possible worlds are now excluded from the realm of safety, it follows that the agent must also have a safe belief that she is not a brain-in-a-vat (or, indeed, the victim of any skeptical hypothesis). The reason for this is that there will be no possible world within the realm of safety in which this proposition is false, and thus, in every world in the realm of safety in which she believes this proposition (which, plausibly, is all of them), her belief is true. Accordingly, skepticism is evaded and closure, as least as it functions in skeptical and anti-skeptical reasoning, is retained.

Moreover, the adoption of safety as a necessary condition on knowledge is also able to speak to the core relevant alternatives

thought that we saw in §2. As expressed there, the thought was that one should be able to have knowledge without having to consider far-fetched, and therefore irrelevant, error-possibilities. This was supposed to be the intuition that Dretske was trying to accommodate with his modal account of knowledge, but, as we saw, he in fact ended up with a slightly different view which did make far-fetched error-possibilities relevant to knowledge, albeit only to knowledge of the denials of skeptical hypotheses (note that it was this element of the view that resulted in the denial of closure). In contrast, a modal interpretation of the relevant alternatives approach that is much closer to this core intuition is that knowledge (*any* knowledge) is only dependent upon one’s ability to track the truth in the relevant range of nearby possible worlds, not also in worlds far away. Accordingly, if the skeptical possibility is indeed far-fetched then it ought to be unable to influence one’s knowledge of everyday propositions or, for that matter, one’s knowledge of the denials of skeptical hypotheses. Safety captures this intuition by allowing agents to have knowledge in both cases provided skeptical possible worlds do not feature in the realm of safety. Dretskean sensitivity, in contrast, violates this intuition by making knowledge dependent not just on the relevant circumstances in nearby worlds but also on the circumstances that obtain in far off worlds (such as skeptical worlds) where the target proposition is false.

Furthermore, since the realm of safety does not vary in response to mere conversational factors, it follows that this is not a semantic contextualist thesis. If the agent does indeed know everyday propositions then, in line with the central Moorean contention, she will also know the denials of radical skeptical hypotheses, and this will be so no matter what

conversational context the agent is in. We thus have a neo-Moorean variety of anti-skepticism which, while keeping to the relevant alternatives spirit of both the Dretskean and the semantic contextualist proposals, lacks the epistemological revisionism of either.³⁸

This is a compelling account of how a neo-Moorean proposal might run, and it certainly does seem to present one very plausible way of reading the core relevant alternatives thesis. It is not quite as novel as it may at first seem, however, since one can trace the beginnings of such view to Gail C. Stine's critical appraisal of the Dretskean thesis back in 1976. She noted that so long as one sticks to the core relevant alternatives conception of relevance then the proper conclusion to be extracted is precisely that we *shouldn't* concede a lack of knowledge of the denials of radical skeptical hypotheses, and thus reject closure. Instead, the conclusion we should draw is that, since such skeptical error-possibilities are indeed modally far off, and thus irrelevant, it follows that we do know their denials after all (to the extent that we know anything much) and thus that closure remains intact. Stine thus reaches a similar conclusion to Sosa's, though without employing the technical modal machinery that Sosa adduces. It is interesting to note, however, that Stine recognised problems with this proposal that Sosa does not mention. In particular, she saw that the difficulty facing this brand of relevant alternatives thesis is to explain how it can be that closure holds and thus that we do know the denials of anti-skeptical hypotheses after all, a problem that we also saw facing the semantic contextualist account above. As Stine herself admits, such a conclusion does indeed "sound odd". In defense of her position she argues that this "oddness" is not due to what is said being false, however (as the Dretske-Nozick line

would suggest), but rather concerns the kind of false conversational implicatures that such a claim to know generate.

Although Stine does not develop this move, it is clearly a maneuver that has a lot of mileage in it since it confronts head-on the worries about the lack of diagnostic appeal of the neo-Moorean approach. Recall that one of the advantages of both the Dretskean and the semantic contextualist line is that they can explain the intuitive appeal of radical skepticism without succumbing to it. The neo-Moorean approach, in contrast, seems to make it a mystery as to why we were ever taken in by this fallacious line of reasoning in the first place. Thus, if it could be shown that the skeptic plays on pragmatic features of our language games with epistemic terms in order to make her arguments superficially plausible, then we could supplement the neo-Moorean theory with a powerful diagnostic account which explains the phenomenology of our engagement with radical skepticism.

Just such an account is offered in Pritchard (2002c; cf. Pritchard 2001e; 2002a). Pritchard claims that we can strengthen Sosa's account of knowledge while keeping true to the basic intuition that drives that notion (he offers an account of knowledge based on the notion of "super-safety"). Moreover, he goes on to show that the kind of mechanisms employed by the semantic contextualist account could just as well be regarded as mechanisms which govern the appropriate assertion of knowledge *claims* rather than as mechanisms which influence the truth-conditions of what is claimed. For example, the thought is that in skeptical conversational contexts it is the standards for *correct assertion* of knowledge claims that is raised (rather than the standards for knowledge), and that this fact explains why it is that we are reluctant to ascribe knowledge in

skeptical contexts even though we are happy to ascribe such knowledge in quotidian contexts. Accordingly, the idea is that semantic contextualism, construed as a thesis about the fluctuating propriety conditions of knowledge claims, could actually be put into the service of the neo-Moorean view to offer the required diagnostic account. Indeed, elsewhere Pritchard (2001e) argues that a similar view can be extracted from some of Wittgenstein's remarks on "hinge" propositions, so this view may well represent the beginnings of a thesis which integrates the two branches of recent work as regards skepticism. Moreover, Pritchard (2002a) also argues that Wright's distinction between transmission and closure can be recast in terms of a distinction between the transference of knowledge across known entailments *simpliciter* (closure), and the transference of knowledge across known entailments where the knowledge retains a certain quality that makes it apt for proper assertion. Wright's intuition about transmission is thus captured within this development of the neo-Moorean model by making use of the pragmatic features of our employment of epistemic terms.

This neo-Moorean approach is so recent that detailed critical assessment of it has not yet appeared. Nevertheless, we can note one possible avenue of discussion here, which is the possibly contentious use that the diagnostic element of this theory makes of the pragmatic/semantic distinction. If this element of the theory could be called into question then this would suffice to scupper the diagnostic component of the theory and thus drastically reduce its plausibility as an anti-skeptical thesis.

9. EPISTEMOLOGICAL EXTERNALISM AND THE NEW SKEPTICS

Before concluding, it is worthwhile to consider those influential figures in the recent epistemological debate who, in contrast to the current mood of optimism that can be found in epistemological discussion of the problem of radical skepticism, are deeply suspicious that any intellectually satisfactory solution could ever be given to this problem. The roots of this movement in the contemporary literature can be traced back to the work of three main figures—Unger (1971; 1975), Stroud (1984; 1989) and Nagel (1986).³⁹ We saw Unger's infallibilist defense of skepticism earlier on, so here we will only summarize the views put forward by Stroud and Nagel, and highlight one way in which this variety of "meta-skepticism" currently informs the skeptical debate, particularly as it figures in more recent work by Stroud (1994; 1996) and Fumerton (1990; 1995).

For both Nagel and Stroud, the thought seems to be that there is something in our philosophical quest for objectivity that inexorably leads us to skeptical conclusions. Nagel argues, for instance, that objectivity involves attaining a completely impartial view of reality, one that is not tainted by any particular perspective. We must, he argues, "get outside of ourselves," and thereby achieve the impossible task of being able to "view the world from nowhere from within it" (Nagel 1986, 76). We realise that the initial appearances present to a viewpoint can be unreliable guides to reality and therefore seek to modify our "subjective" view with a more "objective" perspective that is tempered by reason and reflection. As Nagel points out, however, the trouble with this approach is that

if initial appearances are not in themselves reliable guides to reality, [*then*] why should the products of detached reflection be any

different? Why aren't they . . . equally doubtful . . . ? . . . The same ideas that make the pursuit of objectivity seem necessary for knowledge make both objectivity and knowledge seem, on reflection, unattainable (Nagel 1986, p. 76).

We can reconstruct the argument here as follows. We recognise that our initial unmodified "subjective" experience of the world is unreliable and therefore should be adapted along "objective" lines by eliminating the "subjective" element. For instance, initial appearances tell us, falsely, that straight sticks suddenly become "bent" when placed in water. Accordingly, we modify our initial "subjective" view with the testimony of "objective" scientific investigation which tells us that the stick in fact stays straight, it is just the light that is bending. However, and here is the crux of the matter as far as Nagel is concerned, why do we regard this modified view as being any more reliable than the completely "subjective" perspective that it replaces? Since we cannot eliminate every trace of "subjectivity," it follows that the problematic component of our conception of reality that engendered the pursuit of objectivity in the first place remains. Consequently, we are both aware of the need for objectivity while also recognising that such objectivity is impossible. As a result, according to Nagel, we are condemned to the following pessimistic evaluation of our epistemic capacities:

The search for objective knowledge, because of its commitment to a realist picture, is inescapably subject to skepticism and cannot refute it but must proceed under its shadow. . . . Skepticism . . . is a problem only because of the realist claims of objectivity (Nagel 1986, p. 71).

That is, the problem of skepticism

has no solution, but to recognise that is to come as near as we can to living in the light of truth (Nagel 1986, p. 231).

Moreover, since these "realist" truths concerning objectivity are meant to be inherent in our epistemic concepts, so it is held that this pessimism falls naturally out of any reflective analysis of our epistemic concepts.

Stroud makes similar claims. He writes:

The sceptical philosopher's conception of our position and of his question for an understanding of it . . . is a quest for an objective or detached understanding and explanation of the position we are objectively in. What is seen to be true from a detached "external" standpoint might not correspond to what we take to be the truth about our position when we consider it "internally," from within the practical contexts which give our words their social point. Philosophical scepticism says the two do not correspond; we never know anything about the world around us, although we say or imply that we do hundreds of times a day.

I think we do have a conception of things being a certain way quite independently of their being known or believed or said to be that way by anyone. I think that the source of the philosophical problem of the external world lies somewhere within just such a conception of an objective world or in our desire, expressed in terms of that conception, to gain a certain kind of understanding of our relation to the world. But in trying to describe that conception I think I have relied on nothing but platitudes we would all accept—not about specific ways we all now believe the world to be, but just the general idea of what an objective world or an objective state of affairs would be. If those platitudes about objectivity do indeed express the conception of the world and our relation to it that the sceptical philosopher relies on, and if I am right in thinking that scepticism can be avoided only if that conception is rejected, it will seem that in order to avoid scepticism we must deny platitudes we all accept (Stroud 1984, pp. 81–82).

And note that if responding to skepticism involves denying "platitudes that we would

all accept,” then it follows that any adequate response to the problem of radical skepticism is bound to be intellectually unsatisfactory.⁴⁰

One might wonder, however, exactly how such epistemological pessimism is to undermine the kind of anti-skeptical proposals that we have considered in previous sections. Nearly all of them result in the denial of some key claim that the skeptic makes, and do so on motivated grounds. In what sense, then, must we accept that these proposals, whatever the details, are all going to be intellectually unsatisfactory?

In more recent discussion, one begins to see the emergence of a more definitive pessimistic line, however, which may well be able to give a more compelling expression to this worry. In general, although the point is not always put in these terms, the complaint is that these recent anti-skeptical approaches offer us, at best, an epistemologically externalist response to skepticism when what we wanted was one that functioned within an epistemologically internalist framework.

Simplifying somewhat, we can take epistemological internalism to consist in the claim that there is some substantive necessary condition for knowledge which depends upon facts that the agent is in a position to know by reflection alone. That is, the internalist insists that meeting an appropriate “internal” epistemic condition is necessary for knowledge possession. Externalists, in contrast, demure from this claim and therefore allow that agents might know merely by meeting “external” epistemic conditions. So, for example, externalists tend to allow that small children or unreflective subjects (such as the now notorious, and possibly non-existent, “chicken sexer”) can have knowledge even though they fail to meet an internal epistemic condition. Internalists, in contrast, set the standards for knowledge higher and

thus exclude these agents from possessing knowledge.⁴¹

It is notable that all of the main anti-skeptical accounts that we have discussed so far have tended to side with externalism. For both Dretske and Nozick, for example, the sensitivity condition is a condition that merely needs to be met by the agent for that agent to be a potential knower—it is not further demanded that the agent should have the relevant reflective access to the facts which determine that sensitivity. Similar remarks apply to the semantic contextualist account, with both DeRose and Lewis endorsing externalist views (though Cohen might be an exception in this respect). Finally, both Sosa and Pritchard’s neo-Moorean accounts are described in terms of an externalist epistemology. On closer inspection, it is unsurprising that externalism should be so widely endorsed in this way, though the reasons in each case are different.

The Dretskeans need to be externalists because it is far more problematic to deny closure if knowledge is given an internalist construal. If the denial of closure were not contentious enough, to argue that one could have the kind of “reflective” knowledge at issue in the internalist account of both the antecedent and the entailment and yet lack it of the consequent just seems plain absurd.⁴²

Since both the neo-Mooreans and the contextualists retain closure, it follows that they need not be troubled by this particular motivation for externalism. The driving force behind their adoption of externalism is, in contrast, the fact that they allow (albeit, in the contextualist case, relative to certain contexts) that an agent can have knowledge of the denials of radical skeptical hypotheses. The problem is that this seems to be precisely the sort of knowledge that cannot be possessed under the internalist rubric because skeptical scenarios are,

ex hypothesi, phenomenologically indistinguishable from everyday life. Accordingly, it is plausible to suppose that agents are unable to have the required reflective access to the relevant facts that determine the epistemic status of their beliefs when it comes to these propositions.

Despite the current popularity of externalist accounts of knowledge, however, one might be less sanguine about the prospects for an externalist response to the problem of radical skepticism. The reason for this is that the skeptical puzzle seems to itself strike against the *prima facie* attraction of an externalist account of knowledge. As we saw Craig in effect arguing in §4, if we cannot tell the difference between being in the world that we think we are in now and a skeptical possible world such as the brain-in-a-vat-world, then of what comfort is it to be told that, provided we are in the world that we think we are, we know a great deal? Instead, it would seem that we want some sort of subjective assurance as regards our knowledge that externalist anti-skeptical accounts are not in a position to provide.

It is this line of thinking that ultimately motivates recent work by Stroud (1994; 1996) and Fumerton (1990; 1995). For example, Fumerton makes the following point:

It is *tempting* to think that externalist analyses of knowledge . . . simply remove one *level* of the traditional problems of skepticism. When one reads the well-known externalists one is surely inclined to wonder why they are so sanguine about their supposition that our commonplace beliefs are, for the most part, . . . knowledge. . . . Perception, memory, and induction *may* be reliable processes (in Goldman's sense) and thus given his metaepistemological position we may [*have knowledge of*] the beliefs they produce but, the sceptic can argue, we have no reason to believe that these process are

reliable and thus even if we accept reliabilism, we have no reason to think that the beliefs they produce [*constitute knowledge*] (Fumerton 1990, p. 63).⁴³

In effect, the complaint that Fumerton is giving expression to here is that externalism allows that there are certain conditions on knowledge that we are unable to reflectively determine have obtained. Indeed, Fumerton is more explicit about the focus of his objection when he goes on to write that

the main problem with externalist accounts, it seems to me, just is the fact that such accounts . . . develop concepts of knowledge that are irrelevant. . . . The philosopher doesn't just want true beliefs, or even reliably produced beliefs, or beliefs caused by the facts that make them true. The philosopher wants to have the relevant features of the world directly before consciousness (Fumerton 1990, p. 64).

Presumably, to argue that externalist accounts of knowledge are problematic because they fail to demand that the relevant facts should be "directly before consciousness" is simply to complain that such theories make the satisfaction of non-reflectively accessible external epistemic conditions central to knowledge possession.

Fumerton is not the only one to put forward objections to externalism that run along these lines, though he is perhaps the most explicit about what the complaint that he is giving voice to amounts to. For example, a similar argument against externalism seems to be implicit in the following passages from Stroud:

suppose there are truths about the world and the human condition which link human perceptual states and cognitive mechanisms with further states of knowledge and reasonable belief, and which imply that human beings acquire their beliefs about the physical world through the operation of belief-forming mechanisms which are on the whole reliable

in the sense of giving them mostly true beliefs. . . . If there are truths of this kind . . . that fact alone obviously will do us no good as theorists who want to understand human knowledge in this philosophical way. At the very least we must believe some such truths; their merely being true would not be enough to give us any illumination or satisfaction. But our merely happening to believe them would not be enough either. We seek understanding of certain aspects of the human condition, so we seek more than just a set of beliefs about it; we want to know or have good reasons for thinking that what we believe about it is true (Stroud 1994, p. 297).⁴⁴

It is difficult to understand Stroud's objection here if it is not to be construed along similar lines to that found in the passages from Fumerton cited above. Stroud's thought seems to be that it is not enough merely to meet the external epistemic conditions that give us knowledge, rather we should also have the special kind of internal access to those conditions that the internalist demands (and perhaps even more than that).⁴⁵

A new debate is thus emerging in the literature which is to some extent orthogonal to the key discussions that we have considered so far since it questions the very presuppositions of those discussions.

Future work on the skeptical puzzle will thus enjoin participants to not only meet the skeptical argument in a motivated fashion, but also respond to this meta-epistemological challenge.

10. CONCLUDING REMARKS

As we have seen, there are two main currents in contemporary treatments of the problem of radical skepticism—one that emerges out of the relevant alternatives account of knowledge proposed by Dretske and Nozick, and another which takes its stimulus from remarks made by Wittgenstein in his final notebooks. The recent confluence of these two currents will almost certainly produce interesting work on this topic in the near future. More significantly still, however, is the focussed manner in which the externalism/internalism debate in epistemology—which has itself been a central theme in the contemporary epistemological literature—has begun to influence treatment of this issue. For what this promises to achieve is to bring debates about radical skepticism back into more general epistemological discussion, and this can only be a helpful development.⁴⁶

University of Stirling

NOTES

1. This description can be found in, for example, Williams (1991, chapter 1).
2. For more on this debate, see Yalçın (1992) and Pritchard (2002d). Relatedly, those varieties of skepticism—such as those which employ the “dreaming” skeptical hypothesis—that are not so readily couched in terms of the closure-based template will not be discussed. This point is discussed further below, in footnote 11.
3. For more on this debate, see Pritchard (2001c) and the exchange between Fogelin (1999) and Williams (1999).
4. For example, even if it did follow from the Principle of Charity as Davidson (1983) understands it that it is an *a priori* truth that “belief is by its nature veridical” (314), this would not answer those skeptical challenges which allow widespread truth in our beliefs and focus instead upon the epistemic status of those beliefs. For more on the semantic style of response to (certain

forms of) radical skepticism, see Williams (1988–1989); Craig (1990a); Brueckner (1992b); Christensen (1993); Wright (1994; 2000a); Forbes (1995); Warfield (1998); and the exchange between DeRose (2000) and Williamson (2000b).

5. DeRose (2000) argues for a similar conclusion. Another element of this treatment of the radical skeptical debate that is perhaps worth commenting on is the focus on those skeptical arguments that are directed at our possession of *knowledge* rather than on those skeptical arguments that aim to undermine our putative possession of other epistemic properties for our beliefs, such as justification. The motivation for doing this is twofold. First, because most recent discussions of radical skepticism tend to focus on knowledge in this respect. Second (and this might explain the first point), any radical skeptical argument directed at a subsidiary epistemic property like justification will almost certainly result in a radical skeptical argument against knowledge possession by default and so we can legitimately side-step the complication of expressing the argument in terms of the subsidiary property. That this is so is only on the assumption that the property in question is a necessary component of knowledge, of course, so there is room for a separate skeptical debate where this is not so. However, skeptical arguments that are directed at epistemic properties that are not essential components of knowledge are not nearly so relevant to the issue of radical skepticism than skeptical arguments that are directed at knowledge or some essential subsidiary epistemic property. For more discussion of the issue of the difference between skeptical arguments directed against justification and those directed at knowledge, see Wright (1991, 87–89) and the response by Pritchard (2001d, pp. 101–106). For a general discussion of the constraints that govern cogent anti-skeptical proposals, see Throop (1998) and Pritchard (2001d).

6. It is not important to the argument that it is cashed-out in the first-person in this way, since skeptical arguments are meant to apply to *any* agent (and, indeed, at *any* time). The template conforms to this first-person construal solely for the sake of simplicity.

7. We can plausibly disregard any response to skepticism that depends upon a denial of *modus ponens* since the skeptic could quite reasonably respond to this move by arguing that it merely represents a *reductio* of the anti-skeptical position.

8. Another influential statement of what could be termed a relevant alternatives view which is roughly contemporaneous with Dretske's is Goldman (1976).

9. For more on infallibilism and its role in skeptical arguments, see Craig (1990b) and Pritchard (2000c). Lewis gives a neat account of the infallibilist motivation for radical skepticism in the following passage:

The sceptical argument is nothing new or fancy. It's just this: it seems as if knowledge must be by definition infallible. If you claim that *S* knows that *P*, and yet you grant that *S* cannot eliminate a certain possibility in which not-*P*, it certainly seems as if you have granted that *S* does not after all know that *P*. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory (Lewis 1996, p. 549).

10. One might want to amend this characterization of closure in a number of ways in order to evade certain types of concern that one could raise about it. For example, if one thinks that one could know the antecedent proposition and the entailment without thereby even *believing* the consequent proposition, then one might want to incorporate into this formulation some demand that the consequent proposition was believed. After all, since belief is, plausibly, a necessary condition for knowledge, a failure to believe the consequent proposition in such a case would constitute a counterexample to closure. Crucially, however, this kind of “doxastic” counterexample would be very different from the kind of epistemically motivated counterexamples that are at issue in the contemporary skeptical debate, and thus it is best to keep them apart. Though potentially imprecise, this formulation of closure should suffice for our purposes here.

11. For more on the role of closure in radical skeptical arguments, see Pritchard (2002a; 2002b). As noted above, although it is commonly accepted that it is closure that underlies the radical skeptical argument, there have been some suggestions in the literature that what motivates the skeptical reasoning is in fact a different epistemic principle, what has been called an ‘underdetermination’ principle. For discussion of this view, see Yalçın (1992); Brueckner (1994); Cohen (1998c); and Pritchard (2002d). Nevertheless, although there is a substantial difference between these commentators regarding the exact nature and role of the underdetermination principle in skeptical arguments, there is still a general consensus that denying closure would suffice to meet the skeptical argument.

The reader should also note that not every radical skeptical argument will directly conform to the closure-based template. The reason for this is that some skeptical hypotheses—such as, in particular, the hypothesis that one might be dreaming—do not entail the denials of everyday propositions. It could be, for example, that one is dreaming *and* that one has two hands. Such arguments demand a modification to the template skeptical argument advanced above which captures the fact that although the denial of this hypothesis is not entailed by the *truth* of everyday propositions, it is entailed by an agent’s *knowledge* of these everyday propositions. For discussion of dreaming skepticism in this regard, see Wright (1991) and Pritchard (2001f).

12. Dretske (1970) calls epistemic operators “semi-penetrating.”

13. Dretske actually says very little about whether the entailment is known. We can, however, reasonably suppose that agents can come to know such trivial analytic consequences of what they believe. Besides, if Dretske is to be seen as making a substantive point here then we ought to be able to assume that the entailment is known. That knowledge might not transfer across entailments which are not known is not nearly so controversial a thesis.

14. The reader should note that this characterization of Dretskean sensitivity represents a more simplified version than that which Dretske himself offers. For example, Dretske restricts the range of possible worlds at issue by stipulating that they must be worlds in which the “evidence-base” which produces the belief is the same (a similar move, this time cashed-out in terms of methods, is incorporated into Nozick’s (1981) version of this thesis). Nevertheless, this characterization supplies all the details that we need for the current discussion. Note also that it is of course important that Dretskean sensitivity only applies to *contingent* propositions, because, as regards necessary propositions, there is no nearest possible world in which they are false. Without this restriction, a stubborn belief in a necessary proposition would thus count as knowledge no matter how that belief was formed.

15. Of course, by taking this modal line Dretske’s view in this respect is hostage to a suitable account being given of the semantics of possible worlds. Two problems, in particular, affect any modal characterization of knowledge along these lines. The first is that it is far from clear that possible worlds are ordered in the manner that Dretske and others seem to suppose (the so-called “World Order” thesis). At the very least, there does seem to be a certain vagueness inherent in much of the potential orderings of possible worlds (though this issue is less troubling in this context because skeptical possible worlds are so clearly very different from non-skeptical worlds). The second key problem is that the very idea of a nearest possible world that is at the heart of the Dretskean account is contentious (the so-called ‘World Border’ thesis). For some of the key texts on possible worlds, see the papers collected in Loux (1979).

16. Nozick makes a number of changes to the basic Dretskean thesis, the most fundamental of which is the introduction of an extra modal condition which demands that the agent’s belief should track the truth of the proposition in question in the nearest possible world or worlds in which it continues to be true. An extended discussion of Nozick’s account is not undertaken here because although the Nozickean view has a number of important epistemological ramifications

that are distinct from those generated by the Dretskean thesis, these ramifications do not have any direct bearing on the skeptical argument as it is understood here. For discussion specific to Nozick's account, see the papers collected in the excellent anthology edited by Luper-Foy (1987).

17. For more on Craig's complaint against the Dretske-Nozick line (which was, it should be noted, explicitly focussed on Nozick's account), see Brueckner (1991) and Pritchard (2000c).

18. For an attack on the Nozickean anti-skeptical proposal on grounds that are, primarily, internalist, see Bonjour (1987).

19. For some of the key discussions of the relationship between skepticism and closure, see Thalberg (1974); Stine (1976); Brueckner (1986); Vogel (1990); Williams (1991, chapter eight); Hales (1995); Klein (1995); Vahid (1995; cf. Brueckner 1998); and Pritchard (2002a; 2002b).

20. Interestingly, Klein (1987) also offers a defense of closure that runs along reliabilist (and thus non-internalist) lines.

21. Though, of course, "tracking" here cannot mean the same thing as it does in the Dretske-Nozick account because this conception of tracking is allied to the relevant alternative thesis that results in the denial of closure. We will consider an account of tracking more suited to this alternative relevant alternative approach in §9.

22. For the main discussions of semantic contextualism, see Schiffer (1996); Rieber (1998); Feldman (1999); Heller (1999b); Sosa (1999a; 2000); Vogel (1999); Fogelin (2000); Kornblith (2000); and Pritchard (2001b). See also the survey article by Norman (1999). For a discussion of Cohen's particular version of semantic contextualism (which is subtly different, in important respects, from that advanced by DeRose and Lewis), see the summary offered by Cohen (2000a), the critical responses by Hawthorne (2000), Klein (2000) and Prades (2000), and Cohen's (2000b) response. Earlier forms of epistemological contextualism were not specifically directed at the skeptical problem at all, nor were they expressed in explicitly conversationalist terms. See, in particular, Annis (1978). For discussion of Annis' view, see Airaksinen (1982; cf. Annis 1982) and Henderson (1994). A different form of contextualism that is specifically directed at skepticism, advanced by Williams (1991), is discussed below.

23. Lewis agrees. In Lewis (1979) he briefly applies this line of argument to epistemic terms (and thus to the skeptical problem), and expands upon this thought in Lewis (1996).

24. The issue of what is involved in bringing about a change of conversational context is still somewhat controversial. Most contextualists agree, however, that merely mentioning a skeptical hypothesis will, at least sometimes, suffice.

25. For more on the relationship between arguments for contextualism and arguments for non-closure, see the exchange between Heller (1999a) and Pritchard (2000a).

26. Indeed, Lewis (1995, 551) has been one commentator who has explicitly denied that justification is either a necessary or a sufficient condition for knowledge.

27. For specific discussion of Unger's claim in this respect, see Craig (1990b, appendix). For a critique of the semantic contextualist account that also attacks its "hierarchical" conception of epistemic contexts, see Pritchard (2001b), who argues that while this form of contextualism offers a vivid account of how skepticism can raise the epistemic standards and thereby destroy our knowledge, it is far from clear that it can offer an equally compelling account of how the standards can then be lowered again so that knowledge is restored. Putnam (1998) offers a related criticism of contextualism in this respect, though his point is directed at Williams's (1991) version of the view. A further issue in this regard is whether semantic contextualists would be willing to allow that there are some contexts where the standards are so low as to permit mere true belief

to be knowledge. This clearly would be an unpalatable consequence of the view, but it is far from clear what principled grounds a semantic contextualist could offer to disallow it. I am grateful to an anonymous referee from *American Philosophical Quarterly* for pointing this line of criticism out to me.

28. For a commentary specific to DeRose's proposal, see Williamson (2000b).

29. Although the "hinge" metaphor is the dominant symbolism in *On Certainty*, it is accompanied by various other metaphors such as the following: that these propositions constitute the "scaffolding" of our thoughts (§211); that they form the "foundations of our language-games" (§§401–403); and also that they represent the implicit "world-picture" from within which we inquire, the "inherited background against which [we] distinguish between true and false" (§§94–95).

30. Moreover, it is not as if our belief in hinge propositions is epistemically supported directly, because of some special epistemic access that we have to these propositions (as would be the case if they were foundational propositions). In the relevant sense, then, hinge propositions are neither foundational propositions that do not require evidential buttress nor evidentially grounded and herein lies the challenge that they pose.

31. This may have been the original intention, of course, since in his later writings Wittgenstein was far less inclined to offer solutions to traditional philosophical problems. For a subtle discussion of Wittgenstein's *On Certainty* which interprets him as merely *dissolving*, rather than solving, the problem of radical skepticism, see McGinn (1989). Indeed, on this view hinge propositions are unknowable precisely because, strictly speaking, they are not propositions at all (and so there is nothing to be known). For more on this interpretation, see Conant (1998) and the response by McGinn (2002). For a more neutral interpretation of the Wittgensteinian treatment of skepticism, see Pritchard (2001e; cf. Pritchard 2000d). For an account of how the Wittgensteinian line might influence the debate surrounding skepticism regarding the possibility of religious knowledge, see Pritchard (2000b).

32. For discussion of Wright's so-called "implosion" of dreaming skepticism, see Brueckner (1992a); Tymoczko & Vogel (1992); and Pritchard (2001f). It should be noted that Wright (1985) adopts the basic hinge proposition line noted above that these propositions cannot be known, but he restricts this thesis to certain propositions which, he argues, are "non-factive." These are propositions such as "There is an external world" or "The universe has not just come into existence" that Wright claims are so immune to the tribunal of empirical evaluation that it is difficult to see how they could possibly express truths. Wright thus interweaves the Wittgensteinian thesis with an anti-realist semantics. For more on Wright's general anti-realist approach, see Wright (1992).

33. Pritchard (2000d; 2001e) offers a reading of the Wittgensteinian thesis that preserves the quotidian conception of hinge propositions.

34. It should be noted that Wright attempts to extend this analysis beyond radical skeptical concerns to apply to such problems as the "McKinsey" incompatibility between semantic externalism and first-person authority, and the argument offered by Putnam which purports to show that, necessarily, the claim that we are brains-in-vats is false. For discussion of Wright's proposal, see Pritchard (2002a), and the critical appraisals offered by Hale (2000); Sainsbury (2000); and Suárez (2000), along with Wright's (2000b) reply.

35. Williams (1991, chapters 7 & 8), puts forward a sophisticated set of arguments to the effect that even supposedly non-foundationalist epistemologies such as coherence theory or reliabilism involve an implicit commitment to the basic foundationalist structure.

36. For some of the key papers that assess Williams's inferential contextualist view, see the critical reviews by Rorty (1997) and Vogel (1997), along with Williams's response (1997), and the

papers by Brady (1998) and Putnam (1998). See also the exchange between Stroud (1996) and Williams (1996). Pritchard (2002e) offers a detailed comparison of the two types of contextualist theories.

37. Sosa (1999a, 142) actually offers three different characterizations of this notion, but we can plausibly take the formulation offered here to capture the spirit of all three. For (slightly) more on Sosa's view in this respect, see Sosa (2000a). For critique of this latter paper, and Sosa's response, see Kornblith (2000); Lehrer (2000); Sosa (2000b); and Tomberlin (2000). See also Pritchard (2002c). Vogel (2000) also offers a critique of a position that is roughly configured along the lines that Sosa envisages.

38. Sosa's position as regards skepticism is closely related to his general endorsement of a virtue-theoretic theory of knowledge which analyses knowledge in terms of the reliable functioning of cognitive capacities in suitable environments. For more on Sosa's theory in this respect, and some of the key critical appraisals of the general virtue-epistemic line, see Sosa (1985; 1991; 1993); Greco (1993; 1999); and Axtell (1997). Greco (1994; 2000) offers a subtle account of how a virtue-theoretic epistemology can respond to the problem of radical skepticism which is essentially run along similar lines to that proposed by Sosa, although his does not express his point in the same neo-Moorean terms.

39. This is a very selective list, since one finds such pessimism in recent work by a number of influential figures. Cavell, for example, writes that:

the truth of skepticism, or what I might call the moral of skepticism, . . . [is] that the human creature's basis in the world as a whole, its relation to the world as such, is not that of knowing, anyway not what we think of as knowing" (Cavell 1979, 214).

Similarly, McGinn (1979, p. 115), goes so far as to claim that "a *prima facie* vulnerability to . . . a [*skeptical*] challenge should be regarded as an adequacy condition which any form of realism should be required to meet." More recently still, Heil (1998) has argued that, strictly speaking, skepticism is unanswerable because of the inherent, though necessary, realist assumptions that we must make. In a related fashion, we find even those who are usually completely unsympathetic to the skeptical question claiming, in an off-hand way, that "The Humean predicament is the human predicament" (Quine 1969, p. 72). That is, that there is no real escape from Hume's arguments which purport to show a lack of epistemic support for our most basic and cherished beliefs. Quine's particular "spin" on this Humean pessimism is derived from his underdetermination of theory with respect to data thesis, which leads him to comment on the "hypothesis" of physical objects (Quine 1969; cf. Quine 1960, p. 22). For a discussion of Quine's contribution to this debate see Stroud (1981; 1984, chapter VI; cf. Quine 1981); Hookway (1990); and Williams (1991, pp. 254–265). For an illuminating discussion of the new skeptics, see Williams (1991, chapter 1). For an incisive discussion of some of the recent claims that have been made regarding the relationship between the "realist" intuitions at stake here and the problem of radical skepticism, see Moser (1999; cf. Moser 1993).

40. It is interesting to note that the kind of skeptical worry expressed by Nagel and Stroud seems to also be present in Williams's sympathetic account of the Cartesian skeptical enterprise. he writes:

knowledge does have a problematic character, . . . something in it which offers a standing invitation to scepticism (B. Williams 1978, 64).

This "something" that makes this invitation is simply this inherent subjective element in our knowledge. As Williams goes on to argue, the skeptical threat arises from:

the very basic thought that if knowledge is what it claims to be, then it is knowledge of a reality that exists independently of that knowledge, and indeed (except for the special case where the reality known happens to be some psychological item) independently of thought or experience. Knowledge of what is there *anyway* (Williams 1978, p. 64).

41. For the key discussions of the externalism/internalism distinction, see the excellent anthology edited by Kornblith (2001). For a discussion of how this distinction relates to metaepistemological skepticism, see Pritchard (2001g). For related discussions of how the internalist thesis is implicated in the skeptical paradox, see Williams (1991), Hill (1996), and Sosa (1999b).

42. That said, Pritchard (2002b; cf. 2002a) argues that the Dretskean denial of closure is, properly understood, an argument for the denial of an analogue closure principle expressed in terms of the internalist conception of knowledge.

43. Fumerton develops this line of attack at length in Fumerton (1995). For discussion of Fumerton's general position in this respect, see the précis offered in Fumerton (1998a), and the discussions, and Fumerton's response, in Cohen (1998b); Fumerton (1998b); Hookway (1998); and Klein (1998).

44. Stroud was explicitly responding here to Sosa (1994).

45. For more on this kind of complaint against externalist epistemologies, see Foley (1993). Sometimes this point is put in terms of the so-called "iterativity" principle, the principle that (roughly) if one knows then one knows that one knows. The thought is that we do not merely require an analysis of knowledge which ensures that we can know much of what we think we know, but also one that ensures that, at least for the most part, we are able to know that we have such knowledge. It is often contended that only an internalist conception of knowledge can provide support for this principle. For critical discussion of this claim, and of the iterativity principle in general, see Alston (1980); Greco (1990); and Pritchard (2001g; cf. Pritchard 2001a).

46. I am grateful to the following people who have at some point read and commented on various aspects of this research: Lars Binderup; Michael Brady; Edward Craig; Patrick Greenough; Lars-Bo Gundersen; Katherine Hawley; Jesper Kallestrup; Patrice Philie; Stephen Read; Sven Rosenkranz; Leslie Stevenson; Peter Sullivan; Michael Williams; and, especially, Crispin Wright. Various sections of this article have been presented in some form at the following philosophy departments: Edinburgh; Glasgow; Nottingham; Oxford; Reading; St. Andrews; Sheffield; Southampton; and Stirling; as well as at the 23rd annual Wittgenstein symposium in Kirchberg, Austria. I am grateful to all these audiences. Finally, I would like to thank an anonymous referee from *American Philosophical Quarterly* for an extremely useful series of suggestions.

LIST OF REFERENCES

- Airaksinen, T. 1982. "Contextualism: A New Theory of Epistemic Justification?" *Philosophia*, vol. 12, pp. 51–55.
- Alston, W. P. 1980. "Levels Confusions in Epistemology." *Midwest Studies in Philosophy Vol. 5: Studies in Epistemology*. Ed. P. A. French, T. E. Uehling, and H. K. Wettstein. Minneapolis: University of Minnesota Press, pp. 135–150.
- Annis, D. B. 1978. "A Contextualist Theory of Justification." *American Philosophical Quarterly*, vol. 15, pp. 213–219.
- _____. 1982. "The Social and Cultural Component of Epistemic Justification: A Reply." *Philosophia*, vol. 12, pp. 55–60.

- Austin, J. L. 1961. "Other Minds." Reprinted in his *Philosophical Papers*. Ed. J. O. Urmson and G. J. Warnock. Clarendon Press, Oxford, England, pp. 44–84.
- Axtell, G. 1997. "Recent Work on Virtue Epistemology." *American Philosophical Quarterly*, vol. 34, pp. 1–26.
- Baldwin, T. 1990. *G. E. Moore*. London: Routledge.
- _____. 1993. "G. E. Moore." *A Companion to Epistemology*. Ed. J. Dancy and E. Sosa. Oxford: Basil Blackwell, pp. 283–285.
- Brady, M. S. 1998. "Can Epistemic Contextualism Avoid the Regress Problem?" *The Southern Journal of Philosophy*, vol. 36, pp. 317–328.
- Boghossian, P. 1989. "Content and Self-Knowledge." *Philosophical Topics*, vol. 17, pp. 5–26.
- _____. 1997. "What an Externalist can Know *A Priori*." *Proceedings of the Aristotelian Society*, vol. 97, 161–75.
- Bonjour, L. 1987. "Nozick, Externalism, and Skepticism." *The Possibility of Knowledge: Nozick and his Critics*. Ed. S. Luper-Foy. Totowa, N.J.: Rowman & Littlefield, chapter 11.
- Brower, B. W. 1998. "Contextualism, Epistemological." *Routledge Encyclopædia of Philosophy*, vol. 2. Ed. E. Craig. London: Routledge, pp. 646–650.
- Brueckner, A. 1986. "Skepticism and Epistemic Closure." *Philosophical Topics*, vol. 13, pp. 89–117.
- _____. 1991. "Unfair to Nozick." *Analysis*, vol. 51, pp. 61–64.
- _____. 1992a. "Problems with the Wright Route to Scepticism." *Mind*, no. 402, pp. 309–317.
- _____. 1992b. "Semantic Answers to Skepticism." *Pacific Philosophical Quarterly*, vol. 73, pp. 200–219.
- _____. 1994a. "The Shifting Content of Knowledge Attributions." *Philosophy and Phenomenological Research*, vol. 54, pp. 123–126.
- _____. 1994b. "The Structure of the Skeptical Argument." *Philosophy and Phenomenological Research*, vol. 54, pp. 827–835.
- _____. 1998. "Closure and Context." *Ratio* (new series), vol. 11, pp. 78–82.
- Cavell, S. 1979. *The Claim of Reason: Wittgenstein, Skepticism, Morality, and Tragedy*. New York: Oxford University Press.
- Christensen, D. 1993. "Skeptical Problems, Semantical Solutions." *Philosophy and Phenomenological Research*, vol. 53, pp. 301–321.
- Cohen, S. 1986. "Knowledge and Context." *Journal of Philosophy*, vol. 83, pp. 574–583.
- _____. 1987. "Knowledge, Context, and Social Standards." *Synthese*, vol. 73, pp. 3–26.
- _____. 1988. "How to be a Fallibilist." *Philosophical Perspectives*, vol. 2, pp. 91–123.
- _____. 1990. "Skepticism and Everyday Knowledge Attributions." *Doubting: Contemporary Perspectives on Scepticism*. Ed. G. Ross and M. D. Roth. Dordrecht, Netherlands: Kluwer Academic Publishers, pp. 161–169.
- _____. 1991. "Skepticism, Relevance, and Relativity." *Dretske and his Critics*. Ed. B. McLaughlin. Oxford: Basil Blackwell, pp. 17–37.
- _____. 1998a. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy*, vol. 76, pp. 289–306.
- _____. 1998b. "Fumerton on Metaepistemology and Skepticism." *Philosophy and Phenomenological Research*, vol. 58, pp. 913–918.
- _____. 1998c. "Two Kinds of Skeptical Argument." *Philosophy and Phenomenological Research*, vol. 58, pp. 143–159.
- _____. 1999. "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives*, vol. 13, pp. 57–90.
- _____. 2000a. "Contextualism and Skepticism." *Philosophical Issues*, vol. 10, pp. 94–107.
- _____. 2000b. "Replies." *Philosophical Issues*, vol. 10, pp. 132–139.

- Conant, J. 1998. "Wittgenstein on Meaning and Use." *Philosophical Investigations*, vol. 21, pp. 222–250.
- Craig, E. 1989. "Nozick and the Sceptic: The Thumbnail Version." *Analysis*, vol. 49, pp. 161–162.
- _____. 1990a. "Davidson and the Sceptic: The Thumbnail Version." *Analysis*, vol. 50, pp. 213–214.
- _____. 1990b. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Davidson, D. 1983. "A Coherence Theory of Truth and Knowledge." Reprinted in *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Ed. E. LePore, Oxford: Basil Blackwell (1986), chapter 16.
- Davies, M. 1998. "Externalism, Architecturalism, and Epistemic Warrant." *Knowing Our Own Minds: Essays on Self-Knowledge*. Ed. C. J. G. Wright, B. C. Smith, and C. Macdonald, Oxford: Oxford University Press.
- DeRose, K. 1990. "Epistemic Possibilities." *Philosophical Review*, vol. 100, pp. 581–605.
- _____. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research*, vol. 52, pp. 913–929.
- _____. 1995. "Solving the Skeptical Problem." *Philosophical Review*, vol. 104, pp. 1–52.
- _____. 1996. "Relevant Alternatives and the Content of Knowledge Attributions." *Philosophy and Phenomenological Research*, vol. 56, pp. 193–197.
- _____. 1999a. "Contextualism: An Explanation and Defence." *Epistemology*. Ed. J. Greco and E. Sosa. Oxford: Basil Blackwell, pp. 187–205.
- _____. 1999b. "Responding to Skepticism." *Skepticism: A Contemporary Reader*. Ed. K. DeRose and T. Warfield. Oxford: Oxford University Press, pp. 1–26.
- _____. 2000. "How Can We Know that We're Not Brains in Vats?" *The Southern Journal of Philosophy*, vol. 38, pp. 121–148.
- Dretske, F. 1970. "Epistemic Operators." *Journal of Philosophy*, vol. 67, pp. 1007–1023.
- _____. 1971. "Conclusive Reasons." *Australasian Journal of Philosophy*, vol. 49, pp. 1–22.
- _____. 1981. "The Pragmatic Dimension of Knowledge." *Philosophical Studies*, vol. 40, pp. 363–378.
- _____. 1991. "Knowledge: Sanford and Cohen." *Dretske and his Critics*. Ed. B. McLaughlin. Oxford: Basil Blackwell, pp. 185–196.
- Feldman, R. 1999. "Contextualism and Skepticism." *Philosophical Perspectives*, vol. 13, pp. 91–114.
- Fogelin, R. 1994. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press.
- _____. 1999. "The Sceptic's Burden." *International Journal of Philosophical Studies*, vol. 7, pp. 159–172.
- _____. 2000. "Contextualism and Externalism: Trading in One Form of Skepticism for Another." *Philosophical Issues*, vol. 10, pp. 43–57.
- Foley, R. 1993. *Working Without A Net*. New York: Oxford University Press.
- Forbes, G. 1995. "Realism and Skepticism: Brains in a Vat Revisited." *Journal of Philosophy*, vol. 92, pp. 205–222.
- Fumerton, R. 1988. "The Internalism/Externalism Controversy." *Philosophical Perspectives*, vol. 2, pp. 443–459.
- _____. 1990. "Metaepistemology and Skepticism." *Doubting: Contemporary Perspectives on Skepticism*. Ed. M. D. Roth and G. Ross. Dordrecht, Netherlands: Kluwer Academic Publishers, pp. 57–68.
- _____. 1995. *Metaepistemology and Skepticism*. Lanham, Md.: Rowman & Littlefield.
- _____. 1998a. "Précis of *Metaepistemology and Skepticism*." *Philosophy and Phenomenological Research*, vol. 58, pp. 905–906.

- _____. 1998b. "Replies." *Philosophy and Phenomenological Research*, vol. 58, pp. 927–938.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy*, vol. 73, pp. 771–791.
- Greco, J. 1990. "Internalism and Epistemically Responsible Belief." *Synthese*, vol. 85, pp. 245–277.
- _____. 1993. "Virtues and Vices of Virtue Epistemology." *Canadian Journal of Philosophy*, vol. 23, pp. 413–432.
- _____. 1994. "Virtue Epistemology and the Relevant Sense of 'Relevant Possibility.'" *The Southern Journal of Philosophy*, vol. 32, pp. 61–77.
- _____. 1999. "Agent Reliabilism." *Philosophical Perspectives*, vol. 13, pp. 273–296.
- _____. 2000. *Putting Sceptics in Their Place: The Nature of Sceptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge University Press.
- Grice, H. P. 1989. "Logic and Conversation." Reprinted in his *Studies in the Way of Words*. Cambridge, Mass. Harvard University Press.
- Hale, B. 2000. "Transmission and Closure." *Philosophical Issues*, vol. 10, pp. 172–190.
- Hales, S. D. 1995. "Epistemic Closure Principles." *The Southern Journal of Philosophy*, vol. 33, pp. 185–201.
- Hambourger, R. 1987. "Justified Assertion and the Relativity of Knowledge." *Philosophical Studies*, vol. 51, pp. 241–269.
- Hawthorne, J. 2000. "Reply to Cohen." *Philosophical Issues*, vol. 10, pp. 117–120.
- Heller, M. 1989. "Relevant Alternatives." *Philosophical Studies*, vol. 55, pp. 23–40.
- _____. 1999a. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy*, vol. 77, pp. 196–208.
- _____. 1999b. "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives*, vol. 13, pp. 115–130.
- Henderson, D. K. 1994. "Epistemic Competence and Contextualist Epistemology: Why Contextualism is not Just the Poor Person's Coherentism." *Journal of Philosophy*, vol. 91, pp. 627–649.
- Heil, J. 1998. "Skepticism and Realism." *American Philosophical Quarterly*, vol. 35, pp. 57–72.
- Hill, C. 1996. "Process Reliabilism and Cartesian Scepticism." *Philosophy and Phenomenological Research*, vol. 56, pp. 567–581.
- Hookway, C. 1990. *Scepticism*. London: Routledge.
- _____. 1998a. "Normative Concepts and Epistemological Internalism." *Philosophy and Phenomenological Research*, vol. 58, pp. 907–912.
- Klein, P. 1981. *Certainty*. Minneapolis: University of Minnesota Press.
- _____. 1987. "On Behalf of the Sceptic." In *The Possibility of Knowledge*. Ed. S. Luper-Foy. Totowa, N.J.: Rowman & Littlefield, pp. 267–281.
- _____. 1995. "Skepticism and Closure: Why the Evil Genius Argument Fails." *Philosophical Topics*, vol. 23, pp. 213–236.
- _____. 1998a. "Foundationalism and the Infinite Regress." *Philosophy and Phenomenological Research*, vol. 58, pp. 919–926.
- _____. 2000. "Contextualism and the Real Nature of Academic Skepticism." *Philosophical Issues*, vol. 10, pp. 108–116.
- Kornblith, H. 2000. "The Contextualist Evasion of Epistemology." *Philosophical Issues*, vol. 10, pp. 24–32.
- _____. 2001. *Epistemology: Internalism and Externalism*. Oxford: Basil Blackwell.
- Lehrer, K. 2000. "Sensitivity, Indiscernibility, and Knowledge." *Philosophical Issues*, vol. 10, pp. 33–37.
- Lewis, D. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic*, vol. 8, pp. 339–359.
- _____. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy*, vol. 74, pp. 549–567.

- Loux, M. 1979. *The Possible and the Actual*. Ithaca, N.Y.: Cornell University Press.
- Luper-Foy, S. 1987. *The Possibility of Knowledge: Nozick and his Critics*. Totowa, N.J.: Rowman & Littlefield.
- McGinn, C. 1979. "An *A Priori* Argument for Realism." *Journal of Philosophy*, vol. 82, pp. 28–52.
- McGinn, M. 1989. *Sense and Certainty: A Dissolution of Scepticism*. Oxford: Basil Blackwell.
- _____. 2002. "What Kind of Senselessness is This? A Reply to Conant on Wittgenstein's Critique of Moore." *Typescript*.
- McKinsey, M. 1991. "Anti-Individualism and Privileged Access." *Analysis*, vol. 51, pp. 9–16.
- Moore, G. E. 1925. "A Defence of Common Sense." *Contemporary British Philosophy* (2nd series), Ed. J. H. Muirhead. London: Allen and Unwin. Reprinted in his *Philosophical Papers*. London: Allen & Unwin (1959).
- _____. 1939. "Proof of an External World." *Proceedings of the British Academy*, vol. 25. Reprinted in his *Philosophical Papers*. London: Allen & Unwin (1959).
- Moser, P. K. 1993. *Philosophy After Objectivity*. New York: Oxford University Press.
- _____. 1999. "Realism, Objectivity, and Skepticism." In *Epistemology*, ed. J. Greco and E. Sosa. Oxford: Basil Blackwell, pp. 70–91.
- Nagel, T. 1986. *The View from Nowhere*. Oxford: Oxford University Press.
- Norman, A. P. 1999. "Contextualism: Its Past, Present, and Prospects." *Philosophia*, vol. 27, pp. 1–28.
- Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.
- Pollock, J. 1986. *Contemporary Theories of Knowledge*. Totowa, N.J.: Rowman & Littlefield.
- Prades, J. L. 2000. "Scepticism, Contextualism and Closure." *Philosophical Issues*, vol. 10, pp. 121–131.
- Pritchard, D. H. 2000a. "Closure and Context." *Australasian Journal of Philosophy*, vol. 78, pp. 275–280.
- _____. 2000b. "Is 'God Exists' a 'Hinge' Proposition of Religious Belief?" *International Journal for Philosophy of Religion*, vol. 47, pp. 129–140.
- _____. 2000c. "Understanding Scepticism." *SATS: Nordic Journal of Philosophy*, vol. 1, pp. 107–124.
- _____. 2000d. "Wittgenstein, 'Hinge' Propositions, and *On Certainty*." *Rationality and Irrationality: Proceedings of the 23rd International Wittgenstein Symposium*. Ed. Berit Brogaard and Barry Smith. Kirchberg, Austria: Austrian L. Wittgenstein Society, pp. 84–90.
- _____. 2001a. "A Puzzle about Warrant." *Philosophical Inquiry*, vol. 23, pp. 59–72.
- _____. 2001b. "Contextualism, Scepticism, and the Problem of Epistemic Descent." *Dialectica*, vol. 55, pp. 327–349.
- _____. 2001c. "Doubt Undogmatized: Pyrrhonian Scepticism, Epistemological Externalism, and the 'Metaepistemological' Challenge." *Principia: Revista Internacional de Epistemologia*, vol. 5, pp. 187–218.
- _____. 2001d. "Meta-epistemological Constraints on Anti-Sceptical Theories." *Facta Philosophica*, vol. 3, pp. 101–126.
- _____. 2001e. "Radical Scepticism, Epistemological Externalism, and 'Hinge' Propositions." *Wittgenstein-Studien*, pp. 81–105.
- _____. 2001f. "Scepticism and Dreaming." *Philosophia*, vol. 28, pp. 373–390.
- _____. 2001g. "The Opacity of Knowledge." *Essays in Philosophy, Vol. 2: The Internalism/Externalism Debate in Epistemology*. Ed. H. Benjamin Shaeffer. Humboldt, Calif.: Humboldt University Press, pp. 1–10.
- _____. 2002a. "McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments." *Synthese*, vol. 130, pp. 279–302.
- _____. 2002b. "Radical Scepticism, Epistemological Externalism, and Closure." Forthcoming in *Theoria*, vol. 69.

- _____. 2002c. "Resurrecting the Moorean Response to Scepticism." Forthcoming in *International Journal of Philosophical Studies*, vol. 10.
- _____. 2002d. "The Structure of Skeptical Arguments." *Typescript*.
- _____. 2002e. "Two Forms of Epistemological Contextualism." *Typescript*.
- Putnam, H. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- _____. 1992. *Renewing Philosophy*. Cambridge, Mass.: Harvard University Press.
- _____. 1998. "Skepticism." In *Philosophie in Synthetischer Absicht* (Synthesis in Mind). Ed. M. Stamm. Stuttgart, Germany: Klett-Cotta, pp. 239–268.
- Quine, W. V. O. 1960. *Word and Object*. Cambridge, Mass.: MIT Press.
- _____. 1969. "Epistemology Naturalised." Reprinted in *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- _____. 1981. "Reply to Stroud." *Midwest Studies in Philosophy Vol. 6*. Ed. P. A. French, T. E. Uehling, and H. K. Wettstein. Minneapolis: University of Minnesota Press, pp. 473–475.
- Rieber, S. 1998. "Skepticism and Contrastive Explanation." *Noûs*, vol. 32, pp. 189–204.
- Rorty, R. 1997. "Comments on Michael Williams's *Unnatural Doubts*." *Journal of Philosophical Research*, vol. 22, pp. 1–10.
- Rosenberg, R. J. 2000. "Scrutinising a Trade." *Philosophical Issues*, vol. 10, pp. 58–66.
- Sainsbury, R. M. 1997. "Easy Possibilities." *Philosophy and Phenomenological Research*, vol. 57, pp. 907–919.
- _____. 2000. "Warrant-Transmission, Defeaters, and Disquotation." *Philosophical Issues*, vol. 10, pp. 191–200.
- Schiffer, S. 1996. "Contextualist Solutions to Scepticism." *Proceedings of the Aristotelian Society*, vol. 96, pp. 317–333.
- Sosa, E. 1985. "Knowledge and Intellectual Virtue." *The Monist*, vol. 68, pp. 224–245.
- _____. Sosa, E. 1991. "Intellectual Virtue in Perspective." In his *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- _____. 1993. "Proper Functionalism and Virtue Epistemology." *Noûs*, vol. 27, pp. 51–65.
- _____. 1994. "Philosophical Scepticism and Epistemic Circularity." *Proceedings of the Aristotelian Society* (supplementary vol.) 68, pp. 263–290.
- _____. 1999a. "How to Defeat Opposition to Moore." *Philosophical Perspectives*, vol. 13, pp. 141–154.
- _____. 1999b. "Skepticism and the Internal/External Divide." *Epistemology*. Ed. J. Greco and E. Sosa. Oxford: Basil Blackwell, pp. 145–157.
- _____. 2000a. "Skepticism and Contextualism." *Philosophical Issues*, vol. 10, pp. 1–18.
- _____. 2000b. "Replies." *Philosophical Issues*, vol. 10, pp. 38–42.
- Stine, G. C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies*, vol. 29, pp. 249–261.
- Strawson, P. F. 1985. *Skepticism and Naturalism: Some Varieties*. London: Methuen.
- Stroll, A. 1994. *Moore and Wittgenstein on Certainty*. Oxford: Oxford University Press.
- Stroud, B. 1981. "The Significance of Naturalized Epistemology." *Midwest Studies in Philosophy Vol. 6*. Ed. P. A. French, T. E. Uehling, and H. K. Wettstein. Minneapolis: University of Minnesota Press, pp. 455–471.
- _____. 1984. *The Significance of Philosophical Scepticism*. Oxford: Oxford University Press.
- _____. 1989. "Understanding Human Knowledge in General." *Knowledge and Scepticism*. Ed. M. Clay and K. Lehrer. Boulder, Colo.: Westview.
- _____. 1994. "Scepticism, 'Externalism,' and the Goal of Epistemology." *Proceedings of the Aristotelian Society* (supplementary vol.) 68, pp. 290–307.
- _____. 1996. "Epistemological Reflection on Knowledge of the External World." *Philosophy and Phenomenological Research*, vol. 56, pp. 345–358.

- Suárez, A. G. 2000. "On Wright's Diagnosis of McKinsey's Argument." *Philosophical Issues*, vol. 10, pp. 164–171.
- Thalberg, I. 1974. "Is Justification Transmissible Through Deduction?" *Philosophical Studies*, vol. 25, pp. 357–364.
- Throop, W. 1998. "Defeating the Sceptic." *Philosophia*, vol. 27, pp. 321–336.
- Tymoczko, T., and J. Vogel. 1992. "The Exorcist's Nightmare: A Reply to Crispin Wright." *Mind*, no. 403, pp. 543–552.
- Unger, P. 1971. "A Defence of Skepticism." *Philosophical Review*, vol. 80, pp. 198–219.
- _____. 1975. *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
- _____. 1984. *Philosophical Relativity*. Oxford: Basil Blackwell.
- _____. 1986. "The Cone Model of Knowledge." *Philosophical Topics*, vol. 14, pp. 125–178.
- Vahid, H. 1995. "Deductive Closure, Scepticism and Confirmation." *Ratio* (new series), vol. 8, pp. 70–86.
- Valdés-Villanueva, L. M. 2000. "Contextualism and Levels of Scrutiny." *Philosophical Issues*, vol. 10, pp. 72–79.
- Villanueva, E. 2000. "What Has Contextualism to Do with Skepticism?" *Philosophical Issues*, vol. 10, pp. 67–71.
- Vogel, J. 1990. "Are There Counter-Examples to the Closure Principle?." *Doubting: Contemporary Perspectives on Scepticism*. Ed. G. Ross and M. D. Roth. Dordrecht, Netherlands: Kluwer Academic Publishers, pp. 20–23.
- _____. 1997. "Skepticism and Foundationalism: A Reply to Michael Williams." *Journal of Philosophical Research*, vol. 22, pp. 11–28.
- _____. 1999. "The New Relevant Alternatives Theory." *Philosophical Perspectives*, vol. 13, pp. 155–180.
- _____. 2000. "Reliabilism Levelled." *Journal of Philosophy*, vol. 97, pp. 102–124.
- Warfield, T. 1998. "A Priori Knowledge of the World: Knowing the World by Knowing our Minds." *Philosophical Studies*, vol. 60, pp. 76–90.
- Williams, M. 1988a. "Epistemological Realism and the Basis of Scepticism." *Mind*, no. 387, pp. 415–439.
- _____. 1988b. "Scepticism without Theory." *Review of Metaphysics*, vol. 41, pp. 547–588.
- _____. 1988–1989. "Scepticism and Charity." *Ratio* (new series), vols. 1–2, pp. 176–194.
- _____. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Basil Blackwell.
- _____. 1996. "Understanding Human Knowledge Philosophically." *Philosophy and Phenomenological Research*, vol. 56, pp. 359–378.
- _____. 1997. "Still Unnatural: A Reply to Vogel and Rorty." *Journal of Philosophical Research*, vol. 22, pp. 29–42.
- _____. 1999. "Fogelin's Neo-Pyrrhonism." *International Journal of Philosophical Studies*, vol. 7, pp. 141–158.
- _____. 2000. "Is Contextualism Statable?" *Philosophical Issues*, vol. 10, pp. 80–85.
- _____. 2001. "Contextualism, Externalism and Epistemic Standards." *Philosophical Studies*, vol. 103, pp. 1–23.
- Williamson, T. 2000a. "Scepticism and Evidence." *Philosophy and Phenomenological Research*, vol. 60, pp. 613–628.
- _____. 2000b. "Scepticism, Semantic Externalism and Keith's Mom." *The Southern Journal of Philosophy*, vol. 38, pp. 148–157.
- _____. 2000c. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- _____. 2001. "Comments on Michael Williams's 'Contextualism, Externalism and Epistemic Standards.'" *Philosophical Studies*, vol. 103, pp. 24–33.

- Wittgenstein, L. 1969. *On Certainty*. Ed. G. E. M. Anscombe and G. H. von Wright, tr. D. Paul and G. E. M. Anscombe. Oxford: Basil Blackwell.
- Wright, C. 1985. "Facts and Certainty." *Proceedings of the British Academy*, vol. 71, pp. 429–472.
- _____. 1991. "Scepticism and Dreaming: Imploding the Demon." *Mind*, no. 397, pp. 87–115.
- _____. 1992. *Truth and Objectivity*. Cambridge, Mass.: Harvard University Press.
- _____. 1994. "On Putnam's Proof that We Are Not Brains in a Vat." *Reading Putnam*. Ed. B. Hale and P. Clark. Oxford: Basil Blackwell, pp. 216–241.
- _____. 2000a. "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof." *Philosophical Issues*, vol. 10, pp. 140–163.
- _____. 2000b. "Replies." *Philosophical Issues*, vol. 10, pp. 201–219.
- Yalçın, U. D. 1992. "Skeptical Arguments from Underdetermination." *Philosophical Studies*, vol. 68, pp. 1–34.
- Yourgrau, P. 1983. "Knowledge and Relevant Alternatives." *Synthese*, vol. 55, pp. 55–70.

