

Why the Ascriber Contextualist Solution to the Sceptical Paradox is Unnecessary

A Thesis Submitted to the College of
Graduate Studies and Research
in Partial Fulfillment of the Requirements
for the Degree of Master of Arts
in the Department of Philosophy
University of Saskatchewan
Saskatoon

By

Ian MacDonald

PERMISSION TO USE

In presenting this thesis in partial fulfillment of the requirements for a Postgraduate degree from the University of Saskatchewan, I agree that the Libraries of this University may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department or the Dean of the College in which my thesis work was done. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to the University of Saskatchewan in any scholarly use which may be made of any material in my thesis.

Requests for permission to copy or to make other uses of materials in this thesis in whole or part should be addressed to:

Head of the Department of Philosophy
University of Saskatchewan
Saskatoon, Saskatchewan, S7N 5A5
Canada

OR

Dean
College of Graduate Studies and Research
University of Saskatchewan
107 Administration Place
Saskatoon, Saskatchewan, S7N 5A2
Canada

ABSTRACT

A widely discussed philosophical puzzle in contemporary epistemology is the so-called sceptical “paradox.” Ascriber contextualism has taken centre stage as the anti-sceptical theory that purportedly offers the best solution to the sceptical “paradox.” Ascriber contextualists Stewart Cohen (1988, 1999) and Keith DeRose (1995) advertise their anti-sceptical theory as the one that exclusively explains and solves it. This is false advertising, however. My thesis, which has been greatly influenced by the critical work of Michael Williams (1991) and Duncan Pritchard (2005), is that the generation of the sceptical “paradox” depends on whether the epistemologist is an internalist or externalist about knowledge, and that the ascriber contextualist attempt to solve the sceptical “paradox” rests on a long history of mistakes concerning internalist assumptions made by externalists Fred Dretske (1970) and Robert Nozick (1981). By applying the semantic thesis of ascriber contextualism to epistemology, ascriber contextualists seek to emend the rejection of the closure principle made by these externalists. This rejection came from these externalists mistakenly making internalist assumptions when facing sceptical hypotheses. Unfortunately, ascriber contextualists leave much unfixed, and end up inheriting and suffering from the serious mistake about internalist assumptions that had plagued the epistemologies of these externalists and now infects the ascriber contextualist “solution” to the sceptical “paradox.” With the help of hindsight to examine this history and an appreciation of how the adoption of one of these respective views about knowledge makes all the difference for whether the sceptical “paradox” arises, we come to see that the contextualist “solution” to the sceptical “paradox” is unnecessary.

ACKNOWLEDGEMENTS

I would like to sincerely thank my thesis supervisor, Dr. Eric Dayton, whose insights and suggestions have been extremely valuable and useful in writing this thesis. It was wonderful to have so many interesting and enriching philosophical discussions with Dr. Dayton. I would also like to acknowledge my two committee members, Dr. Robert Hudson and Dr. Karl Pfeifer, and my external examiner, Dr. Eric Neufeld, for all of their helpful comments, questions and criticisms. Two friends I would especially like to thank for engaging in many philosophical conversations and letting me try out some of the ideas contained in this thesis on them are Derek Postnikoff and Adam Sopuck. Writing this thesis would have been a struggle, were it not for good, supportive friends. I would also like to acknowledge the generous financial support I received from the College of Graduate Studies and Research at the University of Saskatchewan in the form of a Graduate Teaching Fellowship 2008-2009 and a University Graduate Scholarship 2009-2010.

TABLE OF CONTENTS

PERMISSION TO USE	i
ABSTRACT.....	ii
ACKNOWLEDGEMENTS	iii
TABLE OF CONTENTS.....	iv
INTRODUCTION	vi
CHAPTER 1:	
RELEVANT ALTERNATIVES, CONTEXTUALISM, AND SCEPTICISM	1
1.1 Outline.....	1
1.2 What is the Relevant Alternatives Theory of Knowledge?.....	3
1.3 The Issue of Relevance	5
1.4 Dretske’s Epistemic Operators.....	16
1.5 Stine’s Contextualism and Criticisms of Dretske	31
1.6 How to be a Fallibilist.....	38
1.7 The Sceptical Paradox.....	42
CHAPTER 2:	
CONTEXT-SENSITIVITY, SEMANTIC BLINDNESS, AND INDEXICALITY	53
2.1 Semantic Blindness and Error Theory	53
2.2 The Argument of This Chapter	56
2.3 Is the Word “Knows” an Indexical?	57
2.4 Yourgrau’s Objection.....	63
2.5 Problems with the First Account.....	66
2.6 Cohen’s Recent Account.....	70
2.7 Unger’s Invariantism	72
2.8 The Conversational Score	73
2.9 Problems with Cohen’s Recent Account	77
2.10 In the Grip of “Flatness”-Scepticism	82
2.11 Dretske’s Theory and Relationally Absolute Terms	84
CHAPTER 3:	

THE “SKEPTICAL PROBLEM,” TRACKING, AND CONTEXTUALISM.....	89
3.1 Solving the “Skeptical Problem” and the Argument from Ignorance.....	89
3.2 Nozick’s Subjunctive Conditionals Account of Knowledge	93
3.3 DeRose’s Rule of Sensitivity and the Phenomenology of Scepticism.....	99
CONCLUSION.....	105
WORKS CITED	107

INTRODUCTION

Philosophical theories are judged by how well they can solve and explain puzzles. Explanatory power is the principal theoretical virtue of philosophical theories. In contemporary epistemology, especially because of the influence of ascriber contextualism, the so-called sceptical “paradox” has been a puzzle that has also become a test for the adequacy of anti-sceptical theories, testing specifically whether these theories can solve and explain the sceptical “paradox.” The following three individually plausible but mutually inconsistent propositions are said to issue in the sceptical “paradox”: (1) I know that some ordinary empirical propositions are true; (2) If I know some ordinary empirical proposition, and I know that this ordinary empirical proposition entails the falsity of a radical sceptical hypothesis, then I can know the falsity of a radical sceptical hypothesis; (3) I cannot know the falsity of radical sceptical hypotheses. To reveal the sense of “paradox” and feel the threat of scepticism, consider that there is nothing to prevent proposition (3) from being entered into a variant of the conditional (2) to yield the following argument: If I cannot know the falsity of radical sceptical hypotheses, and that I cannot know the falsity of radical sceptical hypotheses entails that I do not know any ordinary empirical propositions, then I do not know any ordinary empirical propositions. Accordingly, this sceptical argument reveals that I do not know that proposition (1) is true, because I do not know that proposition (3) is false. With the exception of the radical sceptic, everyone finds the denial of proposition (1) to be clearly unacceptable. The puzzle, then, is to explain how these three supposedly plausible propositions can be preserved without them resulting in mutual inconsistency generally and the sceptical “paradox” specifically.

Ascriber contextualists claim to have both solved the sceptical “paradox” and provided the only viable explanation of it. Ascriber contextualism is a *semantic* thesis, which states that the *meaning* of a term can be sensitive to changes in the conversational context of the ascriber. As this semantic thesis *applies* to epistemology, the theory states that the *meaning* of the word “know” varies with changes in the conversational context of the ascriber who claims that “S knows that p” or “I know that p” (in which case the ascriber making this claim is making a self-ascription). Such variation of meaning in the word “know” is caused specifically by error-possibilities being raised in the ascriber’s conversational context, such as by conversational participants considering or mentioning radical sceptical hypotheses in that context. To take a famous example of a radical sceptical hypothesis, an ascriber may consider the possibility of

being tricked by an all-powerful evil demon at this very instant. According to ascriber contextualists, if the ascriber dwells on this radical sceptical error-possibility, then this will determine both the kind of conversational context that the ascriber is in at the time of the ascription or denial and the meaning of the word “know” in any “knowledge”-ascription or “knowledge”-denial made in that kind of conversational context.

The contextualist “solution” of the sceptical “paradox” is to claim that there are different kinds of conversational context with different standards that determine the meaning of the word “know.” Roughly, there is the quotidian context with lax standards and the sceptical context with strict standards. Consequently, propositions (1) and (2) are true and proposition (3) is false in the quotidian context, whereas proposition (3) and the variant of the conditional (2) used to reveal the sense of “paradox” are true and proposition (1) is false in the sceptical context. All of the plausible propositions can be expressed without inconsistency, provided that they are expressed in different kinds of conversational context. According to contextualists, the meaning of the word “know” varies with the context from which the argument employing it is issued. Thus, ascriber contextualism allegedly finds a “solution” by expressing all of the individually plausible propositions in different kinds of context without mutual inconsistency or generating the sceptical “paradox.”

With the exception of proposition (1), many epistemologists challenge the acceptability of these propositions that supposedly issue in the sceptical “paradox.” Epistemologists Fred Dretske (1970) and Robert Nozick (1981), for instance, reject proposition (2), which partially expresses the closure principle.¹ For these epistemologists the question is not whether proposition (2) is independently plausible, but whether it is true. Essentially, Dretske and Nozick claim that agents, in order to have knowledge of ordinary empirical propositions, do not need to know the falsity of radical sceptical hypotheses. Accepting proposition (2) enables the sceptic to launch an argument for closure-based scepticism – an argument similar to the one I used above to reveal the sense of “paradox” and feel the threat of scepticism. The anti-sceptical strategy of these epistemologists is to claim that proposition (2) is false so that agents do not need to know

¹ The principle that *knowledge is closed under known entailment* states that if an epistemic subject S knows an ordinary empirical proposition p and S knows that p entails another proposition q – or S “competently deduces,” to use Timothy Williamson’s (2000) apt term, q from p – then S knows that q or at least is *in a position* to know that q.

the falsity of radical sceptical hypotheses as *entailments* of ordinary empirical propositions in order to know ordinary empirical propositions in the first place.

Ascriber contextualists Stewart Cohen (1988, 1999) and Keith DeRose (1995) oppose Dretske's and Nozick's rejection of proposition (2), since they contend that the propositions making up the sceptical "paradox" are independently plausible and such a move to reject one of them is simply an attempt to evade the sceptical "paradox." After all, contextualists view the sceptical "paradox" as a genuine philosophical puzzle that requires both a solution and an explanation of how the "paradox" arises in the first place. Yet Dretske and Nozick motivate the rejection of proposition (2) by claiming that proposition (3) is true. Dretske and Nozick hold that proposition (3) is true, but also that proposition (2) is false, which allows them to continue to claim that proposition (1) is true.

Curiously, Cohen (1988, 1999) and DeRose (1995) have fully accepted Dretske's (1970) and Nozick's (1981) general claim that (3) epistemic agents cannot know the falsity of radical sceptical hypotheses. Upon close examination, however, Dretske's and Nozick's theoretical motivations for rejecting proposition (2) are found to be extremely problematic. Both Dretske and Nozick make epistemological commitments to *externalism* about knowledge, the view that whatever turns an agent's true belief into knowledge can remain beyond or external to that agent's cognitive awareness. The competing view of knowledge is *internalism*, the view that whatever turns an agent's true belief into knowledge must remain reflectively and cognitively accessible to that agent, in which such knowledge is from within the first-person perspective. Upon close examination we shall see that Dretske and Nozick break their externalist commitments, because the general rejection of proposition (2) comes from them unwittingly making distinctively internalist assumptions about knowledge. More specifically, Dretske's and Nozick's general claim (3) that agents do not know the falsity of radical sceptical hypotheses, rests on a mistake, which results from them confusing externalist commitments with internalist assumptions. Consequently, given that their commitments are to externalism, neither Dretske nor Nozick should have rejected proposition (2) nor should they have rejected the closure principle more generally for the reasons that they did so.

Seeing that ascriber contextualists Cohen and DeRose fully accept Dretske's and Nozick's general claim (3), which rests on a mistake, they accordingly inherit this mistake. If

they do not inherit this mistake directly, at the very least ascriber contextualists suffer from the effects of this mistake, insofar as these ascriber contextualists staunchly defend the truth of proposition (3) by pointing to Dretske's and Nozick's acceptance of it, but nonetheless fail to realize that Dretske and Nozick have made a serious mistake in so doing. My thesis, which has been greatly influenced by the critical work of Michael Williams (1991) and Duncan Pritchard (2005), is that the ascriber contextualist "solution" to the sceptical "paradox" both loses the theoretical motivation for its application in light of these mistakes and lacks the explanatory edge it purports to have over its rival anti-sceptical theories, especially when we realize that the plausibility of proposition (3) depends entirely on which view of knowledge, internalist or externalist, is taken. The consequence of taking an externalist view of knowledge is that there is, in general, no reason why proposition (3) should pose a problem, because the reasons for endorsing the truth of proposition (3) depend on making internalist assumptions about knowledge. Consequently, if an epistemologist is an externalist, then the sceptical "paradox" is never generated, since the "paradox" does not arise unless all of its propositions are accepted. The consequence of taking an internalist view of knowledge, however, is a commitment to the truth of proposition (3), and so the sceptical "paradox" remains a genuine philosophical puzzle for the internalist. Yet even internalists who develop contextualist "solutions" to the sceptical "paradox," such as Cohen (1988, 1999), have tremendous difficulty in finding harmony between internalism and contextualism. Indeed, Cohen's internalist contextualist "solution" proves to be *ad hoc* and ultimately unacceptable.

The first chapter examines the relevant alternatives theory of knowledge, including versions of this theory put forth by Fred Dretske (1970), Alvin Goldman (1976), Gail Stine (1976), and Stewart Cohen (1988). I have conducted this examination because ascriber contextualism historically developed out of, and the sceptical "paradox" emerged from, the relevant alternatives theory of knowledge. A discussion of what ascriber contextualism is occurs when I discuss Goldman's two views of relevance, because Goldman's second view of relevance brought about the development of ascriber contextualism. Any such discussion would be woefully incomplete, however, without voicing DeRose's (1992) critical response to the relevant alternatives theory of knowledge generally and Goldman's two views of relevance specifically. DeRose's critical response ultimately demands that ascriber contextualism be distinguished from the relevant alternatives theory of knowledge. A major focus of the chapter is on Dretske's

rejection of the closure principle and his poorly motivated reasons for doing so, since his mistake proves to be crucial in supporting my thesis when I turn to Cohen's presentation of, and ascriber contextualist "solution" to, the sceptical "paradox."

The second chapter challenges the purported and self-proclaimed explanatory power of ascriber contextualism as a semantic thesis about the context-sensitive meaning of the word "know," since ascriber contextualists apply this semantic thesis to epistemology to solve and explain the sceptical "paradox." The idea that the meaning of the word "know" is context-sensitive is extremely controversial. But ascriber contextualists claim that the context-sensitivity of epistemic vocabulary explains how the sceptical "paradox" arises in the first place. According to contextualists, such an explanation is a *desideratum* for a satisfying resolution of the sceptical "paradox." Specifically, ascriber contextualists claim that ascribers or sceptics are sometimes semantically blind to the context-sensitivity of epistemic vocabulary. This is a form of blindness which explains how it is that the ascriber *believes* that the "knowledge"-ascription conflicts with the sceptical hypothesis or "knowledge"-denial, while the sceptic *believes* that the "knowledge"-denial conflicts with the "knowledge"-ascription, even when the ascriber's "knowledge"-ascription is made in the quotidian context and the sceptic's "knowledge"-denial is made in the sceptical context. From a theoretical point of view, the semantic blindness that the ascriber and sceptic allegedly experience to context-sensitive epistemic vocabulary explains how this could be. According to contextualists, the fact that the ascriber's "knowledge"-ascription and the sceptic's "knowledge"-denial are made within different kinds of context makes them non-conflicting. The main contextualist device for explaining the context-sensitivity of the word "know" is to treat this word as an indexical akin to the words "I," "here," and "now." Contextualists contend that if the word "know" is an indexical, then the apparent conflict between the ascriber's "knowledge"-ascription and the sceptic's "knowledge"-denial is no more in conflict than the merely apparent conflict between the sentences "I am hungry (at t1)" and "I am not hungry (at t1)." The conflict may be merely apparent, given that the indexical "I" in each sentence could refer to a different individual with a different state of appetite from the other. Accounts of the indexicality of epistemic vocabulary, however, are strained and prove to be unacceptable. The findings of this second chapter remove much of the purported explanatory power of the ascriber contextualist "solution" to the sceptical "paradox" by undermining and diminishing that explanatory power on its own terms.

The third chapter reaffirms my main claim that the plausibility of endorsing proposition (3) depends on whether one adopts an internalist or externalist view of knowledge. This will reveal that the sceptical “paradox” is only puzzling to the internalist. A similar story to the one told in the first chapter, although with different characters, is told in this final chapter. Instead of Dretske and Cohen at its centre, however, the story revolves around Nozick and DeRose. I investigate Nozick’s subjunctive conditionals account of knowledge, which proposes a philosophical explanation of radical scepticism. While Nozick is a self-proclaimed externalist, he ends up both breaking his externalist commitments and rejecting the closure principle, because he mistakenly makes an internalist assumption about knowledge in the face of radical sceptical hypotheses. This is an egregious mistake for any externalist to make. Curiously, DeRose’s epistemological orientation is externalist too, but yet he fails to see Nozick’s mistake. The fact that DeRose thinks that Nozick’s explanation of radical scepticism is the right approach but also believes that Nozick should not have rejected proposition (2) or the closure principle, motivates DeRose to contextualize Nozick’s subjunctive conditionals account of knowledge. DeRose does so in order to explain the plausibility of proposition (3) and “solve” what he calls the “Skeptical Problem,” which he admits only differs from Cohen’s sceptical “paradox” with respect to its name and the details. Consequently, he cannot reject proposition (2), as Nozick did, since DeRose needs to find an adequate “solution” to the “Skeptical Problem.” An adequate “solution” requires that all of the propositions, including proposition (2), remain intact. Yet if DeRose is an externalist about knowledge and Nozick makes a serious mistake by making an internalist assumption about knowledge in the face of radical sceptical hypotheses, then DeRose’s motivation for applying ascriber contextualism to handle radical scepticism is as poorly motivated as Nozick’s rejection of the closure principle.

While explanatory power is the principal theoretical virtue of philosophical theories, another important theoretical virtue of such theories is economy. Ultimately, ascriber contextualists posit contexts unnecessarily. The “solution” to the sceptical “paradox” does not require ascriber contextualism, because the “paradox” does not arise if an epistemologist is an externalist. Since proposition (3) will be implausible and unacceptable for the externalist, this will thereby render the application of ascriber contextualism as simply unnecessary. The sceptical “paradox,” however, will continue to haunt the internalist – who in the end may have to appeal to the supposedly “almighty” benefits of ascriber contextualism to “exorcize” the radical

sceptical possibility of an all-powerful Cartesian evil demon. The contextualist “solution” tends to be the layer which obfuscates the point that epistemological differences, in the form of externalism and internalism, make all the difference. I think that I have found a razor-sharp reason to shave off the contextualist layer, as I will now endeavour to show.

CHAPTER 1

RELEVANT ALTERNATIVES, CONTEXTUALISM, AND SCEPTICISM

1.1 Outline

When discussing so many interconnected and complex ideas there is a danger of burying the lead. To avoid this danger, I want to immediately provide the following headline: What is called the sceptical “paradox” is not a paradox and the ascriber contextualist attempt to solve it is unnecessary. What is called the sceptical “paradox” is not a paradox, because the acceptability of the proposition that subjects do not know the falsity of sceptical alternatives crucially depends on which view of knowledge, externalism or internalism, is taken. That proposition is only acceptable if one takes an *internalist* view of knowledge. This renders the ascriber contextualist attempt to solve the sceptical “paradox” unnecessary. There is no need to apply a semantic thesis to try to solve the “paradox,” since the “paradox” dissolves if the epistemologist takes an externalist view of knowledge. Let me try to briefly explain why this is so in this section and then go on to make the arguments for my thesis.

Recall the trinity of propositions that are said to make up the sceptical “paradox.” One proposition of the sceptical “paradox” is that epistemic subjects know many ordinary empirical propositions, such as that they know what time it is, what they had for breakfast, or which baseball team won the world series last year. The relevant alternatives theorist holds that only *relevant* alternatives to these ordinary empirical propositions need to be ruled out in order to know them. The relevant alternatives theorist contends that sceptical alternatives, such as that I may have been tricked by an all-powerful evil demon into thinking that the Boston Red Sox won it, very often remain *irrelevant* alternatives to know ordinary empirical propositions. But sometimes the sceptical alternatives are less radical and harder to say whether they are relevant, such as when, for instance, one hears a familiar bird call (say, the call of what sounds like a sparrow) but cannot rule out the alternative that the call may have come from a mocking bird – a natural imitator of other birds. In this case, the question of how to determine whether this

alternative is relevant or irrelevant becomes a very difficult question to answer precisely. The difficulty of the question comes from trying to figure out what kinds of factors or conditions affect relevance. For instance, in the mocking bird case, does the relevance of this alternative depend strictly on geographical location? Does it depend on that subject's ability to discriminate the call of a mocking bird from that of other birds? Does it depend on a combination of these kinds of factors or conditions? Could it depend on the kind of conversational context of the speaker raising this possibility, as the ascriber contextualist argues? Sections **1.2**, **1.3**, and **1.6** discuss the relevant alternatives theory of knowledge and the issue of relevance. Many ascriber contextualists who are relevant alternatives theorists, such as Stine (section **1.5**) and Cohen (section **1.6**), have held that ascriber contextualism makes it much easier to say when an alternative is relevant and when it is irrelevant. Essentially, being able to settle the issue of relevance would make it much easier to claim that subjects know many ordinary empirical propositions and to say how that is so.

Another proposition in the trinity is that the closure principle holds. Section **1.4** discusses Dretske's rejection of the closure principle. Section **1.5** examines Stine's criticism of Dretske's rejection and her attempt to preserve the closure principle by endorsing ascriber contextualism. Section **1.6** discusses Cohen's attempt to preserve the closure principle by also endorsing ascriber contextualism. Ascriber contextualists, such as Stine and Cohen, seek to preserve the closure principle and think that Dretske *mistakenly* rejects it. Stine charges Dretske with committing some error akin to equivocation when he argues for closure failure, whereas Cohen thinks that it is simply implausible to reject the closure principle. I also argue that Dretske mistakenly rejects the closure principle. But I identify Dretske's mistake as coming from an internalist assumption he makes when he claims that subjects do not know the falsity of sceptical alternatives. Dretske is an externalist. So this is an isolated internalist assumption that he makes. But it is a crucial one, since it causes him to reject the closure principle. I contend that if he gets rid of this assumption and continues to be an externalist, then sceptical alternatives will remain non-threatening and the closure principle will indeed hold. I make these claims in section **1.4**.

The remaining proposition in the so-called sceptical "paradox" is that epistemic subjects do not know the falsity of sceptical alternatives. Relevant alternatives theorists, many of whom are ascriber contextualists, are responding to *closure*-based scepticism. Closure-based scepticism can be radical or moderate in scope. Relevant alternatives theorists hold that only a subset of the

alternatives, the *relevant* alternatives, are the ones that subjects need to rule out. Yet finding a good reason to respond to scepticism has been a persistent difficulty for this theory of knowledge. By rejecting the closure principle, Dretske removes the very *basis* of this form of scepticism. This is a response to scepticism. But his argument against the sceptic will not work, because he mistakenly rejects the closure principle. To remove the *apparent* threat of scepticism, Stine and Cohen endorse ascriber contextualism. Yet ascriber contextualists are unaware that the threat of scepticism is *merely apparent* if the epistemologist is an externalist about knowledge. Generally, whether scepticism is threatening or not will depend on fundamental epistemological commitments involving internalism and externalism.

With this last point in mind, I finally arrive at the sceptical “paradox” in section 1.7. I challenge the general acceptability of the proposition that subjects do not know the falsity of sceptical alternatives. If that proposition is not generally acceptable, then the sceptical “paradox” dissolves, rendering the ascriber contextualist “solution” to it unnecessary. Unfortunately, ascriber contextualists who are externalists do not seem to realize this and continue to unnecessarily propose ascriber contextualist “solutions” to what they consider to be the sceptical “paradox.” However, internalists may continue to be riddled by this situation, because the internalist will claim that subjects know many ordinary empirical propositions, the closure principle holds, and subjects do not know the falsity of sceptical alternatives. There is the further question of whether the ascriber contextualist who is an internalist could find a “solution” to what that theorist considers to be the sceptical “paradox.” Section 1.7 explores this question, since Cohen is such an ascriber contextualist. Yet it appears that Cohen’s “solution” is *ad hoc*. Ultimately, however, ascriber contextualists posit contexts unnecessarily, since the sceptical “paradox” dissolves at the epistemological level of analysis and need not require the application of a semantic thesis about the meaning of the word “knows.”

1.2 What is the Relevant Alternatives Theory of Knowledge?

The relevant alternatives theory of knowledge states that in order for an epistemic subject S to know a particular ordinary empirical proposition p, S must have a true belief that p *and* be able to rule out, on the basis of S’s evidence or perceptual abilities, a subset of the alternatives to p; namely, S must be able to rule out the set of *relevant* alternatives to p. For instance, if S is said to know what time it is on the basis of a good working watch, then it would seem the alternative

that S's watch has stopped and tells the wrong time would remain an *irrelevant* alternative to S's knowing what time it is. If S's watch has been malfunctioning lately, however, then the alternative that S's watch has stopped and tells the wrong time may be a *relevant* alternative that S needs to rule out to know what time it is (again, in this case, using the watch as the basis to know what time it is). Which factors precisely make an alternative *relevant* is a vexed question. There is, however, wide agreement among Relevant Alternatives (RA) theorists that radical sceptical alternatives often remain irrelevant to whether S knows ordinary empirical propositions, such as to whether S knows what time it is or what S had for breakfast this morning.² For example, S does not need to know the falsity of the radical sceptical alternative that the world may have sprung into being five minutes ago, complete with pseudo-memories, in order to know that S had oatmeal this morning.³ If S sometimes has cornflakes or toast instead, then ruling out these (relevant) alternatives would be necessary for S to know what S had for breakfast. Accordingly, the relevant alternatives theory of knowledge often is viewed as a corrective to scepticism. As Patrick Rysiew (2006) remarks, "Just about every RA theorist takes the avoidance of scepticism as a *desideratum* for a satisfactory theory of knowledge."⁴ Yet, to truly avoid scepticism in the most desired way, the RA theorist must confront the sceptic on the issue of whether knowledge is closed under known entailment.

The principle of *knowledge closed under known entailment* states that if an epistemic subject S knows an ordinary empirical proposition p and S knows that p entails another proposition q – or S "competently deduces," to use Timothy Williamson's (2000) apt term, q from p – then S knows that q or at least is *in a position* to know that q.⁵ Since scepticism is being discussed, for the present purposes, the closure principle could be equally presented in the following modified form: if S knows that p and S knows that p entails the negation \sim sa of a

² The paradigmatic example of an ordinary empirical proposition in the literature is the proposition "I know that I have hands." No doubt this is due to Moore (1959). I do not consider this to be a very good example, however, because the question "Do you know that you have hands?" would only arise when examining what one knows in the face of a radical sceptical threat. Unlike that example, the questions "Do you know what time it is?" or "Do you know what you had for breakfast this morning?" are less out of place and more quotidian, which is why I have chosen the answers to them as lead examples of ordinary empirical propositions.

³ This is Bertrand Russell's famous five minute sceptical hypothesis, which he thought was "logically tenable" in accounting for memories, but like all other radical sceptical alternatives, Russell found this hypothesis to be totally "uninteresting" and not to be taken "seriously" (Russell 1921, 132).

⁴ Rysiew 2006, 266 [original emphasis].

⁵ Williamson 2000, 117. The qualification of "is in a position to know" acknowledges that S *need not* know all the entailments of S's belief or body of beliefs to know that P. See: Hawthorne 2000, 119.

particular sceptical alternative – and S competently deduces \sim sa from p – then S knows that \sim sa or is at least in a position to know that \sim sa. The sceptic exploits the closure principle to make the argument for closure-based scepticism, an argument which states that if the closure principle holds, then S must know that \sim sa to know some ordinary empirical proposition p; S does not know that \sim sa, or at least does not seem to know that \sim sa on the basis of S's evidence for the ordinary empirical proposition p; therefore, S does not know that p.⁶ Without reasserting the relevant alternatives theory in the face of this sceptical threat – a theory which no sceptic finds convincing, and at this stage some non-sceptics may find unconvincing too – how can the RA theorist respond to the sceptic's argument for closure-based scepticism? Before I turn to the different answers that RA theorists have given to this question, I want to examine the major issue of relevance, which RA theorists have often had difficulty in addressing, especially to the satisfaction of their critics.

1.3 The Issue of Relevance

In the last section I briefly noted that the question of what makes an alternative relevant or irrelevant is a vexed question. Controversial as it may be, it remains necessary for the RA theorist to elucidate this central notion of relevance if the relevant alternatives theory of knowledge is to remain a viable epistemology. Yet only saying so much understates the problem. Critics of the relevant alternatives theory are not optimistically awaiting an account of relevance, but rather have found all of the accounts or specific elaborations offered so far to be extremely wanting. Some of these critics even consider the RA theorist's task a hopeless one.⁷ Although he does not consider it hopeless, Rysiew (2006) observes the unhappiness that critics of the relevant alternatives theory have felt over this issue of relevance:

It's not the *mere* fact that the RA theorist needs an account of... 'relevance' that has tended to lead people to regard the RA approach with suspicion. In itself, this simply means that the RA theorist has some further work to do; and what theorist doesn't? No; the principal source of scepticism regarding the ability of the RA theorist to come up with

⁶ Duncan Pritchard (2005) calls this the argument the *closure-based scepticism argument* (Pritchard 2005, 27). I shall likewise refer to it in this way. Such scepticism concerns *contingent* propositions known *a posteriori*, and not necessary propositions. The *scope* of such scepticism can be *local* (e.g., there is no knowledge of other minds, or God, or the external world or some particular subject for some particular reason) and moderate or can be *global* (e.g., there is no knowledge of any *contingent* propositions) and radical. See: Dancy 1985, 8.

⁷ Sosa (1986), for instance, contends that the RA theory remains "obscure," because it cannot elucidate the notion of "relevance."

a complete and satisfactory account of knowing stems, rather, from an unhappiness with *the specific elaborations* of the core RA claim that various theorists have offered.⁸

Even if there is considerable disagreement over specific elaborations of the notion of relevance, especially with regard to which factors determine relevance, there is unanimity among RA theorists that the relevant alternatives to ordinary empirical propositions exist within a flexible and variable framework. Dretske (1981) is particularly clear about this kind of variation: “...although knowledge requires the evidential elimination of all relevant alternatives (to what is known), there is a shifting, variable set of relevant alternatives.”⁹ Not all RA theorists, however, agree with Dretske’s (1981) requirement that knowledge requires *evidential* elimination; Stine (1976), for instance, adamantly rejects this as a requirement for knowledge.¹⁰ For expository purposes, a more neutral term, such as “epistemic position,” may be preferred to the contested term “evidential position.” Nevertheless, there are not even murmurings of dissent among RA theorists against this kind of variation from being applied to the framework or set of relevant alternatives. As a theoretical consequence of this kind of variation, an epistemic subject S may be in a strong epistemic position with respect to knowing that p on one occasion within one framework, but may not be in that same epistemic position on another occasion within a different framework, because of the introduction of relevant alternatives into that framework. What causes variation in the set of relevant alternatives is a matter of great controversy, especially between ascriber contextualist RA theorists and their older non-contextualist relatives.

To gain an appreciation of this controversy, it will be helpful to come to it by way of Goldman’s (1976) use of the famous barn façade example.¹¹ In this example, Goldman contrasts two cases to analyze the concept of non-inferential perceptual knowledge, but he also considers two different interpretations or views of the notion of relevance. The first case is a set-up to the second case, which is of true philosophical interest. Goldman’s two views of relevance will be the main focus of this section. After introducing Goldman’s two views of relevance, I will discuss DeRose’s (1992) critical response to Goldman, since DeRose’s response provides a clear way of understanding what ascriber contextualism is and what it is not.

⁸ Rysiew 2006, 259 [original emphasis].

⁹ Dretske 1981, 370.

¹⁰ Stine 1976, 257.

¹¹ Goldman (1976) made this example famous, but he credits it to Carl Ginet.

Consider Goldman's first case of perceptual knowledge:

Henry is driving in the country-side with his son. For the boy's edification Henry identifies various objects on the landscape as they come into view. "That's a cow," says Henry, "That's a tractor," "That's a silo," "That's a barn," etc. Henry has no doubt about the identity of these objects; in particular, he has no doubt that the last-mentioned object is a barn and indeed it is. Each of the identified objects has features characteristic of its type. Moreover, each object is fully in view, Henry has excellent eyesight, and he has enough time to look at them reasonably carefully, since there is little traffic to distract him.¹²

According to Goldman, except in perhaps our more philosophical moments, we have a strong inclination to credit Henry with knowledge in this first case. This should be of no surprise since Goldman constructs the case so that the conditions are such that nothing funny is happening. Goldman goes on to contrast this first case with the following second case, which is of philosophical interest:

Suppose we are told that, unknown to Henry, the district he has just entered is full of papier-mâché facsimiles of barns. These facsimiles look from the road exactly like barns, but are really just façades, without back walls or interiors, quite incapable of being used as barns. They are so cleverly constructed that travelers invariably mistake them for barns. Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the object on that site were a facsimile, Henry would mistake it for a barn. Given this new information, we would be strongly inclined to withdraw the claim that Henry *knows* the object is a barn.¹³

In this second case, Henry *just happens* to be looking at a real barn, which prevents us from crediting him with knowledge; or as Goldman points out, upon our receiving this new information about fake barns, we tend to retract our previous knowledge-claim about Henry issued from our judgment of the first case. In his own analysis of perceptual knowledge and treatment of the second case, Goldman provides a counterfactual, subjunctive knowledge-test, in which Henry would know that he is looking at a real barn only *if* Henry *were* able to discriminate or distinguish the real barn from the fake types, which are, possibly – depending on one's interpretation of relevance – relevant alternatives in the given state of affairs.¹⁴ On Goldman's analysis, this is the counterfactual test that Henry needs to pass so that he would cease to mistakenly believe that the barn was a barn *if* it were a facade. After passing the test, then, it

¹² Goldman 1976, 772.

¹³ Goldman 1976, 773 [original emphasis].

¹⁴ Goldman 1976, 774.

would be no accident if Henry were looking at a real barn, even in a land rife with fakes.¹⁵ For the purposes of this discussion, however, Goldman's analysis of perceptual knowledge is not as important as the two competing views of relevance he proposes for how to specify the set of relevant alternatives in the second case.

In Goldman's presentation of the two views of relevance, the main issue is what bearing a commitment to one of these views has on the semantic content of the word "know." According to Goldman's first view of relevance, "the semantic content of 'know' contains (implicit) rules that map any putative knower's circumstances into a set of relevant alternatives."¹⁶ As I understand Goldman's first view, if Henry is credited with knowing that he is looking at a genuine barn, then the semantic content of the word "know," in the knowledge-claim "Henry *knows* that he is looking at a (genuine) barn," *means* that Henry has ruled out all of the fake barns *as* relevant alternatives in his environmental circumstances, circumstances which are the "terrain" on which the semantic content maps. In contrast, Goldman's second view of relevance denies that the word "know" has a determinate semantic content or mapping-relation (semantic content-to-environmental circumstances). This second view instead positively states that "there are regularities governing the alternative hypotheses a speaker (i.e., an attributer or denier of knowledge) thinks of, and deems relevant."¹⁷ In elaborating on the second view, Goldman writes: "It is not only the circumstances of the putative knower's situation, however, that influence the choice of alternatives. The speaker's own linguistic and psychological context are also important."¹⁸ For example, imagine two philosophy students having a conversation about the radical sceptical possibility of Henry being tricked by an omnipotent evil demon. The students have just come out an epistemology class in which Descartes's first meditation was discussed. The students have radical sceptical alternatives on their minds. One of the students claims that "Henry does not know that he is looking at a barn, because he does not know that he

¹⁵ At this point you may be wondering why Goldman did not just include a condition of *non-accidentality* to the knowledge-conditions. The reason is that Goldman did *not* think that the notion of non-accidentality was itself explicable. The notion of "non-accidentality" was put forth by Unger (1968), but Goldman held that "The notion of 'non-accidentality' itself is not very satisfying, however, for the notion of 'non-accidentality' itself needs explication"(Goldman 1976, 773). Goldman, however, without directly explicating it, achieves this non-accidentality through his subjunctive, counterfactual knowledge-test, since it prevents Henry from *just happening* to be looking at a real barn.

¹⁶ Goldman 1976, 776.

¹⁷ Goldman 1976, 776.

¹⁸ Goldman 1976, 776.

is not being tricked by an omnipotent evil demon.” According to Goldman’s second view of relevance, this radical sceptical alternative may be a *relevant* alternative, especially considering the student’s linguistic and psychological context. In Goldman’s characterization of the second view of relevance, these are the seeds which would eventually germinate into ascriber contextualism – although, depending on who you talk to, may turn into the weeds for later ascriber contextualists to get rid of.¹⁹

Although Goldman indicates some favouritism towards the second view of relevance, partly because the knower’s circumstances are so hard to map (There is what I am calling the problem of cartography: Where are the boundaries to be drawn?), Goldman remains officially non-committal about which view is correct.²⁰ Goldman’s neutrality comes from him being more interested in providing an analysis of non-inferential perceptual knowledge, which does not turn on which view of relevance is adopted. So long as there are ways of explicating the qualifier ‘relevant,’ the “scope of knowledge could be substantial,” according to Goldman; otherwise, “If knowledge required the elimination of all logically possible alternatives, there would be no knowledge (at least of contingent truths).”²¹ Goldman finds two ways of explicating the qualifier ‘relevant,’ and so moves on to analyzing perceptual knowledge. I turn next to the second view of relevance and its storied, but controversial, relationship to ascriber contextualism. DeRose thinks we must take account of how the speaker’s linguistic and psychological context affects the *semantic content* of the word “know.” Yet, against Goldman, DeRose contends that the ascriber contextualist can take account of this *without* invoking the second view of “relevance” or the notion of “relevance” altogether. According to DeRose, ascriber contextualism is exclusively about how the ascriber’s context affects the semantics of the word “know,” and has nothing to do with the issue of “relevance.” Accordingly, the ascriber contextualist *need not* be a relevant

¹⁹ As we shall see, DeRose (1992) plays the role of gardener by getting rid of the weeds of the relevant alternatives theory and discovering the roots of ascriber contextualism.

²⁰ Here is what Goldman says about this problem of mapping the putative knower’s circumstances: “The problem, of course, is to specify when an alternative is ‘idle’ [‘irrelevant’] and when it is ‘serious’ (‘relevant’). Consider Henry once again. Should we say that the possibility of a facsimile before him is a serious or relevant possibility if there are no facsimiles in Henry’s district, but only in Sweden? Or if a single such facsimile once existed in Sweden, but none exist now?” (Goldman 1976, 775). The notion of “relevance” is vague and fuzzy. It is thus a problem to answer this question – of where the boundaries are to be drawn – with precision. Because of this problem, Goldman favours the second view of relevance.

²¹ Goldman 1976, 775.

alternatives theorist. On this particular matter, I think that DeRose is correct, although I do *not* endorse ascriber contextualism.

Goldman's second view of relevance is sometimes seen as an early expression of what has come to be known as *ascriber* or *attributor contextualism*, which in the words of prominent ascriber contextualist Keith DeRose (1992) states that "...the truth conditions of sentences of the form 'S knows that p' or 'S does not know that p' vary in certain ways according to the context in which the sentences are uttered."²² There is considerable, and not altogether tolerable, vagueness in both DeRose's definition of ascriber contextualism and his characterization of how the truth-conditions²³ for "knowledge"-ascriptions, to use his clause, "vary in *certain ways*."²⁴ But vagueness aside, by "the context in which the sentences are uttered," with which the truth-conditions accord, DeRose means the *conversational* context of the ascriber or attributor of the "knowledge"-ascription; that is to say, the ascriber who is the speaker who claims, for instance, that "Henry *knows* that he is looking at a (genuine) barn."²⁵ DeRose's definition reveals that he takes the truth-conditions to be what determine the proposition expressed by the sentences "S knows that p" or "S does not know that p," and that these truth-conditions *vary* with changes in the ascriber's conversational context. According to DeRose, this variation means that the *meaning* of a proposition can be affected by what he calls *contextual* or *attributor factors*, such as the mentioning or considering of an error-possibility in a conversational context, factors which factor into the truth-conditions of "knowledge"-ascriptions.²⁶

The rival position to contextualism, which denies that attributor or contextual factors play a role in determining the meaning of knowledge-claims, is known as *invariantism*, a position famously credited to Unger (1975) and further developed in Unger's (1984) later work. In his characterization of invariantism, DeRose writes: "According to the invariantist, such features of

²² DeRose 1992, 914. Although ascriber contextualism has traditionally been associated with the RA theory, DeRose argues that the two are independent of one another and should be seen as distinct. One need not be a RA theorist to be a contextualist (DeRose 1992, 918-923).

²³ Unfortunately, contextualists tend to use the terms "truth-conditions" and "semantic standards" interchangeably.

²⁴ The scare quotes around the word "knowledge," in "knowledge"-ascriptions, are used to indicate that contextualism is a meta-linguistic thesis about the meaning of the word "knowledge," which goes beyond the object-level of propositional knowledge. Other authors, such as Bach 2005b, have used similar devices to signal to the reader that contextualism concerns the context-sensitive meaning of the word "knowledge." See: Bach 2005b, 59.

²⁵ For contextualists, the words "attributor" and "ascriber" are synonyms.

²⁶ DeRose 1992, 915. See: Unger 1975, 1984. In his later work, Unger (1984) does not argue for invariantism. In this later work, he neither supports invariantism nor contextualism; rather, he holds the position of *semantic relativity*, claiming that there are equally good arguments on both sides.

an utterance of a knowledge attribution do not affect how good an epistemic position the putative knower must be in for the attribution to be true.”²⁷

It may *appear* that DeRose is simply rehashing Goldman’s two views of relevance, since DeRose’s contextualism claims that attributor factors factor into the truth-conditions of “knowledge”- attributions, which appears to be similar to Goldman’s second view about the role that the putative knower’s psychological and linguistic context plays in affecting the semantic content of “knowledge”-claims, whereas, invariantism denies this, which makes invariantism much closer to the first view of relevance.²⁸ Moreover, the RA theorist who holds the first view of relevance would contend that *subject factors* – to borrow DeRose’s term, which he uses in contrast to attributor or contextual factors – such as, for example, the epistemic subject’s environmental circumstances, indeed determine the semantic content of knowledge-claims.²⁹ But before making this rough comparison and hasty generalization, one needs to appreciate a twist in this story.

The twist is that DeRose distances his ascriber contextualism from the relevant alternatives theory of knowledge, because he contends that one can be an ascriber contextualist without having to be a relevant alternatives theorist. For DeRose the “birthmark” of the ascriber contextualist appears when that theorist accepts a contextualist approach to semantics. As this approach is applied to epistemology, true colours are revealed when the theorist accepts that attributor factors can affect the truth-conditions of “knowledge”-claims. Consequently, DeRose argues that the notion of “relevance” has nothing to do with ascriber contextualism *per se*, and that ascriber contextualism can remain independent of the relevant alternatives theory of knowledge.

DeRose’s argument comes in the form of a challenge to Goldman’s assumption that the first view of relevance affects the semantic content of knowledge-claims. DeRose widens the scope of his challenge by challenging Goldman’s fundamental assumption that “relevance” has something important to do with semantic content. To begin with, DeRose claims that Goldman too quickly draws the distinction between these two views of “relevance,” because Goldman

²⁷ DeRose 1992, 915.

²⁸ To be clear, the words “attributor” and “contextual” in the phrase “attributor or contextual factors” is simply a difference of words. They function as synonyms in that phrase.

²⁹ DeRose 1992, 918.

overlooks the reasons which motivate drawing the distinction: “Although Goldman draws the distinction between what I am calling subject factors and attributor factors, he does not explain the importance of the distinction.”³⁰ Crucially, DeRose narrows the point of his challenge by directly challenging Goldman’s assumption that the first view of relevance or the set of subject factors associated with that view can actually affect the semantic content of knowledge-claims:

There seems to be a fairly straightforward and important sense in which one *does* mean something different if the range of relevant alternatives has been changed by attributor factors but does *not* mean something different if the range of relevant alternatives has been changed by subject factors.³¹

To appreciate DeRose’s point, consider what Stine (1976) claims in discussing Dretske’s (1970) zebra example, in which we are to imagine our hero John looking at an animal that looks like a zebra at the public zoo:

In Dretske’s zoo example, the animal’s being a mule painted to look like a zebra is not a relevant alternative. So what one *means* when one says that John knows the animal is a zebra, is that he knows that it is a zebra, as opposed to a gazelle, an antelope, or other animals one would normally expect to find in a zoo. If, however, being a mule painted to look like a zebra became a relevant alternative, then one would literally *mean* something different in saying that John knows that the animal is a zebra from what one meant originally and that something else may well be false.³²

With DeRose’s direct challenge to Goldman in mind, consider the following suggestive series of questions: What caused the painted mule possibility to become a relevant alternative in Stine’s description? Was it caused by a *subject factor* – a subject factor in the sense of a factor in the epistemic *subject’s* environment? For instance, was it caused by the introduction of painted mules into the zebra pen during John’s visit to the zoo, which on Goldman’s first view of relevance would seem to make the cleverly painted mules in the pen relevant alternatives, and which would seem to prevent John from gaining knowledge in these circumstances if he were unaware of such an introduction? Or was it an *attributor factor*, such as the painted mule error-possibility being mentioned or considered by the ascriber or denier in a conversational context? Unfortunately, Stine does not say which kind of factor caused an alteration in the meaning of what was said.

³⁰ DeRose 1992, 919.

³¹ DeRose 1992, 920.

³² Stine 1976, 255 [my emphasis].

DeRose argues, however, that changes in John's environmental circumstances or in the set of subject factors would not make the meaning of the utterance "John knows that there is a zebra" literally mean something different. This is especially apparent if the ascriber were ignorant of changes in John's environmental circumstances. In that case, without the ascriber being aware of these changes in John's circumstances, DeRose contends that that ascriber would not literally mean something different in ascribing "knowledge" by saying that "John knows that he is looking at a zebra."

In stark contrast, DeRose contends that attributor factors can indeed affect the semantic content of "knowledge"-claims. To see this, imagine an ascriber who at first claims that "John knows that there is a zebra" (as opposed to a gazelle, an antelope, or any other zebra-like animals at the zoo) in a conversation about specimens at the public zoo. Now consider that when someone else in the conversation mentions to that ascriber that there might be cleverly painted mules wandering about in the zebra pen, this error-possibility changes the set of attributor factors. According to DeRose, if the ascriber were to repeat the claim that "John knows that there is a zebra" in the face of this error-possibility, then this claim would *now* carry the meaning that John knows that there is a zebra as opposed to a cleverly painted mule in the zebra pen. In DeRose's view, this latter claim supposedly *means* that John is in a good enough epistemic position to rule out the previously mentioned error-possibility. Essentially, DeRose argues that attributor factors – *not* subject factors – cause variation in the semantic content of "knowledge"-ascriptions.

From this DeRose contends that these differences in meaning between the two claims made in different conversational contexts – one context in which the error-possibility was absent and another context in which it was mentioned and became salient – cannot be accommodated or explained by Goldman's first view of relevance. DeRose more boldly rejects the relevant alternatives theory of knowledge altogether, however, because he thinks that so long as one espouses a contextualist approach to semantics, then this espousal alone wedds one to ascriber contextualism.³³ That is to say, a contextualist can drop talk of relevance altogether in favour of simple contextualist talk in terms of what attributor factors affect the truth-conditions – a terminology which does everything that the second view of relevance does without the additional

³³ In the next chapter I discuss issues of semantics as they relate to contextualism.

epistemic baggage. The important variation for DeRose, then, is not the variability within the set of relevant alternatives, but how the truth-conditions of “knowledge”-ascriptions co-vary with attributor factors. On my reading of DeRose, the separation is clear when he claims:

RA’s *basic idea* is not about *contextual variations of meanings*... It is, then going *beyond* this basic idea that RA theorists have, by my lights, gone wrong by tying the meaning of a given attribution too closely to what the range of relevant alternatives is.³⁴

Throughout this discussion I have tried to point out many of the loose ends that have tended to unravel the tie between the relevant alternatives theory of knowledge and ascriber contextualism, especially how the attributor factors alone are causing variation in the semantic content of the word “know.”³⁵

This concludes my examination of the issue of relevance and how ascriber contextualism historically developed out of the relevant alternatives theory of knowledge. DeRose considers this development to be a historical accident, since he contends that the ascriber contextualist need not be a relevant alternatives theorist. In DeRose’s view, the two positions can remain independent of one another. Moreover, a commitment to a contextualist semantics alone makes one an ascriber contextualist. This is because attributor factors – and *not* subject factors – are causing variation in the meaning of “knowledge”-claims. According to DeRose, there is no need for the ascriber contextualist to invoke the notion of relevance in order for ascriber contextualism to be a coherent position. On this matter, I agree with DeRose.

The issue of relevance will return when I examine Cohen’s ascriber contextualist version of the relevant alternatives theory of knowledge in section 1.6.³⁶ Cohen does not think that there is a way of generally specifying the relevance of alternatives, and so holds that the relevance of alternatives depends on the kind of conversational context in effect at the time of the “knowledge”-ascription or “knowledge”-denial. In section 1.5 Stine’s discussion of different kinds of conversational context will be useful to get a sense of why Cohen appeals to different kinds of conversational context to specify the relevance of alternatives.

³⁴ DeRose 1992, 923 Note 20 [original emphasis].

³⁵ To say that DeRose would take Goldman’s second view of relevance, or is simply rehashing Goldman’s distinction, would be not only misleading, it would be inaccurate.

³⁶ Note that Cohen could not have been aware of DeRose’s (1992) criticisms of Goldman (1976), because DeRose’s criticisms came after Cohen. Yet also note that DeRose’s claim is that ascriber contextualists *need not* be relevant alternatives theorists – *not* that they cannot be both. The two positions are not mutually exclusive.

In connection with these discussions, a major problem faced by relevant alternatives theorists is to respond to scepticism. The form of scepticism that they are responding to is *closure*-based scepticism, which can be either radical or moderate in scope. The sceptic will argue that the subject must *first* rule out a sceptical alternative (a sceptical alternative which is *entailed* by an ordinary empirical proposition) to know an ordinary empirical proposition. Responding to this form of scepticism is a task taken up by Dretske, Stine, and Cohen. It is important to note that both Stine and Cohen find *closure*-based scepticism threatening, and aim to preserve the closure principle and mitigate that threat by using ascriber contextualism. Seeing that Dretske is *not* an ascriber contextualist, he does not claim that the closure principle holds relative to a kind of conversational context; rather, he rejects the closure principle. Dretske claims that the environmental conditions play a major role in determining whether alternatives are relevant or remain irrelevant to what a subject is said to know.

In the next section, **1.4**, I examine the status of the closure principle or what Dretske calls, to be historically accurate, *the thesis of penetrability*. Dretske rejects the thesis of penetrability, because he contends that epistemic operators are only semi-penetrating and that subjects do not know the falsity of sceptical alternatives. According to Dretske, because the closure principle fails to hold, *closure*-based scepticism is not threatening. Consequently, subjects only need to rule out the *relevant* alternatives to know ordinary empirical propositions. But I challenge Dretske's motivation for rejecting the closure principle. I argue that his rejection of the closure principle is poorly motivated, seeing that it rests on a dubious *internalist* assumption.

The following section will show that there is no need to reject the closure principle. It will also reveal that sceptical alternatives are not threatening if an epistemologist continues to be an externalist. Much of the work in section **1.4** will figure prominently in section **1.7**, in which I challenge the general acceptability of the proposition that subjects do not know the falsity of sceptical alternatives. This is the second proposition in the so-called sceptical "paradox," which depends for its acceptability on an *internalist* view of knowledge. Without this proposition being generally acceptable, the so-called sceptical "paradox" does not arise. Any supposed "solution" to the sceptical "paradox" will thereby be unnecessary, as I will show in section **1.7**. But first I need to discuss Dretske.

1.4 Dretske's Epistemic Operators

According to Dretske, a sentence or statement can be affixed by what he generally calls a sentential operator.³⁷ For instance, if we affix 'it is true that' to the sentence "Albert lost the election," then the affixed 'it is true that' is the operator which penetrates to the sentence "Albert lost the election."³⁸ Some operators, such as 'it is true that,' when affixed to a sentence, penetrate to *all* of the consequences of that sentence. For instance, a (necessary) consequence of the sentence "Albert lost the election" is that "someone lost the election," given that Albert is someone and the nature of elections requires that someone loses. According to Dretske, operators can be distinguished from one another in terms of penetrability: an operator is *penetrating* insofar as that operator reaches all of the consequences of the sentence *or* an operator is *nonpenetrating* insofar as that operator may not even reach some of the elementary consequences of the sentence.³⁹ In a slightly different terminology, penetrating operators could be described as having the property of being *closed*, whereas nonpenetrating operators could be described as lacking that property. The following operators are fully penetrating: 'it is true that,' 'it is necessary that,' 'it is a fact that,' 'it is possible that.'⁴⁰ In stark contrast to these fully penetrating operators, consider, for instance, the nonpenetrating operator 'it is lucky that.' Let us affix this nonpenetrating operator to the sentence "I hit the bull's-eye," and suppose that a necessary consequence of this sentence is that "I either hit the bull's-eye or the side of the barn." The nonpenetrability of this operator is clear: if "it was lucky that" "I hit the bull's-eye," then "it was lucky that" "I either hit the bull's-eye or the side of the barn." In this case, the operator 'it was lucky that' is nonpenetrating, because, when affixed to the main sentence "I hit the bull's-eye," that operator fails to penetrate to a necessary consequence of the main sentence; namely, the operator 'it was lucky that' fails to even penetrate to that main sentence's elementary (necessary) consequence "I either hit the bull's-eye or the side of the barn."⁴¹ Translating this into a non-technical language, it may have been lucky that I hit the bull's-eye, but it was not a

³⁷ Dretske 1007, 1970.

³⁸ Dretske 1007, 1970.

³⁹ Dretske 1007, 1970. A bit of clarification is needed. By nonpenetrating operators Dretske does not mean that these operators have no penetrability whatsoever, but rather means that their penetrability is extremely limited: "They [nonpenetrating operators] fail to penetrate to some of the most elementary consequences of a proposition. I shall refer to this class of operators as *nonpenetrating operators*. I do not wish to suggest by this label that such operators are totally impotent in this respect (or that they are all uniform in their degree of penetration)"(Dretske 1970, 1009[original emphasis]).

⁴⁰ Dretske 1970, 1007.

⁴¹ The bull's-eye example is Dretske's (Dretske 1970, 1008).

matter of luck that I either hit the bull's-eye or the side of the barn – for even the worst player can usually hit at least the (broad) side of the barn.

But this classification of operators, according to differences in penetrability, remains incomplete until we take account of *semi-penetrating* operators, which possess a higher degree of penetrability than nonpenetrating operators, but a lower degree of penetrability than fully penetrating operators. Dretske examines a subclass of the semi-penetrating operators known as *epistemic operators*, such as the operators 'S knows that' or 'S has reason to believe that.'⁴² In these epistemic operators "S" stands for an epistemic subject or agent, and since the kind of knowledge with which Dretske is concerned is *propositional* knowledge, these operators will be affixed to a sentence, statement, or proposition by way of a "that"-clause, as in "S knows *that* P."

With these basic points in mind, Dretske issues a warning about the application of epistemic operators: "When we are dealing with the epistemic operators, it becomes crucial to specify whether the agent in question knows that P entails Q."⁴³ One way of understanding Dretske on this point is to realize that for the agent to know that P it is necessary for that agent to believe that P, since meeting the belief or doxastic condition is a necessary condition for knowledge. Likewise, for the agent to know the entailment Q, that agent needs to believe that Q. With this first warning still reverberating, Dretske next warns epistemologists against a tendency to *idealize* epistemic agents: "Were we all ideally astute logicians, were we fully apprised of all the necessary consequences (supposing this to be a well defined class) of every proposition, perhaps then the epistemic operators penetrate to all the *known* consequences of a proposition."⁴⁴ The reason why the first warning continues to echo through to the second is because, as John Hawthorne (2000) points out, "...one doesn't know all the logical consequences of one's beliefs."⁴⁵ Given that the non-idealized, concrete, epistemic agent does not know all of the logical consequences of her beliefs, it is crucial to specify whether that agent knows that P entails Q, as a particular consequence of her belief that P. It is worth noting, however, that defenders of the closure principle, such as Hawthorne (2005), among many others, who claim that an agent can, in principle, know all of the logical consequences as known entailments,

⁴² Dretske 1970, 1009.

⁴³ Dretske 1970, 1010.

⁴⁴ Dretske 1970, 1010 [original emphasis].

⁴⁵ Hawthorne 2000, 119. For some different versions of the closure principle, see Hawthorne's debate with Dretske: Hawthorne (2005) 26-43.

carefully make the following qualification on the closure principle: so long as the agent is *in a position* to know the entailment, then the closure principle holds. That is to say, were the agent to come to believe that entailment, retain knowledge of the main proposition, and “competently deduce” the entailment from the main proposition, then that agent would indeed have propositional knowledge of the entailment.⁴⁶ Furthermore, defenders of the closure principle do not think that making this qualification at all idealizes the agent in Dretske’s sense of treating the agent as an ideally astute logician.

Epistemic operators are the focus of his examination, because Dretske argues that traditional sceptical arguments exploit the semi-penetrability of these operators. In his characterization of the sceptical exploitation of epistemic operators, Dretske writes:

The traditional sceptical arguments exploit precisely those consequences of a proposition to which the epistemic operators do not penetrate, precisely those consequences which distinguish the epistemic operators from the fully penetrating operators.⁴⁷

To stop this form of exploitation, Dretske challenges *the thesis of penetrability*, viz., “If you do not know whether Q is true or not, and P cannot be true unless Q is true, then you (obviously) do not know whether P is true or not,”⁴⁸ since “Almost all sceptical objections trade on it.”⁴⁹ To gain a preliminary appreciation of the point of Dretske’s challenge, consider that philosophers who assume the thesis of penetrability (for example, G.E. Moore) could argue as follows: if one knows that the object on the windowsill is a (real) tomato, and if one knows that if the object on the windowsill is a (real) tomato then one knows that that object is not made of plastic, then one knows that the object is not made of plastic. In reply to these philosophers, however, the sceptic can make the following objection: if one does not know that the object on the windowsill is made of plastic, and if one does not know that the object on the windowsill is plastic then one does not know that the object is a (real) tomato, then one does not know that the object is a (real) tomato. This first argument made by some philosophers is similar to the logically valid pattern of *modus ponens*, whereas the sceptic’s objection is similar to the logically valid pattern of *modus*

⁴⁶ Hawthorne 2000, 199.

⁴⁷ Dretske 1970, 1012.

⁴⁸ Dretske 1970, 1011.

⁴⁹ Dretske 1970, 1011.

tollens.⁵⁰ In an attempt to break the stalemate between a *modus ponens* argument and *modus tollens* objection at the level of abstract argument, Moore (1959) claimed that if one knew something, then one carried the concrete certainty of common sense. When even sceptical alternatives become radical and introduce the error-possibilities of dreams, demons, illusions, and fakes, Moore thought that common sense spoke against their plausibility. Most contemporary philosophers do not equate knowledge with certainty, as Moore did, but some contemporary philosophers still find his *form* of argument philosophically attractive.⁵¹ Dretske does not directly question Moore's appeal to commonsense against the sceptic, but rather questions the thesis of penetrability on which both Moore's argument, and the sceptic's argument, trade.⁵² Dretske especially questions the theoretical commitment that endorsing the thesis of penetrability carries; namely, that it commits one to the view that epistemic operators are *fully* penetrating operators.

The most critical observation that Dretske makes about epistemic operators is that they do not operate on or penetrate to *presuppositions* made by agents. Consider the following passage:

...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.⁵³

The presuppositions that Dretske is discussing in the above passage are *not logical* or semantical presuppositions, but rather are ones made by agents (and *need not* be logical consequences of statements or propositions). The presuppositions that Dretske is dealing with here are what the

⁵⁰ Dretske holds that "they [the arguments] both assume that if S knows this is a P, and knows that every P is Q, then S knows that this is a Q. The only difference is that the skeptic performs a *modus tollens*, Moore a *modus ponens*" (Dretske 1970, 1011). But talking about these arguments as patterns of *modus ponens* and *modus tollens* has the danger of misleading his reader into thinking that he rejecting *principles of logic* rather than ones of epistemology. Recently, he claims that "It is important to distinguish closure from *modus ponens*, a principle of logic that it superficially resembles" (Dretske 2005, 13). I think that is an important to make this observation. The closure principle includes the word "knows," whereas the word "knows" *need not* be included in *modus ponens* or *modus tollens* arguments, since such arguments need not be epistemological.

⁵¹ Some contemporary philosophers, such as Ernest Sosa (1999) and Duncan Pritchard (2002, 2005), call themselves Neo-Mooreans, because they make an argument similar to the argument resembling *modus ponens* that Moore makes against the sceptic. So, these philosophers are committed to the closure principle. But these philosophers, unlike Moore, do not equate knowledge with certainty, but think that the counterfactual *safety* principle is the key to knowledge.

⁵² Dretske 1970, 1011.

⁵³ Dretske 1970, 1014 [original emphasis].

agent thinks of as epistemically relevant or irrelevant. It is best to demonstrate what Dretske means by “presuppositions” by using an example. According to Dretske, when I *state* that “I know that the wine in the bottle is red,” I leave *unasserted* that *it is* wine – of the liquid in the bottle that looks red to me. That *it is* wine has been left *unasserted* and rather is presupposed in my stating that “I know that the wine in the bottle is red.” Dretske contends that epistemic operators do not penetrate to these presuppositions of my statement. It is important to point out, however, that he is not denying that I may have evidence for, say, the liquid in the bottle being red wine or that I find it plausible that there is red wine in the bottle. I may have checked by smelling, looking, or even having a glass of the liquid, or I may have another reason to believe that it is plausible that it is wine and not, say, water with red food-colouring in it. But Dretske thinks that often we do not make any such checks and make presuppositions instead. Moreover, the set of plausible reasons or evidence we have when making a statement often remains at the level of presupposition too.⁵⁴ Notice that all of the presuppositions are made by the agent and have to do with what the agent knows from within that agent’s perspective. This is a very important theoretical move that Dretske makes, and I will reveal why in a moment.

Dretske thinks that sceptical objections rely on what he calls “contrast consequences.” Contrast consequences are predicates incompatible with the predicate contained in the main statement under examination. For instance, consider the main statement “the wall *is red*.” This statement contains the predicate ‘is red,’ which could be contrasted with the predicates ‘is white.’ The predicate ‘is red’ could even be contrasted with the predicate ‘(but) is white cleverly illuminated to look red.’⁵⁵ This latter predicate is of the kind that is often contained in sceptical alternatives. The sceptic, accordingly, contends that the agent needs to know the falsity of a sceptical alternative in order to know an ordinary empirical proposition. These sceptical alternatives always contain contrast consequences. Endorsing the thesis of penetrability, then, would mean that the following conditional in, say, the red wall case, would need to be met in order to have knowledge: “...if anyone knows that the wall is red he must know that it is not

⁵⁴ Dretske 1970, 1014.

⁵⁵ Dretske 1970, 1015.

white cleverly illuminated to look red.”⁵⁶ Generalizing from this case means that the same kind of conditional would need to be met in order to know ordinary empirical propositions.⁵⁷

Recall the important point that sometimes our evidence for statements is left at the level of presupposition and based simply on what we think is most *plausible*. For instance, the fact that the bottles of wine advertise a Merlot or Pinot Noir often supplies ample evidence to believe them to contain red wine. Yet Dretske thinks that *ordinary*, often presupposed, evidence and the often good enough reasons we have for knowing, say, that the wall is (painted) red, are not good enough reasons to eliminate the sceptical possibility that the room is set up with elaborate trick lighting to make a white wall look red. Consider Dretske’s view on the matter:

We seldom acquire any *special* reasons for believing the lighting normal although we can talk vaguely about there being no reason to think it unusual. The fact is that we habitually take such matters for granted, and although we normally have *good* reasons for making such routine assumptions, I do not think these reasons are sufficiently good, not without special precautionary checks in the particular case, to say of the particular situation we are in that we *know* conditions are normal.⁵⁸

Notice that Dretske is not claiming that this sceptical alternative is plausible or one that is frequently realized; rather, he knows that there is a bit of Moorean common sense within each of us, which very often refuses to take such possibilities seriously. The implausibility of the case, however, is not good enough to know that the sceptical alternative is false.⁵⁹ Nevertheless, Dretske resists the idea that we always need to make *special* precautionary checks to have knowledge, since endorsing that idea would be tantamount to capitulating to the sceptic.

Dretske does not capitulate to the sceptic, because he holds that we can continue to know ordinary empirical propositions, such as that the wall is red or that the animal in the pen is a zebra, on the basis of ordinary evidence, despite failing to know the falsity of sceptical alternatives. That it is to say, sceptical alternatives often remain irrelevant and non-threatening, since knowledge often only requires ordinary evidence and rarely requires the agent to make special precautionary checks. This exposition puts us in a position to truly appreciate Dretske’s

⁵⁶ Dretske 1970, 1015.

⁵⁷ Dretske 1970, 1015.

⁵⁸ Dretske 1970, 1015 [original emphasis].

⁵⁹ Dretske 1970, 1016.

rejection of the thesis of penetrability (or the closure principle). Consider what Dretske thinks he has shown:

What I am suggesting is that we simply admit that we do not know that some of these contrasting “sceptical alternatives” are not the case, but refuse to admit that we do not know what we originally said we knew. My knowing that the wall is red certainly entails that the wall is red; it also entails that the wall is not white, in particular, it entails that the wall is not white cleverly illuminated to look red. But it does not follow from the fact that I know that the wall is red that I *know* that it is not white cleverly illuminated to look red. Nor does it follow from the fact that I know that those animals are zebras that I know that they are not mules cleverly disguised to look like zebras. These are some of the contrast consequences to which epistemic operators do not penetrate.⁶⁰

This is a denial of the closure principle, because Dretske holds that I can know that the animal is a zebra on the basis of ordinary evidence, and I can also know that the animal is a zebra *entails* that the animal is not a cleverly disguised mule, but yet I can still fail to know that the animal is *not* a cleverly disguised mule, because I lack special evidence to know this and epistemic operators are only semi-penetrating.

This completes my exposition of Dretske (1970). We are now in a position to see the serious mistaken internalist assumption that Dretske makes, which causes him to reject the closure principle. This internalist assumption is in regard to the *agent's* presuppositions. As Michael Williams's argues, Dretske makes a serious mistake when he conflates the presuppositions that the *agent* makes about the proposition with the *logical* presuppositions of that proposition (which are logical consequences). Before I point out this serious mistake, I want to briefly focus on the good part of Dretske's argument for closure failure. Recently, Dretske has received a barrage of harsh criticism for denying the closure principle. Some critics hold that the closure principle has an “axiomatic status,”⁶¹ others contend that “to many it seems crazy to actually deny the principle,”⁶² and still others point out that denying the closure principle makes Dretske's epistemology open to abominable conjunctions of the form “John knows that that is a zebra, but he does not know that that is not a painted mule.”⁶³ Instead of adding to the barrage, I

⁶⁰ Dretske 1970, 1016-1017.

⁶¹ Cohen 1999, 68.

⁶² Cohen 2000, 102.

⁶³ This quotation comes from Stanley 2005, 19. “Abominable conjunctions” are discussed by DeRose 1995, 27-29. DeRose, however, has Nozick (1981) in mind when he brings up these abominable conjunctions. Like Dretske, Nozick also denied the principle of closure.

want to tell you about both the good part as well as the bad part of Dretske's argument for rejecting the closure principle. Hopefully, I can find some balance in my assessment.

The good part of Dretske's argument is that it prevents a kind of bootstrapping in epistemology. By "bootstrapping" I mean that Dretske prevents epistemic subjects from being aware of the reliability of their own perceptions and environmental conditions. This prevents subjects from making inferences from these perceptions in order to expand the scope of their knowledge. Dretske especially wants to prevent subjects from knowing – from a first-person perspective – the falsity of sceptical alternatives by making deductions from what these subjects take themselves to already know. With Dretske (1970) in mind, Duncan Pritchard (2005) states:

In effect, [Dretske's argument for] the failure of closure simply reflects the 'anti-bootstrapping intuition' that one cannot come to know the empirical truths that are presupposed in one's empirical knowledge simply by reflecting upon that knowledge.⁶⁴

Recently, Dretske (2005) claims that any act of perception carries with it what he calls "heavyweight implications" that cannot be similarly perceived.⁶⁵ Heavyweight implications imply that there is an external world and that scepticism is false.⁶⁶ For instance, I see that there is a computer in front of me, which implies that there is a physical object in front of me. But that does not, according to Dretske (2005), mean that I similarly *see* (or perceive in the same manner) that a physical object is in front of me.⁶⁷ That there are physical objects is a heavyweight implication of my seeing that there is a computer in front of me. Dretske claims that all heavyweight implications entail the falsity of radical sceptical alternatives.⁶⁸ For example, my seeing that there is a computer in front of me entails the falsity of the radical sceptical alternative that, say, I am being tricked by an omnipotent evil demon to believe that there is a computer in front of me; although, if I were being tricked by an omnipotent evil demon, presumably, I would have no idea of it. Dretske denies that I can know the falsity of this radical sceptical alternative from my particular act of perception. To say otherwise would allow for bootstrapping. Since all ordinary empirical propositions known by perception entail heavyweight implications, Dretske thinks that

⁶⁴ Pritchard 2005, 40. In the same vein, Michael Williams (1991) writes: "We have taken note of the deep intuition underlying denials of closure: that no piece of knowledge can guarantee the fulfilment of its reliability conditions. If it could, we could bootstrap ourselves into knowing that we know certain propositions to be true"(Williams 1991, 346).

⁶⁵ Dretske 2005, 20.

⁶⁶ Dretske 2005, 20-24.

⁶⁷ Dretske 2005, 14.

⁶⁸ Dretske 2005, 23.

accepting the closure principle either enables bootstrapping in epistemology or has sceptical-breeding consequences.⁶⁹

In his (2005) recent argument against the closure principle, I find that Dretske continues to respond to the sceptic as well as to Moore. Recall that Dretske (1970) is just as critical of Moore for accepting the thesis of penetrability as he is of the sceptic for doing so. Essentially, in Dretske's view, Moore's response to scepticism has about as little plausibility as when Doctor Johnson, in a legendary demonstration of "common sense" against Bishop Berkeley's idealism (a philosophical position which was apparently firmly grounded in common sense too, according to Berkeley), kicked a stone and then said that he had refuted the good Bishop thus. In my opinion, the anti-bootstrapping "intuition" that these commentators and critics of Dretske have noted is reflected in Dretske's resistance to those philosophers who use the same form of argument that Moore used to know a heavyweight implication from, say, raising one hand and then the other, observing that hands are physical objects, and finally declaring that there can be no scepticism about the external world – *thus*. To argue in this manner often requires a trust or faith in the reliability of one's own perceptions and reasons which is not altogether conducive to good philosophy.⁷⁰ With these philosophers in mind, Dretske writes, "Despite its philosophical credentials (Moore, 1959), that sounds like chutzpah, not philosophy, to me."⁷¹ Instead of choosing between two unpalatable options that come from accepting the closure principle, either a bootstrapping response or radical scepticism, Dretske (1970, 2005) chooses the remaining option, which is to deny the closure principle. Dretske then carefully mitigates the force of scepticism by claiming that epistemic operators are only semi-penetrating, and claiming that only relevant alternatives thereby have to be ruled out in order to have knowledge. I consider Dretske's attempt to prevent bootstrapping in epistemology to be the good part of his argument.

The bad part of Dretske's argument is that he is confused about the relationship between the agent's presuppositions and the closure principle (which concerns *logical* presuppositions). Michael Williams (1991) is particularly clear about what is the matter: "...Dretske conflates the presuppositions of a proposition P with the presuppositions of my knowing that P. As a result, he

⁶⁹ Dretske 2005, 23.

⁷⁰ Consider again the quotation from Michael Williams in Note 64 above.

⁷¹ Dretske 2005, 24.

treats as my failure to know that P what is in fact only my failure to know that I know.”⁷² As evidence of Williams’s charge of Dretske’s conflating the two, consider again that Dretske holds that “...it does not follow from the fact that I know that the wall is red that I *know* that it is not white cleverly illuminated to look red.”⁷³ To point out the serious mistake that Dretske makes, I will use his example of the red wall. Let P stand for “the wall is (painted) red” and Q stand for “the wall is not white cleverly illuminated to look red.” First consider that if I am able to rule out Q, then I must, according to Dretske, make special precautionary checks; that is to say, I must make sure that the environmental conditions are normal. Even lacking such checks, however, Dretske maintains that I can still know that P. Given that I have no special evidence, it only means that I am unable to rule out Q. Moreover, this may be so despite my knowing that P entails Q. This is not, however, a case of closure failure, since it merely shows I do not know that Q is false, although I may know that P entails Q. What Dretske needs to provide is a case in which I know that P, I know that P entails Q, but still fail to know that Q.

The problem with Dretske’s argument for closure failure is precisely that *my* presuppositions (the presuppositions I make as an *agent*) are not found among the known logical consequences of the statement I am said to know.⁷⁴ My knowing from a first-person perspective that the wall is red *does not entail* my knowing from a first-person perspective that the wall is not white cleverly illuminated to look red. Let me be very clear. It is certainly true that P entails Q, no one disputes this; but it is false that my knowing that P *entails* my knowing that Q. Essentially, I can know that P entails Q, but I cannot know that P entails *my knowing* that Q using the closure principle. The closure principle is silent about *my* presuppositions, because they are not among the logical consequences of the statement. Accordingly, my presuppositions fall outside of the scope of epistemic operators. Crucially, that does not mean that the closure principle does not hold, for the closure principle concerns the logical consequences of statements.⁷⁵ Dretske does not separate the logical consequences or entailments of the statement from the presuppositions that the agent makes about the statement. He rather conflates the two, which is why Williams’s criticism is so powerful. All of the cases Dretske adduces for closure failure rest on this type of confusion.

⁷² Williams 1991, 332.

⁷³ See again Dretske’s example of the red wall in the quotation above for more details (Dretske 1970, 1016).

⁷⁴ Williams 1991, 332.

⁷⁵ Williams 1991, 335.

I want to focus on issues of epistemological externalism and internalism and their bearing on Dretske's conflation, since these issues will figure prominently in my diagnosis of the error which will dissolve, in a non-contextualist manner, Cohen's so-called sceptical "paradox." A number of prominent critics, including Williams (1991) and Pritchard (2005), have focused on Dretske's commitment to "epistemological externalism," and how, after he makes this commitment, Dretske's argument for closure failure becomes poorly motivated. Yet before I define some terminology, I feel I must issue a brief warning.

Although both Williams (1991) and Pritchard (2005) criticize Dretske for breaking "externalist commitments," their criticisms are anachronistic. Dretske wrote "Epistemic Operators" in 1970, but the term "externalism" about *knowledge* was not coined until three years later by David Armstrong in 1973. Moreover, almost a decade after Dretske (1970), discussions of externalist *justification* began with Alvin Goldman's (1979) work and Laurence Bonjour's (1980) subsequent criticism of that work. There is always a danger of one's misconstruing or "reading into" another's theory when that theorist could not have had the privilege of hindsight and the debates have advanced since the construction of the theory. Despite these criticisms being anachronistic, they are too important for this danger to stop their expression.

Externalist accounts of *knowledge* claim that knowledge is a state in which there is a law-like relationship or connection which holds between the subject's belief and truth – often understood as a non-accidental, nomological connection between belief and the way that the world is.⁷⁶ In the words of David Armstrong (1973) who first stated this account of knowledge:

According to "Externalist" accounts of non-inferential knowledge, what makes a true non-inferential belief a case of *knowledge* is some natural relation which holds between the belief-state, Bap [an agent a believes that p], and the situation which makes the belief true. It is a matter of a certain relation holding between the believer and the world. It is important to notice that, unlike "Cartesian" and "Initial Credibility" theories, Externalist theories are regularly developed as theories of the nature of knowledge *generally* and not simply as theories of non-inferential knowledge.⁷⁷

As Armstrong's statement indicates, externalism is largely a reaction against "Cartesian" theories or traditional accounts of epistemic justification, which limit knowledge and epistemological

⁷⁶ Williams 1991, 319. I emphasize "of knowledge" or "about knowledge," because there are also externalist accounts of epistemic *justification*, which have gained considerable popularity lately, especially in the form of Goldman's (1979) process reliabilism.

⁷⁷ Armstrong 1973, 157 [original emphasis].

analysis to the first-person perspective of the believer.⁷⁸ The project most closely associated with internalism, which often starts from the first-person perspective, is called *foundationalism*, which states that there are basic, non-inferentially non-doxastically justified beliefs upon which the edifice of knowledge rests.⁷⁹ There has always been a close relationship between internalism and foundationalism, with some critics, such as Williams (1991), even contending that “...foundationalism and internalism are two names for the same thing.”⁸⁰ Epistemological *Internalism* states that the subject’s reasons which turn a true belief into knowledge *must* all be cognitively or reflectively accessible to (that is to say, internal to) that subject.⁸¹ Internalists often hold that knowledge is gained through “a priori reasoning, introspective awareness, or through a memory of the knowledge that has been gained by either of these ways.”⁸² Importantly, internalists deny that any of the factors that convert belief into knowledge – besides truth – remain *external* to the subject’s cognitive awareness or grasp. Externalism denies internalism by stating that whether the subject has knowledge can depend on whether that subject’s belief *in fact* stands in relation to the truth, especially as that belief stands in relation to what the world is like. According to externalists, whatever turns the subject’s true belief into knowledge can remain external to that subject’s cognitive awareness, and thereby externalists deny that there is any special accessibility (or distinctively internalist) requirement for knowledge.⁸³

The reason why issues of externalism and internalism about knowledge are pertinent to this discussion of Dretske (1970) is that Dretske explicitly makes externalist commitments to

⁷⁸ Kornblith 2001, 4. Importantly, accounts of knowledge, some of which are anachronistically referred to as “externalist,” were also put forth to solve the Gettier problem, see: Gettier (1963); Goldman (1967); Dretske (1971); Nozick (1981).

⁷⁹ For the “building blocks” of foundationalism, which are basic beliefs, the term “basic” denotes the property of being self-justified; “non-inferential” denotes the property of not being based on an inference; “nondoxastic” denotes the property of basic beliefs not receiving their justification from other beliefs. According to Matthias Steup (1996), “Foundationalists would say that they [basic beliefs] receive their justification from a *nondoxastic* source—that is, not from other beliefs, but from somewhere else: perception or introspection, for example (Steup 1996, 90 [original emphasis]).

⁸⁰ Williams 1991, 323. The full quotation is the following: “What I find especially puzzling is the way nonclosure links externalism with concessions to the sceptic. For I have argued that scepticism depends on foundationalism and that foundationalism and internalism are two names for the same thing” (Williams 1991, 323).

⁸¹ Furferton 2006, 56.

⁸² Pritchard 2005b, 188.

⁸³ Recently, the debate between externalists (mainly, reliabilists) and internalists has been cast in terms of epistemic *justification*, with *justificatory* externalists contending that *reliably* produced beliefs are justified – beliefs produced by often perceptual reliable mechanisms, reliable in the sense that they yield a higher ratio of true beliefs than false ones – whereas internalists deny this and hold that justification must be internal. The most famous works in this literature include: Goldman 1979, 1980; BonJour 1980.

knowledge, but compromises these commitments by making an internalist assumption when he denies that subjects can know the falsity of sceptical alternatives. Dretske makes a commitment to externalism by holding that environmental conditions often are normal and that subjects do not normally have to know what the environmental conditions are like to have knowledge. As Williams (1991) claims, “We can see right away that Dretske’s treatment of his examples commits him to externalism. Which alternatives I need to rule out to know that the animals are zebras depends on the facts of the situation... not on what I know.”⁸⁴ Dretske holds that subjects, however, do not know the falsity of sceptical alternatives, a verdict which comes from the limited penetrability of epistemic operators, the special evidence that that subject lacks, but especially from considering the subject’s presuppositions and what that subject can and cannot know from within the *first-person perspective*. As an externalist, Dretske should not, however, claim that a subject needs to know, from a first-person perspective, what the environmental conditions are like to know that sceptical alternatives do not obtain.⁸⁵ At best, as an externalist, Dretske is only entitled to claim that subjects do not know that they know that sceptical alternatives do not obtain. But this claim has an extremely odd ring to it, for externalists simply deny the claim that subjects need to know that they know in order to have knowledge. This claim may, however, ring true for the internalist.

Both Williams (1991) and Pritchard (2005) contend that only internalists directly confront the problem of scepticism, since internalists claim that subjects must have cognitively accessible reasons to have knowledge, even when that knowledge is supposedly of the falsity of sceptical hypotheses. Internalists sometimes claim, more generally, that subjects must have reasons about whether the environmental conditions are normal or to trust the reliability of their perceptual abilities in order to have knowledge.⁸⁶ That is to say, for the internalist, *my* knowing that the animal is a zebra will sometimes require *my* knowing that the environmental conditions are conducive to *my* knowing this, and sometimes I must, therefore, rule out the possibility that I am looking at a cleverly disguised mule to know that I am looking at a zebra. Since the internalist is committed to the view that all of the subject’s reasons must gain their justification from within a first-person perspective, which means that they must be reasons of which the

⁸⁴ Williams 1991, 331.

⁸⁵ Williams 1991, 333.

⁸⁶ See: Williams 1991, 323-325; Also: Pritchard 2005, 44.

subject is cognitively aware, the sceptic will raise precisely those error-possibilities which would be, to use Pritchard's (2005) apt term, "subjectively indistinguishable" to that subject, if the error-possibility were to obtain.⁸⁷ As Jonathan Vogel (1990, 1999) argues, however, an internalist can hold that the subject indeed knows that the animal is a zebra and *not* a cleverly disguised mule by knowing that public zoos exhibit genuine specimens, by having ample statistical evidence against the possibility of fraud or hoax.⁸⁸ In this *particular* example, I happen to agree with Vogel that this statistical evidence and certain well-known facts about public zoos could be the subject's reasons which could count against the possibility of fraud or hoax. But any comfort here is short-lived for the internalist, because the sceptic is simply unrelenting. Using an arsenal of error-possibilities, including dreams, demons, illusions, and fakes, the sceptic will bombard the subject with these error-possibilities, a subject who is seemingly defenceless to sceptical assault, because of the extreme vulnerability of the first-person perspective. When the (radical, closure-based) sceptic is on the prowl, the rather tame error-possibility of a cleverly disguised mule quickly turns into the zoo visitors, as putative knowers, having to also rule out the wild and radical error-possibility that they are not being deceived by an omnipotent evil demon bent on deceiving anyone who claims knowledge. *Radical* scepticism remains a major problem *for* the internalist, precisely because of the susceptibility of the first-person perspective.⁸⁹ Yet radical scepticism would not be a problem for Dretske, if he were to steadfastly adhere to his externalist commitments.

In summary, Dretske's most serious mistake is that he conflates the agent's presuppositions with the logical consequences of a statement. My presuppositions are *not* operated on by epistemic operators. This is precisely because, in making a statement, my presuppositions fall outside of the scope of epistemic operators and thereby, contra Dretske, are *not logical* consequences of statements. Dretske's argument for closure failure gives the appearance of closure failure, but this is merely an appearance – my knowing that P entails Q

⁸⁷ Pritchard 2005, 24. By the term "subjectively indistinguishable," Pritchard means that that subject would not be able to tell the difference if that subject were, say, either a brain-in-a-vat, or in the matrix, or the victim of deception by an all-powerful evil demon, or one of many such scenarios that the sceptic uses as error-possibilities.

⁸⁸ Vogel 1990, 15; Vogel 1999, 160.

⁸⁹ I distinguish radical and global scepticism from moderate and local scepticism below when I discuss the so-called sceptical "paradox" and Cohen's (1988) account. Michael Williams (2001) defines *radical scepticism* as "the thesis that we never have the slightest justification for believing one thing rather than another. The important point about radical scepticism is that it is formulated in terms of justification" (Williams 2001, 59). This fits in with our discussion, because internalist theories of knowledge are specifically about justification, traditionally understood.

does *not* mean that *my* knowing that P entails *my* knowing that Q. Therefore, Dretske's examples of the non-penetrability of the agent's presuppositions cannot be counterexamples to the closure principle. In addition to this conclusion, I am suggesting that Dretske's (1970, 2005) complaint about bootstrapping in epistemology reflects his dissatisfaction with what are really *internalist* "solutions" to scepticism, in the sense that Dretske vehemently denies that the subject can know, from a first-person, internalist perspective, the falsity of sceptical alternatives. Dretske's (2005) latest purported examples of closure failure tend to support my suggestion as well.⁹⁰ Dretske's dissatisfaction is common among critics of internalism who share his anti-bootstrapping "intuition," and who also claim that radical scepticism is indeed a real problem for the internalist. This further suggests that only internalists who accept the closure principle inherit the problem of radical scepticism.⁹¹ But critics of Dretske (1970), such as Williams (1991) and Pritchard (2005), stress that Dretske's externalist commitments are compromised by his insistence both that subjects do not know the denials of sceptical alternatives and that the closure principle does not hold. The externalist should not demand, however, that the subject needs to know, from a first-person perspective, that sceptical alternatives do not obtain. Minimally, the externalist should claim that so long as a relation between the subject's belief and the world holds, then that subject has knowledge, even in the face of sceptical threat. Radical scepticism is *not* a problem for the externalist, or at least there is no obvious reason to think so. I fear, however, that Dretske (2005) continues to labour under mistaken assumptions about speaker presuppositions, about the scope of epistemic operators, and especially about issues pertaining to internalism and externalism. Without realizing these mistaken assumptions, Dretske continues to try to argue for closure failure. Ultimately, Dretske's argument for closure failure will be to no effect.

⁹⁰ There is a dangerous ambiguity that Dretske flirts with in many of his examples of perception and heavyweight implications, which involves the term "sees." The term is *factive*, insofar as if one properly is said to (visually) "see" something, then that something which one sees must *in fact* exist. I suppose one does not need to be explicitly conscious of the object one (visually) "sees" to (visually) "see" it, although an ascription of the term "sees" would be factive. But notice that there is another sense of the term "see," in which one indeed needs to be explicitly conscious of whatever it is that one (cognitively) "sees." Often the second sense takes a "that"-clause, as in *S sees that p*, which denotes a kind of understanding, which is also perhaps properly epistemic. Of course, there are expressions which refer to the latter sense without the "that"-clause, as in "I see what you mean." However, Dretske's examples of closure failure all have to do with this latter cognitive sense of the term "sees." For example, I see *that* there is a computer in front of me does not imply that I see *that* there is a physical object in front of me. But if I (visually) see a computer in front of me, there has to be a physical object for me to (visually) see it; otherwise, I may be hallucinating or "seeing things," where "things" stand for nothing in particular and hence indicate a non-factive state. Dretske is again here concerned with what can be known from the first-person perspective, and especially with preventing bootstrapping. Crucially, externalism, properly understood, does not allow for bootstrapping.

⁹¹ Pritchard (2005) is extremely clear on this point. See: Pritchard (2005) 32-66.

1.5 Stine's Contextualism and Criticisms of Dretske

Stine tries to improve upon Dretske's relevant alternative theory by contextualizing it in order to avoid an "equivocation" that Dretske (1970) commits in his argument for closure failure. After making this emendation, Stine claims that her contextualist version of the relevant alternatives theory finds a suitable corrective to scepticism.⁹² She also makes some extremely intriguing criticisms of Dretske that suggest her own externalist leanings.⁹³ In the end, however, it is difficult to square Stine's externalist leanings with her supposed corrective to scepticism.

According to Stine, a glaring problem with Dretske's argument for closure failure is that he commits "...some logical sin akin to equivocation."⁹⁴ To see this, consider again that on Dretske's relevant alternatives theory, I can know that the animal in the pen is a zebra, I can know that if the animal in the pen is a zebra then this entails that the animal is not a cleverly disguised mule, and yet I can fail to know that the animal is not a cleverly disguised mule, because this is a sceptical alternative to which epistemic operators do not penetrate. But, as Stine keenly notes, the term "know" is equivocal in this argument, because the set of alternatives for the antecedent is different from that of the set for the consequent. For Dretske, in order to know the consequent "the animal is not a cleverly disguised mule," the subject would need to rule out the sceptical alternative that the animal is not a cleverly disguised mule. This particular alternative remains *irrelevant*, however, to whether the subject knows the antecedent "the animal is a zebra," and which is thereby not a part of the set of relevant alternatives associated with the antecedent. This exposes Dretske's "equivocation."

To "absolve" Dretske of his "equivocation," Stine claims that "...to be justified in claiming to know p (or simply to know p) it is sufficient to be able to rule out alternatives *relevant to that context*."⁹⁵ On this point, Stine crucially contends that there are different criteria appropriate to a given context, which Dretske fails to keep separate or even recognize. In identifying these different criteria, Stine claims:

⁹² Jason Stanley (2005) claims that "Contextualism was originally advanced both as an interpretation (Goldman 1976: 776-7) and as an improvement (Stine 1976) of Dretske's *Relevant Alternatives Theory* (Dretske 1970)"(Stanley 2005, 17).

⁹³ Notice that Stine is writing in 1976, three years after Armstrong (1973) wrote his famous book. Stine (1976) actually quotes Armstrong in the paper, see: Stine 1976, 251. Moreover, she provides the quotation in her section titled *Skepticism*, which I find extremely intriguing. I shall briefly discuss this in the body of the text.

⁹⁴ Stine 1976, 256.

⁹⁵ Stine 1976, 249 [my emphasis]. Later in her paper, Stine again makes the same claim, see: Stine 1976, 259.

It is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts. It is one thing in a street encounter, another in a classroom, another in a law court – and who is to say it cannot be another in a philosophical discussion?⁹⁶

As I understand Stine, when John, a visitor to the public zoo, for instance, claims to know that he is looking at a zebra, then he only needs to rule out the other animals at the zoo which look like zebras, such as gazelles or antelopes. Yet, when Sherlock, an animal inspector who inspects animals at a private zoo, say, claims to know that the animal is a zebra, he may have to rule out not only the alternatives that John needs to rule out, but also the alternative that the animal is not a mule cleverly disguised to look like a zebra. The criteria appropriate to John's context are loose, whereas the criteria appropriate to Sherlock's context are much tighter. In Stine's view, the relevance of alternatives is largely determined by context. But notice Stine's provocative question at the end of the above quotation. Philosophical discussions are sometimes about radical scepticism. Although she finds the sceptic's criteria to be "perverse," nevertheless, because the concept of knowledge is so elastic as to allow for differences of criteria relative to different contexts, Stine maintains that "...we cannot legitimately go so far as to say that their [sceptical philosophers'] perversity has stretched the concept of knowledge out of all recognition – in fact they have played on an essential feature of the concept."⁹⁷ That essential feature of the concept of knowledge is contextualism, with its tighter criteria appropriate to different contexts. Consequently, even radical sceptical alternatives are sometimes relevant alternatives, at least according to the rarely legitimate criteria of philosophical discussions about radical scepticism.

Stine's remarks about the essential characteristic of our concept of knowledge, with different criteria for different contexts of use, however, too strongly suggest that the application of "knowledge" depends simply on different social uses. If Stine indeed holds this view that she is strongly suggesting, then she has things backwards. The concept of knowledge straightforwardly conceptually constrains these different social uses, and even excludes some of them that misapply the concept. To my mind, this makes it an open question whether the concept of knowledge should be applied to philosophical discussions of radical scepticism.

⁹⁶ Stine 1976, 254.

⁹⁷ Stine 1976, 254.

An instance of a type of social use to which the concept knowledge is misapplied, because it uses criteria that conceptually fall short of those compatible with the concept of “knowledge,” would be the criteria used in, say, trivia game shows. On such shows, when a contestant is said to “know” the answer to a trivia question, this is simply a matter of that contestant giving the right answer or supplying correct information. But sometimes contestants make guesses, and sometimes these guesses turn out to be lucky. According to the criteria for “knowing” in this particular social use, the contestant still “knows” the answer, despite making what was really a lucky guess. Evidence of this is that the contestant continues to enjoy the prize money. The concept of knowledge, however, is incompatible with a lucky guess, which is a kind of happy or fortunate accident. This is an example, then, in which the concept of knowledge should conceptually constrain the social use, and not the other way around: the concept of knowledge should not stretch out to accommodate the social use. From an epistemic view, we cannot consider the contestant making a lucky guess to have knowledge, because lucky guesses do not count as knowledge on any standard analysis of the concept of knowledge.

I find it strange that Stine labels the criteria of sceptical contexts “extreme” and “perverse,” but yet finds the criteria appropriate to these contexts compatible with our concept of knowledge. The criteria of philosophical discussions of radical scepticism require subjects to achieve knowledge *as* certainty by having indubitable and infallible true beliefs. Stine may be right that our concept of knowledge is not stretched totally out of recognition by including some social uses with these extremely tight criteria, but if there are social uses in which the criteria are too loose to fall within the conceptual compass of knowledge, as in the example of trivia game shows, then I see no obvious reason to make concessions to the sceptic on the sole basis of observing that there are some philosophical discussions with extremely tight criteria. It may happen that these criteria for philosophical discussions of scepticism rest on a conceptual confusion similar to the too loose criteria of the trivia game shows, which may be the very reason why such criteria for philosophical discussions of radical scepticism are described as “extreme” and “perverse.” Stine needs to provide an argument, rather than simply making observations of different social uses and their applications of the word “knowledge,” to make a persuasive case for the inclusion of sceptical contexts and the appropriateness of these applications. It is not obvious to me that the sceptic is playing on *our* concept of knowledge, with the difference being merely one of degree in terms of tighter criteria or thresholds of justification; or whether the

sceptic has in mind distinct conceptual requirements, such as psychological certainty (in which any doubt by the subject counts against knowledge), which go *beyond* our concept of knowledge.

To be clearer about Stine's contextualism, the kind of contextualism that she has in mind is *ascriber* contextualism, since it concerns how the tightness of the criteria is determined by the error-possibilities raised in the conversational context of the speaker ascribing "knowledge," and how these error-possibilities affect the very meaning of the "knowledge"-claim – although, as we have seen, DeRose (1992) finds Stine somewhat unclear on the matter, because she does not declare which of Goldman's two views of relevance she takes.⁹⁸ Stine holds that the literal meaning of John's claim that "I know that the animal is a zebra" means one thing in a conversation during his leisurely visit to the public zoo, but would mean another thing if John were to make the claim in a conversation about the identity of the animal with, say, Sherlock, the animal inspector, in which conversational context the alternative that the animal is a cleverly disguised mule is a relevant alternative.⁹⁹ The tighter criteria peculiar to that context make this so. Stine claims that unless we hold the set of relevant alternatives relative to a conversational context, we will remain prone to committing an equivocation of the kind that Dretske makes.

Stine severely criticizes Dretske for his evidential requirements on knowledge; particularly, for his requirement that subjects need to make special precautionary checks to know that sceptical alternatives do not obtain. As we have seen, this eventually leads Dretske to reject the closure principle when he pairs this claim about special evidence with his claim about speaker presuppositions (while making an internalist assumption in the process). In response to Dretske's evidential demands, Stine charges that "Dretske is deluded by the fact that many knowledge claims require evidence on the part of the knower into thinking that all knowledge claims require evidence."¹⁰⁰ But Stine should be careful here, because, although Dretske holds that subjects need evidence to have knowledge, he distinguishes between two kinds of evidence, ordinary and special evidence. Unlike Dretske, however, Stine claims that subjects can indeed know that sceptical alternatives do not obtain *without* evidence. To justify this claim, Stine writes: "...if the negation of a proposition is not a relevant alternative, then I know it – obviously,

⁹⁸ Although these papers came out in the same year, Stine (1976) is very aware of Goldman's (1976) two views of "relevance," see: Stine 1976, 252; also: DeRose 1992, 920.

⁹⁹ Stine 1976, 255.

¹⁰⁰ Stine 1976, 257.

without needing to provide evidence....”¹⁰¹ For instance, if I know that the animal in the pen is a zebra, then, Stine is claiming, I know that the animal is not a cleverly disguised mule, because this alternative is an *irrelevant* alternative or one that has already been ruled out, given that to know that the animal in the pen is a zebra in the first place requires that I have ruled out all of the *relevant* alternatives to that proposition. Consequently, in Stine’s view, I can know that the animal is not a cleverly disguised mule without evidence and simply on the basis of this alternative being irrelevant. Yet this should not suggest that the subject is privy to the evaluation of knowledge and knows which alternatives are relevant and which are irrelevant. To demand so much would be overly demanding, because it would force the subject to have to know when that subject has ruled out all of the alternatives to the main proposition. Fortunately, Stine has another way of making her claim about how it is that subjects sometimes gain knowledge of the falsity of sceptical alternatives.

Intriguingly, Stine suggests that the reliability of the subject’s belief depends in large part on being in the right circumstances, and that that subject in those circumstances can, accordingly, know that the animal is a zebra, without having any particular evidence against the cleverly disguised mule alternative. Stine even cites a famous passage by David Armstrong (1973) to make this point.¹⁰² Combining Stine’s above suggestion with her view that Dretske distorts “the evidence picture of knowledge,” strongly indicates her externalist leanings.¹⁰³ But Stine’s contention that the subject can know that a sceptical alternative does not obtain, because of the irrelevance of the alternative, has been met with stiff resistance. Palle Yourgrau (1983), for instance, contends that there are some knowledge-claims that do not require evidence, but that these are usually knowledge-claims based on “self-presenting states,” such as being in pain.¹⁰⁴ At best, Yourgrau claims that Stine is revealing the implausibility of sceptical alternatives. But Yourgrau holds that the implausibility of sceptical alternatives is insufficient to know the falsity of sceptical alternatives.¹⁰⁵ So much is certainly true from an internalist point of view. Moreover,

¹⁰¹ Stine 1976, 258.

¹⁰² Stine cites the following passage from Armstrong (1973): “It is not a conclusive objection to a thermometer that it is only reliable in a certain sort of environment. In the same way, reliability of belief, but only within certain sort of environment, would seem to be sufficient for the believer to earn the accolade of knowledge if that sort of environment is part of his boundary-conditions”(Armstrong 1973, 174). See: Stine 1976, 251.

¹⁰³ Stine 1976, 258.

¹⁰⁴ Yourgrau 1983, 185.

¹⁰⁵ Yourgrau 1983, 185. Dretske would happen to agree with Yourgrau that Stine cannot rule out sceptical alternatives, because of their implausibility. On this point, Dretske writes: “But the question here is not whether this

Yourgrau is borrowing the term “self-presenting states” from Roderick Chisholm (1973a) who is an internalist about knowledge, who counts self-presenting states as a basis for knowledge on the ground that such states, by their very nature, remain evident and internal to the putative knower.¹⁰⁶ Granted it is true that knowing the falsity of a sceptical alternative is considerably different from knowing because one is in a self-presenting state; but Yourgrau’s criticism would not stick if Stine makes externalist commitments about knowledge.

A deeper problem, however, is to square Stine’s externalist leanings with her concession to the radical sceptic that sometimes the conversational context is such as to allow for radically tight criteria for “knowledge”-ascriptions – or “knowledge”-denials, as it may be. For an externalist there is no good reason to concede to the radical sceptic that subjects do not know the falsity of radical sceptical alternatives, since all that the sceptic shows is that the subject cannot know from within a first-person perspective that these alternatives do not obtain. From an externalist perspective, this has no tendency, however, to show that subjects are not *in fact* in circumstances favourable to gaining knowledge. Moreover, the closure principle holds in these circumstances, since the proposition I know – “know” in the externalist sense of knowledge – remains closed along with the set of entailments that this proposition generates; and although I may not know all of these entailments as such, I will at least be *in a position* to know them. There is no “equivocation” in making closure-based arguments, if this externalist conception of knowledge is maintained.

I submit that the above incongruity between Stine’s externalist leanings and her concession to the radical sceptic results from her peculiar understanding (or perhaps misunderstanding) of the concept of knowledge as being extremely elastic. Against Stine I have argued that it remains an open question, however, whether the radical sceptic is actually playing on our concept of knowledge or going beyond it. This is so given that Stine is justifying her inclusion of the radical sceptical context as a legitimate context on the basis of an observation that there is such a context and that the concept knowledge is elastic enough to encompass this context without stretching the concept of knowledge beyond recognition. But seeing that there are other social uses that indeed do stretch the concept of knowledge beyond recognition, but yet

alternative is plausible, not whether it is more or less plausible than that there are real zebras in the pen, but whether *you know* that this alternative hypothesis is false”(Dretske, 1016 [original emphasis]).

¹⁰⁶ Chisholm 1973a, 29-34.

are ones that we nevertheless observe as having too loose criteria for the application (or misapplication, as it were) of “knowledge,” should make us wary of judgments made about the compatibility of the radical sceptical context with our concept of knowledge. The bolder suggestion I am making, however, is to include Stine among those ascriber contextualists who do not realize that the application of ascriber contextualism, in order to preserve the closure principle in the face of radical sceptical hypotheses, is simply unnecessary, if, like some ascriber contextualists, Stine begins by making externalist commitments.

In section **1.5** we have seen that Stine preserved the closure principle and diminished the threat of (radical) closure-based scepticism by endorsing ascriber contextualism. According to the contextualist, so long as the sceptic remains quarantined to the sceptical conversational context, then subjects can continue to enjoy knowledge of ordinary empirical propositions in other kinds of context. In the sceptical context, the sceptic can claim that we suffer gross ignorance; and in that context, the sceptic is correct, according to the ascriber contextualist. Yet, as I discussed in section **1.4** and will further discuss in section **1.7**, there is no good reason for the ascriber contextualist to make such a concession to the sceptic, especially if the ascriber contextualist is an externalist. The acceptability of the proposition that subjects do not know the falsity of sceptical alternatives precisely depends on an epistemological commitment to internalism. I will examine this claim further in section **1.7**.

In the following section, **1.6**, the issue of relevance comes up again. Cohen holds that there are two conditions of relevance, but which do not enable the theorist to generally specify when an alternative is relevant or irrelevant. This worries Cohen, since, like Stine, he seeks to both preserve the closure principle and remove the threat of closure-based scepticism. To do so, Cohen endorses ascriber contextualism, which he thinks relieves much of the worry of having to specify when an alternative is relevant or irrelevant, since this will depend on the kind of conversational context in effect at the time of the “knowledge”-ascription or “knowledge”-denial. Yet it is important to point out that Cohen is not simply interested in removing the threat of scepticism; he also desires to explain it. On Cohen’s view, reasons are probabilistic and ultimately fallibilistic in nature, which means there will always be a chance that our beliefs that are based on reasons will be false. In Cohen’s view, the *explanation* of why scepticism seems threatening is that the sceptic underscores the probabilistic nature of reasons whenever that

sceptic raises sceptical alternatives or error-possibilities. Cohen claims that contextualism does justice to the apparent threat that sceptical alternatives pose. Cohen's desire to explain the apparent threat of scepticism will come up again when I examine the sceptical "paradox" in section 1.7. The so-called sceptical "paradox" is just up ahead. But before I examine it, I need to discuss the aspects of Cohen's account that led him to formulate it.

1.6 How to be a Fallibilist

Cohen (1988) takes up the onerous task of explicating the notion of relevance. Eventually he comes to the conclusion that in order to remove the threat of scepticism, the relevant alternatives theorist must endorse ascriber contextualism. With Goldman's (1976) first view of relevance in mind, Cohen claims that there are two kinds of conditions on relevance: external and internal conditions, both of which can be considered Cohen's way of making modifications and refinements to Goldman's first view of relevance.¹⁰⁷ Cohen asserts that the external conditions partially determine relevance, with these conditions embodying the same kind of factors that DeRose (1992) would later call "subject factors," because Cohen's external conditions take account of the particular environmental circumstances of the subject.¹⁰⁸ Cohen contends, however, that many relevant alternatives theorists have overlooked the importance of internal conditions, which evaluate the strength of the epistemic subject's reasons or evidence for the proposition under consideration. On Cohen's view, the internal conditions also play a critical role in determining relevance.

With internal and external conditions in mind, Cohen further interprets Goldman's first view of relevance as suggesting "something like a probability conditional," in which "an alternative is relevant if S's evidence and certain features of the circumstances constitute a reason to believe h [h is an alternative to some ordinary empirical proposition q]."¹⁰⁹ According to Cohen, the probabilistic criterion is best reflected in the external conditions in the following way:

¹⁰⁷ Cohen 1988, 102.

¹⁰⁸ Cohen 1988, 102.

¹⁰⁹ Cohen explicitly makes reference to Goldman (1976). Cohen 1988, 95.

An alternative (to q) h is relevant, if the probability of h conditional on reason r and certain features of the circumstances is sufficiently high (where the level of probability that is sufficient is determined by context).¹¹⁰

But Cohen is keenly aware of the fact that playing this probability game brings up what we might call the problem of cartography: Where are the boundaries to be drawn? To avoid this problem of vagueness, Cohen tries to clarify the notion of “relevance” using *both* internal and external conditions.¹¹¹

Cohen contends that many relevant alternatives theorists, unfortunately, have overlooked internal conditions: “Both Stine and Dretske overlook the fact that S’s evidence against h plays a role in S coming to know q because they think this evidence is not sufficient for S to know not-h.”¹¹² In Cohen’s formulation of the probabilistic criterion, the internal conditions are also affecting the probability of when an alternative becomes relevant by making that probability depend on the subject’s evidence as well as certain features of that subject’s circumstances. Doubtless, those “certain features” in his formulation of the probabilistic criterion, include whatever may be creating a knowledge-defeating environment for the subject, such as – depending on where the subject is and what that subject’s evidence is – fake barns, cleverly disguised mules, or wine bottles containing merely water with red food-colouring in them.

The set of internal conditions is thus encapsulated in the following internal criterion: “An alternative (to q) h is relevant, if S lacks sufficient evidence (reason) to deny h, i.e., to believe not-h.”¹¹³ The question of when evidence is sufficient to rule out an alternative is particularly difficult for Cohen to answer, however, since he espouses a *fallibilistic* view with respect to reasons, according to which “S may know q on the basis of r, even though there is a proposition h, compatible with r but incompatible with q.”¹¹⁴ Consequently, Cohen claims that reasons lack a

¹¹⁰ Cohen 1988, 102.

¹¹¹ Cohen is very sensitive to the “considerable vagueness here” in talking about the external conditions as pointing to a (inductive) probability criterion: For example, how many barn replicas must there be? The specific questions that this general question generates include: How many facades need to be in the immediate area to prevent Henry from knowing that he is looking at a genuine barn? Could facades near to whatever we are calling the immediate area count against Henry knowing that he is looking at a genuine barn? What about if there were facades in the immediate area at one time but there are none there now? The considerable vagueness in the general question seeps into the answers given to it. Contextualists often claim that they can remove this vagueness by talking in terms of the standards that govern the conversational context rather than the circumstances of the subject. See: Cohen 1988, 95.

¹¹² Cohen 1988, 102.

¹¹³ Cohen 1988, 103.

¹¹⁴ Cohen 1988, 93.

normative component and are really statistical and probabilistic in nature, which means that there will always remain a chance that one can be mistaken in what one believes on their basis.¹¹⁵

An alternative view to fallibilism, which would make it very easy to specify when a reason is strong enough to convert a true belief into knowledge and when one has ruled out an alternative, would be to endorse *the entailment principle*, which states that “S knows q on the basis of reason r only if r entails q.”¹¹⁶ But given that reasons or evidence for contingent propositions do not have the property of being factive or entailing, means that endorsing the entailment principle is tantamount to the scepticism, since there will be virtually nothing that we can know on this basis.¹¹⁷ I see a close parallel between Cohen’s thoughts on the nature of reasons and those of Roderick Chisholm’s (1973b), a prominent epistemologist who also rejected the notion of entailing evidence. In Chisholm’s view, “Any adequate theory of evidence must provide for the fact that a proposition e may make evident a proposition h for a subject S even though e does not entail h.”¹¹⁸ Cohen claims that “fallibilistic theories reject this entailment principle thereby avoiding this immediate skeptical result,”¹¹⁹ while Chisholm holds that “we reject the sceptical view according to which there is no reason to believe that the premises of an inductive argument even confer evidence upon the conclusion.”¹²⁰ Cohen’s position on the nature of reasons and evidence seems aligned with Chisholm’s position, especially insofar as rejecting the entailment principle precludes an “immediate sceptical result,” and both positions state that evidence is, at bottom, statistical.¹²¹

But how can a fallibilist who is a RA theorist, such as Cohen, respond to scepticism, if neither the external conditions nor the internal conditions are sufficient to determine relevance? Cohen comes to the ultimately unsatisfactory, but temporary, conclusion that there is no general specification for either of these kinds of conditions, given that specifying the external conditions faces the problem of cartography and playing the probability game, and that the internal

¹¹⁵ Cohen 1988, 108.

¹¹⁶ Cohen 1988, 91.

¹¹⁷ I suppose a reason for the proposition “there are reasons” is an entailing reason, although I doubt very much that this could be an Archimedean point from which to overcome cognitive limitations and defeat scepticism, if one were to endorse the entailment principle.

¹¹⁸ Chisholm 1973b, 232.

¹¹⁹ Cohen 1988, 91.

¹²⁰ Chisholm 1973b, 232.

¹²¹ I find the parallels striking, although Cohen does not acknowledge Chisholm.

conditions reveal that reasons are, at bottom, statistical and ultimately fallibilistic. Cohen makes his verdict in the following passage:

The very same considerations apply to the internal criterion. So there will be no general specification of what constitutes sufficient evidence to deny an alternative in order for it not to be relevant, and as such, no general specification of what constitutes sufficient evidence to know q.¹²²

This conclusion is extremely worrisome to Cohen, because he claims that the goal of the relevant alternatives theory of knowledge is to “rescue fallibilism” from scepticism, the antagonist position which now seems to be virtually invincible, since the sceptic can always raise sceptical alternatives without us having powerful enough reasons to rule them out; namely, because, in Cohen’s view, we simply have no way of specifying what a powerful *enough* reason would be.¹²³ It follows that there is no way to consider sceptical alternatives as irrelevant, because Cohen has not worked out a precise account of relevance. But Cohen cannot be blamed for this, since his discovery is that the central notion of relevance *per se* cannot be defined precisely.

Like Stine (1976), who despite considering sceptical criteria “perverse,” nevertheless permitted them in certain philosophical discussions, Cohen thinks that his account of “knowledge” too leaves some room for sceptical entry. On Cohen’s view, when the sceptic raises an error-possibility, that sceptic essentially underscores the statistical nature of reasons, which is not altogether inappropriate by the sceptic, considering the nature of reasons.¹²⁴ On this important point, Cohen initially claims that there are no independent considerations – in the form of, say, reasons or evidence against the error-possibility – to deny or dismiss the threat of scepticism:

The fact that the denial of the hypotheses is not supported by any independent considerations, underscores our ultimate fallibility. The skeptic calls our attention to the fact that S has no such independent consideration, thereby focusing our attention on the chance of error. This leads us to consider these alternatives relevant.¹²⁵

To diminish the sceptical threat, Cohen makes a key theoretical move, claiming that a specification of what counts as a relevant alternative “...will depend on the context in which the

¹²² Cohen 1988, 103.

¹²³ Cohen 1988, 101.

¹²⁴ Cohen 1988, 108.

¹²⁵ Cohen 1988, 114.

attribution of knowledge occurs.”¹²⁶ The clearest, most positive, statement of Cohen’s theory is the following:

An essential aspect of the version of the theory of relevant alternatives I wish to defend is that the standards that govern relevance are context-sensitive. How probable an alternative must be in order to be relevant will depend on the context in which the knowledge attribution is made. To say simply that skeptical alternatives are not relevant is to fail to do justice to the apparent threat skeptical arguments pose to our knowledge claims.¹²⁷

Accordingly, the sceptic’s alternatives will count as relevant when the sceptic calls attention to the ultimate fallibility of reasons, which makes that alternative salient and shifts the quotidian context to the sceptical context. Essentially, Cohen holds that there are different standards of relevance for the quotidian context from those of the sceptical context. By putting the sceptic’s closure-based argument in conversational context, Cohen removes the threat of scepticism and rescues fallibilism, thereby achieving his goal. Cohen also thinks that ascriber contextualism does justice to the apparent threat of sceptical arguments by revealing that these arguments underscore the fallibilistic nature of reasons. This exposition puts us in a position to now examine the so-called sceptical “paradox.”

1.7 The Sceptical Paradox

The sceptical “paradox” is one of the most widely discussed topics in the literature.¹²⁸ According to contextualists, the supposed sense of “paradox” is pronounced when we consider the sceptic’s argument for closure-based scepticism.¹²⁹ Where “K” stands for “knows that,” “s” for an epistemic subject, “p” for an ordinary empirical proposition, and “q” for an entailment of that proposition in the form of the falsity of a sceptical alternative, the *argument for closure-based scepticism* states:

$$[Ksp \ \& \ Ks(p \rightarrow q)] \rightarrow Ksq$$
$$\sim Ksq$$
$$\sim Ksp$$

¹²⁶ Cohen 1988, 103.

¹²⁷ Cohen 1988, 96.

¹²⁸ Cohen 1986, 1988, 1991, 1999, 2000, 2005, 2008; DeRose 1992, 1995; Feldman 1999; Heller 1999; Neta 2003; Pritchard 2001, 2005; Reiber 1998; Schiffer 1996; Sosa 1999; Vogel 1999 – along with many others.

¹²⁹ Cohen 1999, 67-70.

In words, and making the right qualifications on the closure principle, the above argument states that if S knows that p and S knows that p entails q – or competently deduces q from p – then S knows that q or is in a position to know that q; S does not know that q; therefore, S does not know that p. This argument closely resembles the logically valid pattern of *modus tollens*. Many would accept this argument as valid, but few would consider it sound *without qualification*, because, given that “p” can stand for any contingent proposition, this argument can be generalized to result in a massive sceptical victory. That is to say, if this argument is sound, then this means that the argument is valid and its premises are true, which would deductively necessitate the conclusion that S does not know a given contingent proposition, which is tantamount to S not knowing any contingent propositions. This conclusion has very serious implications for the (now apparently dimming) prospects of knowledge. After laying out this argument, then, the sense of “paradox” becomes most pronounced for ascriber contextualists when we see that they claim that the premises of the argument are acceptable and that the argument is valid. Most contextualists seek to preserve the closure principle and rather freely accept this argument’s second premise, which states that we do not know that sceptical alternatives do not obtain. Given that the premises of the argument for closure-based scepticism are acceptable – according to contextualists – makes it extremely difficult for anyone to answer the sceptic. Cohen (1988, 1991, 1999, 2005, 2008) and DeRose (1995) claim that our extreme sense of difficulty in figuring out what is wrong with the sceptic’s argument for closure-based scepticism, and having our “intuitions” come into conflict when we are facing this argument, reveals that we are dealing with a “paradox.”

Cohen (1988) defines a “paradox” as: “a set of inconsistent propositions all of which have considerable independent plausibility.”¹³⁰ Cohen (1991) nicely lays out the set of individually plausible propositions, or conflicting “intuitions” that issue in the “paradox,” as follows:

¹³⁰Cohen 1988, 94. I have some worries about whether this definition is *broad* enough to be a good definition of a “paradox” (which is part of my reason for putting the word “paradox” in the scare quotes). I think the problem with Cohen’s definition is that it is too *narrow* to cover certain paradigmatic examples of paradoxes, such as the Liar’s paradox. In the Liar’s paradox, the statements are mutually inconsistent. However, each statement seems to lack independent plausibility. For instance, it is unclear how the statement “This next statement is false” could have independent plausibility – rather, it appears to have a definite sort of *dependent* plausibility. That is to say, we can only assess its plausibility when we find out what the next statement is and whether it is indeed false. I suppose for the sake of the example we could *treat* each statement as plausible in the Liar’s Paradox. But *treating* each statement *as* plausible is quite different from each statement *being* independently plausible or *having* independent plausibility.

- (1) We know that some ordinary empirical propositions are true.
- (2) We do not know that skeptical alternatives are false.
- (3) If s knows q, and s knows that q entails not-h, then s knows not-h [in which “s” is an epistemic agent, “q” is an ordinary empirical proposition, and “h” is a sceptical alternative].¹³¹

Before I canvass the different non-contextualist anti-sceptical theories which attempt to explain away the “paradox,” note that Cohen (1988) claims the following as a *desideratum* for any resolution: “...a satisfying resolution requires an explanation of why the paradox arises—an explanation of why we have the intuitions that saddle us with the paradox.”¹³² Moreover, if these propositions indeed express “intuitions,” then a satisfying resolution of the “paradox” will need to keep all of them intact; otherwise, a resolution that gives up one of these propositions will be too costly and simply counterintuitive.

To begin with, no one but the sceptic denies proposition (1). So rejecting proposition (1) is out of the question, if we mean to respond to the sceptic. Moore and the Neo-Mooreans, which include Ernest Sosa (1999) and Duncan Pritchard (2004), deny proposition (2).¹³³ Yet Cohen (1988) and DeRose (1995) argue that Moore and the Neo-Mooreans fail to explain how the “paradox” arises in the first place and how our “intuitions” saddle us with the “paradox.” They do not, then, meet Cohen’s *desideratum* for a satisfying resolution. Furthermore, if proposition (2) indeed expresses an “intuition,” then its rejection, for that very reason, would make such an attempt at a “resolution” counterintuitive. Dretske rejects proposition (3). But contextualists argue that his rejection of closure suffers from a host of problems: it rests on an “equivocation” (Stine 1976) or leads to “abominable conjunctions” (DeRose 1995) or is simply counterintuitive (Cohen 1988, 1999). Most contextualists, along with many non-contextualists, find the closure principle to be thoroughly plausible. Many philosophers find it plausible that knowledge can be extended through deduction.¹³⁴ The closure principle is an expression of this idea. So a rejection of the closure principle may be viewed as implausible for denying that knowledge can be

Cohen’s definition, however, states that each proposition will have independent plausibility. But that definitional clause – it seems to me, anyway – is precisely what makes his definition of a “paradox” too narrow.

¹³¹ Cohen 1991, 20.

¹³² Cohen 1988, 94.

¹³³ See note 51 above for more information on the Neo-Mooreans.

¹³⁴ See: Williamson 2000; also see: Hawthorne 2005.

extended through deduction. But the best reason to resist Dretske's rejection of the closure principle is to realize that he mistakenly conflates the agent's presuppositions with the logical consequences of a statement. The most effective strategy of countering Dretske, then, is to point out his serious mistake in rejecting the closure principle by using his own terms to do so, as I have done. Even if contextualists do not realize this as the most effective strategy of countering Dretske, they nevertheless maintain that it is extremely doubtful that Dretske's rejection of proposition (3) can be the way to escape the "paradox." A possible last non-contextualist option is to claim, as Stephen Schiffer (1996) does, that the concept of knowledge is, on examination, "implicitly incoherent."¹³⁵ The "implicit incoherence" of the concept of knowledge, which cannot be seen on the surface, may meet Cohen's *desideratum* for a satisfying resolution to the sceptical "paradox." But for many Schiffer's claim that the concept of knowledge is incoherent is just as unsatisfying as scepticism is. Moreover, Schiffer's claim for the "implicit incoherence" of knowledge comes from him finding the sceptical "paradox" to be a genuine paradox.¹³⁶ So, Schiffer's claim cannot resolve the sceptical "paradox."

The contextualist resolution to the sceptical "paradox" is to claim that all of the propositions and supposed "intuitions" can be expressed at the same time within different kinds of conversational context, but *never* within the same kind of conversational context.¹³⁷ Contextualists argue that the standards which determine the truth of "knowledge"-ascriptions in the quotidian context are sufficiently different from those standards which determine the truth of "knowledge"-denials in the sceptical context. These standards make the "knowledge"-ascriptions made in the quotidian context *mean* something different from any "knowledge"-ascriptions made in the sceptical context (if they were repeated and made in this kind of conversational context).¹³⁸ According to contextualists, the key to unlocking the sceptical "paradox" is to see that the "knowledge"-ascription does not contradict the "knowledge"-denial when these are issued in different kinds of context.¹³⁹ Accordingly, the following argument which closely resembles *modus ponens* is sound in the quotidian context. Where "K" stands for "knows that," "s" for an

¹³⁵ Schiffer 1996, 333.

¹³⁶ Schiffer 1996, 329.

¹³⁷ Cohen 1999, 65-67.

¹³⁸ See: Cohen 1988, 108; also: DeRose 1992; also: DeRose 1995.

¹³⁹ Cohen 1988, 97.

epistemic subject, “p” for an ordinary empirical proposition, and “q” for an entailment in the form of the falsity of a sceptical alternative, the argument for closure-based knowledge states:

$$[K_{sp} \ \& \ K_s(p \rightarrow q)] \rightarrow K_{sq}$$

K_{sp}

K_{sq} ¹⁴⁰

In this argument, which is the opposite of the sceptic’s argument for closure-based scepticism, this argument’s conclusion, which it comes to on the basis of S knowing a given contingent proposition, such as that the animal is a zebra, states that S can know that a sceptical alternative, such as that the animal is not a cleverly disguised mule, does not obtain. This argument, then, expresses “intuitions” (1) and (3), while it suppresses the “intuition” (2) that we do not know that sceptical alternatives are false. In contrast, in the sceptical context, the contextualist claims that the argument for closure-based scepticism is sound. Consequently, in the sceptical context, that argument expresses “intuitions” (2) and (3), while it suppresses “intuition” (1) that we know that some ordinary empirical propositions are true. According to contextualists, the closure principle (3) holds *relative* to the context in which the argument employing it is made. Everything depends, then, on which kind of conversational context the ascriber or denier is in at the time of making the “knowledge”-ascription or “knowledge”-denial. The contextualist resolution to the sceptical “paradox” apparently manages to keep all of the “intuitions” intact, although these “intuitions” are not all expressed within the same kind of context.

But how does this contextualist resolution meet Cohen’s *desideratum*, which states that a satisfying resolution must not only keep all of the “intuitions” intact but must also explain how the sceptical “paradox” arises in the first place? Cohen’s (1988) explanation is that, “...attributions of knowledge are context-sensitive” and “the apparent closure failures are illusions that result from inattention to contextual shifts.”¹⁴¹ In his explanation, Cohen is contending that the sceptic exploits the fallibilistic nature of knowledge generally by raising error-possibilities which the subject cannot rule out, and that these error-possibilities tend to raise

¹⁴⁰ This argument also needs the same qualifications made on the sceptic’s argument for closure-based scepticism; that is to say, for S to know the conclusion, S must have made a “competent deduction” from the second premise using the closure principle stated in the first premise.

¹⁴¹ Cohen 1988, 111.

the context-sensitive standards of “knowledge” by affecting attributor factors so that the kind of conversational context shifts and changes. Given that the *meaning* of “knowledge”-claims vary according to the conversational context from which they are issued and that speakers can sometimes be inattentive to shifts in contextual standards, contextualists are theoretically committed to the hypothesis that speakers are sometimes blind to the semantics of “knowledge”-ascriptions or “knowledge”-denials.¹⁴² Contextualists often talk about this blindness being caused by the indexicality of the word “knowledge.”¹⁴³ To understand Cohen’s explanation of how the sceptical “paradox” arises, recall that the contextualist resolution states that the closure principle holds relative to the kind of conversational context that the speaker is in at the time of the ascription or denial. So, for instance, if the ascriber’s context shifts from the quotidian context to the sceptical context without that ascriber noticing, then the ascriber’s claims to both know a contingent proposition and the falsity of a sceptical alternative (as an entailment of that contingent proposition) will not go through, because these two claims, as a premise and conclusion of the above argument for closure-based knowledge, are false in the sceptical context. Consequently, this appearance of closure failure comes from being inattentive to contextual shifts. Which conversational mechanisms cause shifts in the context and affect the standards of “knowledge,” however, is a matter of debate among contextualists, with changes in relevance (Cohen 1988), the salience of error-possibilities (Lewis 1979; Cohen 1999) or the rule of sensitivity (DeRose 1995), all being candidates for the mechanisms which cause shifts in the context or affect the standards of “knowledge.” Nevertheless, failing to track changes in the conversational context generally explains how ascribers or sceptics get caught in the “paradox.”

Despite seemingly being able to accommodate propositions (1) and (3) with some ease, proposition (2), we do not know that sceptical alternatives are false, of the sceptical “paradox,” causes tremendous difficulty for ascriber contextualists to accommodate. The main reason for this difficulty is that contextualists do not realize that proposition (2) is plausible and acceptable only if there are internalist demands on knowledge, such as that the subject must be cognitively aware of the reasons or evidence that justify that subject’s belief from a first-person perspective. It should be of no surprise, then, that ascribers will be unwilling to ascribe “knowledge” of the

¹⁴² The next chapter is devoted to semantic issues, including subjects being “semantically blind” to the meaning of context-sensitive terms, the purported indexicality of the word “knows,” and the supposed lack of contradiction between “knowledge”-ascriptions and “knowledge”-denials made in different kinds of conversational context.

¹⁴³ I talk about the indexicality of “knowledge” at length in chapter two.

falsity of sceptical alternatives to subjects if these ascribers also do not distinguish between internalist and externalist demands on knowledge, and often assume that, in facing sceptical challenges, the demands to meet such challenges are distinctively internalist.

I asserted earlier, in the above discussion about Dretske breaking his externalist commitments, that the problem of radical scepticism is only a problem for the internalist, because the problem is posed in internalist terms. For Cohen (1988, 1999) the problem of radical scepticism becomes deep enough as to threaten to make his versions of contextualism theoretically incoherent. The reason why this is so is that Cohen's (1988) contextualist version of the relevant alternatives theory of knowledge is committed to internalism insofar as there remains a requirement for the subject to meet internal conditions on relevance in order for ascriptions of "knowledge" to be properly applied. In his recent version of contextualism, however, Cohen (1999) drops any talk of relevance in favour of talk in terms of strength of evidence, which is staunchly internalist in outlook: "My internalist account construes the strength of one's epistemic position as to a large part determined by the strength of one's reasons or evidence."¹⁴⁴ The deep problem and threat is that Cohen's contextualism commits him to the view that subjects can indeed know the falsity of radical sceptical alternatives when the ascriber ascribes "knowledge" to that subject in the quotidian context. But what kind of reasons or evidence, internalistically construed, can a subject have against sceptical alternatives, especially radical sceptical ones, such as the possibility of being deceived by an all-powerful evil demon? This is a tremendously difficult question for Cohen to answer because of his internalism. Yet his contextualism theoretically commits him to the view that such reasons or evidence can be given against radical sceptical alternatives by subjects in at least the quotidian context, since the closure principle is said to hold relative to quotidian context and some of the entailments within that closed set will be ones that prove the falsity of radical sceptical alternatives.

In my discussion of Dretske, I accepted Vogel's (1990; 1999) claim that internalists can supply some evidence and reasons against the cleverly disguised mule alternative, such as – using the following as a reason or piece of evidence – the well-known fact that public zoos do not orchestrate such elaborate pranks. I immediately pointed out, however, that the same kind of reasons – for instance, the statistical improbability of this event – cannot be given to know the

¹⁴⁴ Cohen 1999, 70.

falsity of *radical* sceptical alternatives, because the radical sceptic uses an arsenal of radical sceptical error-possibilities which cannot be ruled out from within the first-person perspective, so that any comfort for the internalist is at best short-lived. The difference between the sceptical alternative that the animal may be a cleverly disguised mule and the sceptical alternative that one may be being tricked by an all-powerful evil demon at this very instant is primarily a difference in scope.¹⁴⁵ The relevant distinction is to view the first sceptical alternative of the cleverly disguised mule as being local and moderate, since there can be some evidence or reasons of a certain kind that can be supplied to count against that local sceptical alternative, whereas the second alternative of the all-powerful evil demon is global and radical, since there supposedly cannot be any evidence – internalistically construed – against that radical sceptical alternative.¹⁴⁶

Cohen's (1988, 1999) contextualism easily manages to deal with local sceptical alternatives, since he holds that subjects can have enough statistical evidence against the possibility of zookeepers playing pranks to rule out this possibility in the quotidian context.¹⁴⁷ But Cohen admits that radical sceptical alternatives cannot be dealt with in this same manner, since these alternatives hang at a level of possibility far removed from local sceptical alternatives, in which no evidence counts against them.¹⁴⁸ On Cohen's view, the scant evidence that subjects have against local sceptical alternatives is indeed sufficient to allow ascribers to properly ascribe "knowledge" to subjects in the quotidian context, in which lax "knowledge" standards are in effect. But this same kind of treatment cannot be given to radical sceptical alternatives for the very reason that there is no evidence against them. There can be no evaluation of the subject's strength of evidence in relation to the quotidian context, since there is simply no evidence to evaluate. Thus, contextualism seems to fail Cohen's internalist commitments. But Cohen does not despair. He contends that "intrinsically rational" beliefs can be internally justified in the quotidian context when radical sceptical alternatives are under consideration.

¹⁴⁵ Dancy 1985, 8-10

¹⁴⁶ Cohen 1988, 111-115; Cohen 1999, 67-70.

¹⁴⁷ Cohen 1988, 104.

¹⁴⁸ Cohen 1988, 111.

In the absence of evidence or reasons against radical sceptical alternatives, Cohen (1988) appeals to beliefs having the property of being “intrinsically rational” as the saving grace for his internalist contextualism. By beliefs being “intrinsically rational,” Cohen claims:

I am referring to a way in which it can be rational (or reasonable) to believe a proposition without possessing evidence for the belief. We can call beliefs of this sort intrinsically rational. While we may concede to the skeptic that we lack evidence against radical skeptical hypotheses, I do not think we should be willing to concede that it is not rational to deny these hypotheses (believe they are false). If so we can view the denials of these hypotheses as intrinsically rational.¹⁴⁹

Doubtless, we make certain presumptions against the radical sceptic, which speak to the irrationality of radical scepticism. In Cohen’s view, to say the opposite would be irrational: “We think it would be crazy to believe radical skeptical hypotheses. This suggests that we think it is rational to deny such hypotheses.”¹⁵⁰ Yet it is important to put Cohen’s contention in the broader perspective of his theoretical commitments. He is contending that subjects can indeed know the falsity of *radical* sceptical alternatives in the quotidian context, provided that these subjects believe that these alternatives are false because it is “intrinsically rational” to believe them not to obtain. I do not question that supposedly “intrinsically rational” beliefs may be internally justified, although it may be difficult to evaluate that justificatory property, since it likely remains at the less-than-explicit level of presumption. I would imagine that Cohen’s anti-sceptical strategy, however, would only infuriate other internalists who have sought to answer the sceptic more directly and without making presumptions against the sceptic. Cohen is keenly aware that his answer is likely to be seen as question-begging against the sceptic.¹⁵¹ But Cohen counters by accusing the sceptic of begging the question against putative knowers. On this score, he concludes that at best there is a trade off on the issue of question-begging.

Cohen realizes that many will remain dissatisfied with his response to the sceptic. But he reminds us of the greater goal of his project: “To resolve the paradox is not to demonstrate to the skeptic that we know. Rather it is to demonstrate to ourselves that we can claim to know without paradox.”¹⁵² Looking at his project at large, however, what is precisely the problem with Cohen’s internalist “resolution” is that it appears to be completely *ad hoc*, for he refuses to admit

¹⁴⁹ Cohen 1988, 112.

¹⁵⁰ Cohen 1988, 112.

¹⁵¹ Cohen 1988, 113.

¹⁵² Cohen 1988, 113.

that subjects can know the falsity of radical sceptical alternatives by using “intrinsically rational” beliefs to gain knowledge in any other context but the quotidian context. Moreover, it is *only* when radical sceptical alternatives are under examination that Cohen permits “intrinsically rational” beliefs as counting against these alternatives. “Intrinsically rational” beliefs cannot be used to know the falsity of *local or moderate* sceptical alternatives in the quotidian context; rather, reasons or evidence, internalistically construed, are required to know the falsity of these alternatives. Thus, “intrinsic rationality” is not an accredited way of knowing, even by Cohen’s own lights. To my mind, this makes a glaring exception of “intrinsic rationality,” which merely serves Cohen’s internalist commitments and saves his contextualism from turning incoherent in the face of *radical* sceptical alternatives. His internalist contextualist “solution,” which relies on “intrinsically rational” beliefs, is *ad hoc* and simply theory-preserving.¹⁵³

Proposition (2) of the sceptical “paradox” is acceptable only if it is construed in a distinctively internalist manner. A number of critics, including Williams (1991) and Pritchard (2005), argue that radical scepticism is only a problem for internalism.¹⁵⁴ Contextualists, many of whom are externalists, including DeRose (1995) and Lewis (1996), make similar mistakes to those of Dretske (1970), especially by compromising their externalist commitments in thinking that radical sceptical alternatives pose a problem. Pritchard (2005) is abundantly clear on the matter: “Primarily, however, the problem is that it is impossible to be *internalistically justified* in believing the denials of sceptical hypotheses.”¹⁵⁵ Pritchard (2005) concludes:

Since externalist theories do not regard internalist justification as being a necessary ingredient of knowledge, it follows that our in principle lack of internalist justification for our beliefs in the denials of sceptical hypotheses will not immediately translate into a lack of knowledge of these propositions... On the face of it, then, externalist theories of knowledge can resist the closure-based sceptical challenge.¹⁵⁶

I agree with Pritchard’s finding. Unless internalist justification is demanded for our beliefs, then there is no good reason to conclude that we do not know that radical sceptical alternatives are false. Accordingly, proposition (2) of the sceptical “paradox” is unacceptable, or is acceptable only *if* it is restricted to internalism about knowledge; at the very least, proposition (2) does not

¹⁵³ For criticisms along these lines, see: Stanley 2005, 128-129; also: Pritchard 2005, 58-61.

¹⁵⁴ See: Williams 2004, 456-475.

¹⁵⁵ Pritchard 2005, 44 [original emphasis].

¹⁵⁶ Pritchard 2005, 46.

express an “intuition” that we all share. This dissolves the sceptical “paradox.” That the sceptical “paradox” continues to riddle internalists is no surprise, especially if Pritchard is right that it is impossible to be internalistically justified in believing the denials of radical sceptical hypotheses. It continues to riddle ascriber contextualists too, such as Cohen (1988; 1999), because of his internalist commitments. But it riddles even those contextualists who are externalists because they labour under the mistaken internalist assumption that radical sceptical alternatives threaten knowledge and must be ruled out from within a first-person perspective.¹⁵⁷

¹⁵⁷ I examine DeRose’s (1995) externalist contextualism in chapter three, which will prove that he makes this mistaken assumption.

CHAPTER 2

CONTEXT-SENSITIVITY, SEMANTIC BLINDESS, AND INDEXICALITY

2.1 Semantic Blindness and Error Theory

A central task for ascriber contextualists is to explain how the *meaning* of the word “knows” is sensitive to the conversational *context* of the ascriber or denier at the time of the “knowledge”-ascription or “knowledge”-denial.¹⁵⁸ It is not obvious that either the word “knows” or an instance of the schema “S knows that p” is context-sensitive in this manner, although contextualists have argued for at least one of these options.¹⁵⁹ Contextualists are aware that many find their claims for the context-sensitivity of epistemic vocabulary unobvious or perhaps even false. In acknowledging the initial resistance many feel towards the thesis that “knowledge”-ascriptions are context-sensitive, Cohen (1999) observes that, “...those who do accept the thesis, generally do so only as a result of being convinced by philosophical reflection.”¹⁶⁰ This tendency of ordinary speakers to resist or deny the context-sensitivity of epistemic vocabulary could be what ultimately prevents such speakers from appreciating contextualism’s alleged discovery that the word “knows” is context-sensitive. Nevertheless, given that many ordinary speakers have this tendency to resist the contextualist’s thesis, the contextualist’s task of elucidating the context-sensitivity of epistemic vocabulary is made all the more demanding.

As we have seen from an examination of the so-called sceptical “paradox,” contextualists contend that, from the point of view of ordinary speakers ascribing-“knowledge” or denying-“knowledge,” the context-sensitivity of the word “knows” remains hidden from these speakers.¹⁶¹ Consequently, these speakers suffer from what John Hawthorne (2004) aptly terms

¹⁵⁸ The term “context” is highly ambiguous, even after it has been qualified as a “*conversational context*,” as opposed to, say, an extra-conversational context or the environment of the epistemic subject. As Kent Bach (2005a) notes, the word “context” could refer either to the parameters whose values delimit the semantic values of expressions that are context-sensitive with variable contents or to the set of salient shared beliefs or presumptions of the speakers of a particular conversation (Bach 2005a, 15-44). Unfortunately, ascriber contextualists often gloss over this difference by simply talking about the “conversational context.” That contextualism is supposedly a semantic thesis about the *meaning* of the word “knows” would strongly indicate that contextualists are working with the former, more linguistic, understanding of “context,” which is prominent in semantics. However, care should be taken, since, I think, contextualists sometimes dangerously switch between these two rather different senses of the term “context.”

¹⁵⁹ Stanley 2005, 16. Notice that the contextualist who argues that the word “knows” is context-sensitive holds that the schema “S knows that p” is likewise context-sensitive, since the schema contains the word “knows” in it.

¹⁶⁰ Cohen 1999, 78.

¹⁶¹ Specifically, I have in mind, speakers who are unaware of ascriber contextualism as a thesis about the context-sensitivity of the meaning of the word “knows.”

“semantic blindness,” which is being blind to the purported context-sensitive meaning of the word “knows.”¹⁶² That speakers suffer from this semantic blindness reveals that contextualism is committed to what Stephen Schiffer (1996) calls an *error-theory*.¹⁶³ Contextualists are committed to an error-theory because, from a *theoretical* perspective, they claim that ordinary speakers make an *error* when these speakers think that a proposition expressed in a quotidian context conflicts with another proposition expressed in a sceptical (or in another extraordinary) conversational context.¹⁶⁴ According to contextualists, however, there is no actual conflict between the “knowledge”-ascription and “knowledge”-denial when these are made in different kinds of context. That is why contextualists contend that ordinary speakers make an error when these speakers think that there is conflict.

The contextualist is not seeking to correct this error, however. The error-theory is rather a theoretical explanation of how the error arises, which identifies semantic blindness as the cause.¹⁶⁵ Moreover, the explanatory value of semantic blindness is indispensable to the contextualist, because it explains how the sceptical “paradox” arises. For if speakers were cured of their semantic blindness and able to see the context-sensitivity of epistemic vocabulary and hence realize which semantic standards were in effect at the time of making the “knowledge”-

¹⁶² Hawthorne 2004, 107-111. Note that Hawthorne is a critic of contextualism. Contextualists have accepted the aptness of this term; see Cohen 2008; also: DeRose 2006, 2008. In coining the term “semantic blindness,” Hawthorne acknowledges Stephen Schiffer’s (1996) essay, which questions the plausibility of contextualism’s explanation of how competent speakers could fail to appreciate the context sensitivity of epistemic vocabulary.

¹⁶³ Note that Schiffer is another critic of contextualism. To be clear about what Schiffer means by an error-theory, consider his description: “...the claim that people uttering certain knowledge sentences in certain contexts systematically confound the propositions their utterances express with the propositions they would express by uttering those sentences in certain other contexts”(Schiffer 1996, 325).

¹⁶⁴ Contextualists often contrast these two kinds of conversational context. This is mostly because ascriber contextualism has been developed with an eye toward solving the so-called sceptical “paradox,” so that naturally contextualists only discuss these two kinds of context to compare them. This can dangerously mislead readers into thinking that the meaning of the “knowledge”-ascription and “knowledge”-denial are well-defined and that something like ambiguity is what the contextualist is pointing out. But that is not so. There can be intermediate contexts too. Each “knowledge”-ascription can be subtly different from others, a difference which can depend on both the number and kind of error-possibilities raised in the ascriber’s context and how the semantic standards are being set at the time of the “knowledge”-ascription. My point here is that a difference in relevant error-possibilities in different contexts, however slight, is enough to ensure that two different propositions will result and thereby there will be a difference of meaning. See DeRose (1992) for an important discussion of this topic regarding semantic standards, truth-conditions, and how changes in the relevance of error-possibilities changes the very proposition expressed.

¹⁶⁵ I realize that this statement contains an *anachronism*, since Schiffer (1996) first talks about an error-theory and then Hawthorne (2004) comes up with the term “semantic blindness.” Nevertheless, for the purposes of understanding, I find this the most useful way of connecting the two ideas. Making the connection, however, is not at all novel, since, as Hawthorne admits, he came up with this term of “semantic blindness” by considering Schiffer’s criticism and description of contextualism as a certain kind of “error-theory.”

ascription or “knowledge”-denial, then ascribers or deniers would neither mistakenly ascribe “knowledge” to epistemic subjects in sceptical contexts nor would they mistakenly deny “knowledge,” as the sceptic does, to epistemic subjects in quotidian contexts.¹⁶⁶ The generation of the sceptical “paradox,” then, requires this kind of confusion on the part of ascribers and deniers. Thus, without the explanatory roles that semantic blindness and this error-theory play, contextualism would no longer offer an adequate solution to the sceptical “paradox,” because contextualism would not explain how the “paradox” arises in the first place. This theoretically commits the contextualist both to the claim that speakers suffer from semantic blindness and to an error-theory about how speakers sometimes mistakenly make “knowledge”-ascriptions in sceptical contexts and mistakenly make “knowledge”-denials in quotidian contexts. This error-theory is manifest in the ascriber and the sceptic believing – mistakenly believing, according to contextualists – that there is a *real* disagreement between them. From the theoretical perspective of contextualism, there is allegedly no genuine disagreement between them, because the semantic standards of the ascriber, which determine the meaning of the “knowledge”-ascription, do not conflict with the semantic standards of the sceptic, which determine the meaning of the “knowledge”-denial.

To account for the context-sensitivity of epistemic vocabulary, some contextualists, such as Cohen (1988) and DeRose (1992), treat the word “knows” as an *indexical*. On this account, the word “knows” is to be grouped with other context-sensitive words, such as “I,” “here,” and “now,” which are known as indexicals.¹⁶⁷ On a more recent account, Cohen (1999, 2000, 2008) claims that schemas of, “S knows that p” at t1 (relative to a particular standard determined by a conversational context), are like schemas of, “surface X is flat” at t1 (relative to a particular standard determined by a conversational context), in the sense that both schemas exhibit the same *type* of indexicality.¹⁶⁸ Cohen (1999, 2000, 2008) contends that the word “knows” remains

¹⁶⁶ I do not mean to suggest by this statement that contextualists have machinations to perpetuate semantic blindness in ordinary speakers in order to prove their thesis or anything like that. Like all other claims of contextualism, this statement too should be read at the level of theoretical explanation. To relieve any lingering worries about the contextualists’ intentions, however, observe that a proper treatment of a disease first requires its proper diagnosis. Contextualists take themselves to be giving the proper diagnosis of this “semantic disease,” and thus would be allowing for the opportunity of proper treatment.

¹⁶⁷ The contextualist also claims that the word “knows” is to be included in a class of indexicals known as *demonstratives*, which includes as members the words “this” and “that.”

¹⁶⁸ Blauw (2005) discusses and distinguishes these two *types* of indexicality; that is to say, Blauw distinguishes Cohen’s (1988) and DeRose’s (1992) old account from Cohen’s (1999, 2000, 2008) recent account. (Blauw 2005, 131).

context-sensitive, but with indexical standards which determine its correct application; that is to say, the standards are implicit and ultimately refer back to the conversational context of the speaker.

Both of these accounts take up the demanding task of elucidating the context-sensitivity of epistemic vocabulary. Cohen (1988) and DeRose (1992) do this by claiming that, because the word “knows” is an indexical, there is no contradiction between “knowledge”-ascriptions and “knowledge”-denials, and so there is no real disagreement between ascribers and deniers. Cohen (1999, 2000, 2008) does this by claiming that, because the “knowledge”-sentences and “flatness”-sentences take the above schemas and exhibit the same type of indexicality, the contextualist can explain purported cases of semantic blindness, explain away the apparent contradiction between “knowledge”-ascriptions and “knowledge”-denials that have these schemas, and also explain away disagreements between ascribers and deniers who make utterances of “knowledge”-sentences that have these schemas.

2.2 The Argument of This Chapter

What the contextualist owes us is both a *semantic* account of the context-sensitivity of epistemic vocabulary and an *explanation* of how fluent and competent speakers could systematically fail to see the context-sensitivity of epistemic vocabulary. The contextualist’s explanation of semantic blindness should be closely connected to the contextualist’s account of the context-sensitivity of the word “knows” in the epistemic vocabulary, because the contextualist is claiming that something about the semantics of the epistemic vocabulary is the cause of semantic blindness in ordinary speakers. Contextualists contend that indexicality holds the most promising explanation of this context-sensitivity. Ultimately, then, contextualists should present genuine cases of speakers’ being similarly semantically blind to words that are not a part of the epistemic vocabulary; otherwise, contextualists risk their explanation of speakers’ being semantically blind to words in the epistemic vocabulary appearing exceptional and *ad hoc*.

In this chapter, I shall argue that both of these contextualist accounts fail to adequately explain the context-sensitivity of the word “knows” – the most central word in the epistemic vocabulary. Accordingly, I contend that both of these accounts fail to deliver a suitable explanation of the indexicality of the word “knows.” Consequently, I claim that contextualists do not present genuine cases of speakers’ being semantically blind to epistemic vocabulary – or, to

put it differently, these supposedly analogous cases of speakers' being semantically blind to other non-epistemic but nevertheless context-sensitive terms turn out to only have a superficial similarity to the purported cases of speakers' being semantically blind to epistemic vocabulary. Lastly, I will try to meet Cohen's (2008) most recent challenge of having to explain, without using ascriber contextualism, the cases of "flatness"-scepticism that he cites as genuine cases of semantic blindness for non-epistemic terms.

2.3 Is the Word "Knows" an Indexical?

Cohen (1988) and DeRose (1992) treat the word "knows" as an indexical. Paradigmatic examples of indexicals, or of what David Kaplan (1989) more specifically classifies as pure indexicals, include the words "I," "here," and "now."¹⁶⁹ An indexical is a word which depends for its content on features of the (semantic) context in which it is used, so that these features establish an index for the word.¹⁷⁰ An index is whatever completely specifies the content of the indexical, such as, in Kaplan's pure indexicals, the person uttering it, the time at which it is uttered, or the place where it is uttered. When an indexical is used in a sentence, the content of the indexical in that sentence needs to be specified for that sentence to express a proposition.¹⁷¹ That is to say, for sentences containing pure indexicals, we need to know to whom the indexical "I" refers, to where the indexical "here" refers, and to when the indexical "now" refers, in order to give the indexical an index and assign the sentence containing it a truth-value. In short, indexicals need content and indices to express propositions.

According to Kaplan's (1989) treatment of indexicals, an indexical has a context-*variable* content but a context-*invariable* linguistic character.¹⁷² The linguistic *character* of the indexical "I," for instance, is a first-person reflexive pronoun, which invariably refers to whoever utters it.¹⁷³ This is true simply as a matter of the word's linguistic meaning. The *content* of the indexical "I," however, depends on the particular agent who sincerely uses it.¹⁷⁴ This is because

¹⁶⁹ Stanley (2005) points out that "...most contextualists have claimed that 'know' is, like 'I', 'here', and 'now', an indexical expression"(Stanley 2005, 16).

¹⁷⁰ See: DeRose 1992, 921-922; Also: Blaauw 2005, 130; Also: Predelli 2005, 40. Note that in this definition "context," or what I indicated as a *semantic* context, would be the parameters whose values delimit the semantic values of expressions that are context-sensitive with variable contents, see: Bach 2005a.

¹⁷¹ See: Feldman 1999, 73. Also see: Feldman 1999, 83, especially Note 25.

¹⁷² Kaplan 1989, 490.

¹⁷³ Likewise, the character of the word "now" is the time of the context and the character of the word "here" is the place of the context. The character of the indexical linguistically constrains the indexical's content.

¹⁷⁴ Martijn Blaauw 2005, 130.

the indexical “I” has a content that can vary from agent-to-agent. Again, we need to specify the content of the indexical for a sentence containing it to express a proposition and be assigned a truth-value. This would naturally happen when a particular speaker sincerely utters a sentence containing an indexical, because the indexical would then gain its content.¹⁷⁵

Affirmative sentences containing the same type of indexical need not contradict negative sentences also containing the same type of indexical, provided that there is a difference of content or indices for the relevant indexicals. To see this, consider that the sincere utterances of the sentences “I am hungry (at t1)” and “I am not hungry (at t1)” need not contradict one another, because the utterance of the first sentence could be said by one speaker, say, by David Kaplan, having an index such as David Kaplan is hungry at noon on January 1st 2000, whereas the utterance of the second sentence could be said by another speaker, say, by Stewart Cohen, having an index such as Stewart Cohen is not hungry at noon on January 1st 2000. In this case, the content of the indexical “I,” in the utterance of the sentence “I am hungry (at t1),” refers to David Kaplan, and this provides the indexical with an index, whereas the content of the indexical “I,” in the utterance of the sentence “I am not hungry (at t1),” refers to Stewart Cohen, and this provides the same type of indexical with an index distinct from the other one. Therefore, in this example, since David Kaplan’s utterance of the sentence “I am hungry (at t1)” has the index of David Kaplan is hungry at noon on January 1st 2000, whereas Stewart Cohen’s utterance of the sentence “I am not hungry (at t1)” has the index of Stewart Cohen is not hungry at noon on January 1st 2000, which has an index distinct from the other one, there is no contradiction between these utterances of these sentences.

Speakers sometimes – albeit, rarely, in my opinion – fail to appreciate this lack of contradiction between utterances of sentences containing indexicals, especially in situations of the above sort in which affirmative and negative utterances of sentences contain the same type of indexical. This is what contextualists find so attractive about viewing the word “knows” as an indexical. After all, if the word “knows” were an indexical, then this would account for the lack of contradiction between “knowledge”-ascriptions and “knowledge”-denials issued in different

¹⁷⁵ I suppose that in *mentioning* the sentence *I am hungry*, for instance, that there are truths about it – such as that the third word of the mentioned sentence begins with the letter “h.” But I think the point stands that when the sentence is *used*, the sentence’s content needs to be specified for it to be assigned a truth-value.

kinds of context. Moreover, it may enable the contextualist to explain away the apparent disagreement between the ascriber and sceptic, and also may account for semantic blindness.

Cohen (1988) explicitly advertises the attractiveness of this lack of contradiction for utterances of sentences containing indexicals when he declares that, "...the theory I wish to defend construes 'knowledge' as an indexical. As such, one speaker may attribute knowledge to a subject while another speaker denies knowledge to that same subject, without contradiction."¹⁷⁶ Unfortunately, he does not elaborate on how exactly the word "knowledge" is an indexical. Fortunately, DeRose (1992) provides some much needed clarification of Cohen's passage. DeRose clarifies Cohen's passage and elaborates on how the word "knows" is an indexical by way of the following Intercom example:

This lack of contradiction is the key to the sense in which the knowledge attributor and knowledge denier mean something different by 'know.' It is similar to the sense in which two people who think they are in the same room but are in fact in different rooms and are talking to each [other] over an intercom mean something different by 'this room' when one claims, 'Frank is not in this room' and the other insists, 'Frank is in this room—I can see him!' There is an important sense in which both do mean the same thing by 'this room,' in which they are using the phrase in the same sense. But there is also an important sense in which they do not mean the same thing by the phrase; this is the sense by which we can explain the lack of contradiction between what the two people are saying.¹⁷⁷

In DeRose's Intercom example, I take it that the speakers think they are contradicting one another, and there is, therefore, a disagreement between these speakers. DeRose holds that the source of confusion comes from a failure on the part of these speakers to appreciate the context-sensitivity of their utterances of sentences containing the same type of indexical, both speakers thinking that the sentences "Frank is in this room (at t1)" and "Frank is not in this room (at t1)" contradict one another. Like the above example of "I am hungry," however, this is merely the appearance of contradiction, even if the speakers continue to disagree with one another.

As DeRose notes, it is useful to employ Kaplan's terminology and description of indexicals to explain what is happening in the Intercom example.¹⁷⁸ In DeRose's example, the relevant indexical is the word "this," which is a true demonstrative, according to Kaplan's

¹⁷⁶ Cohen 1988, 97.

¹⁷⁷ DeRose 1992, 920-921.

¹⁷⁸ DeRose 1992, 921.

classification of indexicals.¹⁷⁹ Recall that Kaplan draws a distinction between the character and the content of an indexical. When DeRose observes “the important sense in which both do mean the same thing by the phrase ‘this room,’” the shared sense between the speakers of the phrase “this room” is its invariable character or linguistic meaning. Where the meaning of the phrase “this room” diverges is with regard to the context-sensitive meaning of the phrase “this room,” a context-sensitive phrase which is to be explained by a difference in the variable *content* of the indexical “this.” The content of the indexical “this” depends on the context in which it is used and perhaps also on a demonstration by the speaker who makes a demonstration to aid communication.

In the Intercom example, what the first speaker means by the phrase “this room” is the room where she is situated and Frank is absent, which could, with some further specification, be treated as an index of the first instance of the indexical “this,” whereas, what the second speaker means by the phrase “this room” is the room where Frank appears and which can be seen on the intercom’s screen, which could be treated, with some further specification, as a distinct index for the second instance of the indexical “this.” Accordingly, there is no contradiction between the first speaker’s utterance of the sentence “Frank is in this room (at t1)” and the second speaker’s utterance of the sentence “Frank is not in this room (at t1),” since these two propositions differ with respect to their content. Nevertheless, as this example suggests, speakers sometimes – albeit, rarely, in my opinion – fail to appreciate the context-sensitivity of utterances of sentences that contain indexicals, and thereby mistakenly think that their utterances conflict with one another. For contextualists this kind of confusion on the part of speakers is a further attractive feature for treating the word “knows” as an indexical, because this would enable contextualists to explain how speakers could be semantically blind to the context-sensitivity of epistemic vocabulary. Essentially, if the word “knows” were an indexical, then this would rather easily enable contextualists to meet a *desideratum* for a satisfying resolution of the sceptical “paradox,” which is to explain how the “paradox” arises in the first place.

¹⁷⁹ Kaplan (1989) draws a distinction between two different kinds of indexicals, separating them into the categories of “pure indexicals,” which include words such as “I,” “here,” “now,” “tomorrow,” “actual,” and “present,” and of “true demonstratives,” which include words such as “he,” “she,” “him,” “her,” “this,” and “that.” In order to clarify the content of the true demonstrative such as “this,” a demonstration as an “aid in communication,” such as gesturing or pointing, is sometimes required, whereas, “the reference of a pure indexical is secured by the contextual parameters such as the agent, time, and location of the utterance...” (Corazza 2004, 138).

DeRose applies Kaplan's (1989) analysis of indexicals to the word "knows" to allegedly reveal its context-sensitivity. On this point, DeRose writes:

To make use of the character/content distinction, the 'character' of 'S knows that p' is, roughly, that S has a true belief that p and is in a *good enough* epistemic position with respect to p; this remains constant from attribution to attribution.¹⁸⁰

According to DeRose, attributor or contextual factors, such as the error-possibilities that the attributor is considering in the conversational context at the time of the "knowledge"-attribution, "...set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true: They [the attributor or contextual factors] affect *how good an epistemic position the putative knower must be in to count as knowing*."¹⁸¹ Attributor or contextual factors, then, specify the content of the word "knows" and *set* the semantic standards for what the meaning of the word "knows" means. These factors also specify the epistemic position that an epistemic subject or putative knower must be in for the "knowledge"-ascription, in the form of the sentence "S knows that p," to be true. Accordingly, attributor or contextual factors have a rippling-effect, in that they set semantic standards which will determine both the content of indexicals and the truth-values of "knowledge"-sentences.¹⁸² There remains considerable vagueness, however, in the phrase "good enough." Yet this vagueness, according to DeRose, is not pernicious; rather, it is indispensable for DeRose's purposes, since it leaves ample room for the purported variable, context-sensitive, content of the word "knows." On this point, DeRose asks: "But how good is good enough?" to which he answers, "This is what varies with context. What the context fixes in determining the 'content' of a knowledge attribution is how good an epistemic position S must be in to count as knowing that p."¹⁸³ DeRose's phrase "how good an epistemic position S must be in to count as knowing that p" again emphasizes that ascriber contextualism is about how contextual or attributor factors, such as considering or mentioning error-possibilities in the conversational context, *set* the semantic standards for determining the meaning of the purportedly context-sensitive word "knows" and the epistemic position S must be in for the "knowledge"-attribution "S knows that p" to be true.

¹⁸⁰ DeRose 1992, 922 [original emphasis].

¹⁸¹ DeRose 1992, 921 [original emphasis].

¹⁸² I say "allow" because a further consideration of features of the *epistemic subject* may be required. This will be clearer when I discuss a quotation from Bach (2005b) in the next section.

¹⁸³ DeRose 1992, 922.

It is especially important to notice that contextualists, such as DeRose (1992; 1995; 2001; 2004; 2005) and Cohen (1988; 1999; 2000; 2008), are not taking a stance on how epistemic subjects need to *meet* necessary conditions for knowledge. Contextualists are not interested in whether the subject meets epistemic conditions, such as the justification condition or whatever necessary epistemic conditions need to be met for the subject's belief to be non-accidentally connected to truth in order for that subject to have knowledge. Contextualists rather are concerned exclusively with the issue of *setting* semantic standards. Once these standards are set, they will determine the meaning of the sentence "S knows that p." To their credit, contextualists are particularly clear on this point. In his description of how contextualists are concerned with setting standards, Kent Bach (2005b) notes:

Contextualists do make clear that the context they have in mind is not the epistemic context of the subject of the knowledge attribution. Everybody agrees that what it takes for George to know that he has hands or, to put it more accurately, for the sentence 'George knows that he has hands' to be true, can depend on George's epistemic situation. This is a matter not of setting standards but of meeting them.¹⁸⁴

What Bach (2005b) is pointing out when he makes the uncontroversial claim that everybody agrees that the truth of the sentence "George knows that he has hands" can depend on George's epistemic situation is that if, say, George's senses have been unreliable and failing him lately, then, whether the "knowledge"-ascription or the sentence "George knows that he has hands" is true or not will depend on – or better yet, in my opinion, will require taking into consideration – George's rather extraordinary epistemic situation. I presume Bach (2005b) is using the word "epistemic" to emphasize that the meeting of standards is a matter of the epistemic subject's situation rather than the ascriber's conversational context. Bach is correct to employ the word in this manner; although, the word "epistemic" appears somewhat out of place, considering George's extraordinary situation, which may not only prevent him from knowing that he has hands, but may also prevent "knowledge"-ascriptions made to him from being true. This indeed makes Bach's (2005b) point; but calling George's rather extraordinary situation "epistemic" has a rather odd ring to it.¹⁸⁵ Nevertheless, the contextualist is not somebody who disagrees with what Bach has pointed out, but rather claims that attributor or contextual factors play a crucial role in the setting of semantic standards. The semantic standards have to be set before

¹⁸⁴ Bach 2005b, 56.

¹⁸⁵ I am not submitting this as criticism, but rather point it out to prevent any confusion; see Bach's (2005b) excellent essay on ascriber contextualism: Bach 2005b, 51-90.

determining the truth of the sentence, because the sentence “George knows that he has hands” contains the purportedly context-sensitive word “knows,” which supposedly does not have a meaning independent of its content. This brings us back to the main question of whether the word “knows” is an indexical. Not everybody, however, agrees with DeRose’s (1992) claim that the word “knows” is an indexical.¹⁸⁶ DeRose’s claim, then, requires further examination.

As in the Intercom example, in which the first speaker’s utterance of the sentence “Frank is in this room (at t1)” meant something different from the second speaker’s utterance of the sentence “Frank is not in this room (at t1),” DeRose similarly is contending that the utterances of the sentences “S knows that p (at t1)” and “S does not know that p (at t1)” can also have different meanings, since the word “knows,” like the word “this,” is supposed to be an indexical. Presumably, the index of the alleged indexical “knows” would be a specification of the features of the attributor’s conversational context and the semantic standards in force at the time of the “knowledge”-attribution, for, unlike the character of the word “knows,” these vary from attribution-to-attribution. According to contextualists, these contextual features determine the content of the “knowledge”-attribution and set the standards for how good an epistemic position a subject must be in to count as having “knowledge.” Only after specifying these contextual features, then, does an utterance of a sentence containing the word “knows” express a proposition; after all, whether a proposition is expressed or not depends on content and contextualists argue that the word “knows” has a *variable* content. This makes the word “knows” appear to be an indexical.

2.4 Yourgrau’s Objection

DeRose (1992) contends that his account of the indexicality of epistemic vocabulary parries a seemingly powerful and popular objection to contextualism, which was first raised by Yourgrau (1983) in the form of what DeRose (2000) now calls the “now you know it, now you don’t” objection to contextualism.¹⁸⁷ More than simply parrying Yourgrau’s objection, however, DeRose (1992) thinks that his account of the indexicality of the word “knows” exposes the faulty

¹⁸⁶ A short list of philosophers who disagree with DeRose include: Feldman 2001, 2004; Rysiew 2001; Blaauw 2004; Hawthorne 2004; Williamson 2005; Bach 2005b; Stanley 2005; Schiffer 1996.

¹⁸⁷ For a thorough discussion of his defence against the “Now you know it, know you don’t” objection, see DeRose’s paper which goes by the same name: DeRose 2000, 91-106.

assumptions upon which objections of this kind rest. Yourgrau constructs the following dialogue to underscore the seeming incoherence of contextualism:

A: Is that a zebra?

B: Yes, it is a zebra.

A: But can you rule out its being merely a cleverly disguised painted mule?

B: No, I can't.

A: So, you admit you didn't know it was a zebra?

B: No, I did know then that it was a zebra. But after your question, I no longer know.¹⁸⁸

Yourgrau concludes from this dialogue that “something is amiss.”¹⁸⁹ His objection to contextualism is that it permits a kind of flip-flopping of knowledge-claims.¹⁹⁰ From what is happening in this dialogue, one is tempted to agree with Yourgrau that if speaker B claims to know that the animal is a zebra on the one occasion, but subsequently admits to not knowing this claim after an error-possibility has been explicitly raised on another occasion, then speaker B should not have been counted as knowing the original claim, since, given that the particular error-possibility recently raised apparently destroys knowledge, speaker B never had the justification necessary to know that the animal was a zebra in the first place.¹⁹¹ Accordingly, if newly raised error-possibilities could overturn old knowledge-claims, then this would lead to the unpalatable phenomenon of “lost knowledge,” which would not be caused by forgetfulness on the part of the subject, but rather by potent error-possibilities destroying knowledge.¹⁹²

This *appears* to be a problem for the contextualist, because the contextualist is happy to allow for newly introduced error-possibilities raised in the speaker's conversational context to become relevant and affect the semantic standards. The important question to ask, then, is: Does speaker B's “knowledge”-denial at the end of the dialogue conflict with speaker B's original “knowledge”-claim? Yourgrau contends that there is indeed conflict between speaker B's knowledge-denial at the end of the dialogue and speaker B's original knowledge-claim, which seems to reveal that contextualism permits a kind of flip-flopping and remains open to what

¹⁸⁸ Yourgrau 1983, 183.

¹⁸⁹ Yourgrau 1983, 183.

¹⁹⁰ I have omitted the quotation marks from the word “knowledge”-claim here on purpose, because Yourgrau does not seem aware that contextualism is a meta-linguistic thesis which requires meta-linguistic formulations.

¹⁹¹ Yourgrau 1983, 183.

¹⁹² See: DeRose 2000, 92.

DeRose (2000) now calls the “now you know it, now you don’t” objection. This makes contextualism seem incoherent.

DeRose (1992) counters by claiming that Yourgrau fails to appreciate the contextualist’s claim for the context-sensitivity of the word “knows.”¹⁹³ DeRose parries Yourgrau’s objection to contextualism by arguing that speaker B’s “knowledge”-claim was originally made in a conversational context with contextual parameters different from the conversational context in which the current dialogue is taking place, given that a new error-possibility has become salient.¹⁹⁴ By treating the word “knows” as an indexical, DeRose contends that this change in speaker B’s conversational context would be analogous to the following situation: speaker C sincerely claims “I am here” when he is at the office. But when questioned later – say, by his boss over the telephone – after having left the office, speaker C would speak falsely if he were to reassert “I am here,” with the word “here” standing for him being at the office.¹⁹⁵ So, if speaker C were to say “I am not here” when he is not at the office, then this claim would not conflict with his original claim of “I am here” when he was at the office. Contextual factors – using this term loosely – in both cases, then, supposedly determine the content of the indexical under consideration. To return to the above dialogue with this analogy in mind, presumably, since speaker A’s error-possibility has been recently introduced, speaker B’s original “knowledge”-claim was made in a conversational context with different semantic standards, in which “knowledge”-claims in this kind of context are often true. But this is not the whole story, because contextual or attributor factors, such as the newly raised error-possibility by speaker A, affect the semantic standards of the “knowledge”-claim in such a way that if speaker A’s error-possibility becomes relevant, it follows that if speaker B were to assert what appears to be the same “knowledge”-claim in the form of a self-ascription, then this would actually *mean* something different from speaker B’s original “knowledge”-claim or so DeRose claims. To be clearer on this point, if speaker B were to assert the new “knowledge”-sentence – if Speaker B asserts that “I know that the animal is not a cleverly disguised mule” – then this claim would not be the original claim, since speaker B’s new claim, as a self-ascription, would mean that speaker B *qua* epistemic subject would be (or would have been) in a good enough epistemic position to rule out the error-possibility that the animal is a cleverly disguised mule. According to DeRose, speaker

¹⁹³ DeRose 1992, 925.

¹⁹⁴ DeRose 2000, 92.

¹⁹⁵ DeRose 1992, 925.

B's "knowledge"-denial would be of this new "knowledge"-claim – which speaker B does not and did not make – since it would carry the implication that speaker B *qua* epistemic subject is (or was) in a good enough epistemic position to eliminate that potentially relevant error-possibility. Crucially, then, speaker B's current "knowledge"-denial does not conflict with the speaker B's original "knowledge"-claim. The *meanings* of the "knowledge"-denial and the "knowledge"-claim do not conflict. Therefore, DeRose concludes that speaker B is *not* saying anything *contradictory* in asserting both the original "knowledge"-claim and the "knowledge"-denial at the end of the dialogue.

I think DeRose is correct that Yourgrau's objection to contextualism does not go through, because his objection to contextualism is aimed at knowledge at the object-level and not the meta-linguistic thesis that contextualism advances. Essentially, Yourgrau fails to appreciate that contextualism is not about meeting standards, but is about *setting* the standards which determine the context-sensitive meaning of the word "knows" and how contextual and attributor factors play a decisive role in determining what would be a good enough epistemic position for the subject to be in to count as "knowing." Accordingly, Yourgrau fails to appreciate the context-sensitivity of the word "knows" when he makes his objection to contextualism. But DeRose can only answer Yourgrau in this manner if DeRose is right about the claim that the word "knows" is an indexical. DeRose's account hinges on the truth of this claim. I think that DeRose's account, however, is on the verge of collapse. I will give you my reasons for thinking so in the next section.¹⁹⁶

2.5 Problems with the First Account

A lot turns on whether the word "knows" is an indexical. If it is, then the contextualist can accomplish the following: the contextualist can readily explain away seemingly contradictory instances of "knowledge"-ascriptions and "knowledge"-denials and disagreements between speakers who think that their statements result in a contradiction; the contextualist can expose the faulty assumptions made in the popular "now you know it, now you don't" objection, of which Yourgrau's objection to contextualism is an instance; and the contextualist might also be able to explain how speakers sometimes remain semantically blind to epistemic vocabulary. The

¹⁹⁶ In the following section I will sometimes talk generally about "the contextualist" rather than DeRose (1992). This is *not* merely for stylistic reasons. The reason why I am doing so is to be inclusive, since Cohen (1988) may also have had in mind the same type of indexicality as DeRose had in mind.

contextualist claims that the word “knows” is an indexical, because contextual or attributor factors of the speaker’s conversational context set the semantic standards of “knowledge”-ascriptions or “knowledge”-denials. After these factors have been specified, the contextualist claims that the variable content and meaning of the “knowledge”-ascription or “knowledge”-denial becomes fixed. This appears to be similar to the way in which indexicals gain their content, because specifying the context in which the indexical is used fixes that indexical’s content. It is principally on this basis that the contextualist claims that the word “knows” is also an indexical.

The contextualist, however, has presented a weak case for viewing the word “knows” as an indexical. To make my case against the contextualist on this matter, consider first that fluent and competent speakers rarely confuse the reference of pure indexicals or true demonstratives in the way the contextualist claims that they do. Consider again the two speakers in DeRose’s Intercom example. In the Intercom example, the exchange strikes us as comical. The reason for this is not the recognition of some familiar kind of folly that speakers often make, but rather the exchange is comical because the speakers are in the same room and foolishly think that they are disagreeing with one another about Frank’s location, supposedly unaware that the phrase “this room” could refer to different rooms.¹⁹⁷ This kind of folly is rare and unexpected.¹⁹⁸ Moreover, it is easy to imagine a third party in the form of a conciliator intervening in the dispute and being able to resolve the matter by explaining to the two speakers that there is no genuine disagreement between them, because they are simply referring to different rooms. In this case, the speakers, if they were reasonable, would realize their confusion, perhaps shrink with embarrassment, and promptly end the dispute. It would be surprising if things were to go even this far and the intervention were needed, however, because the indexical “this” is a true demonstrative.¹⁹⁹ A true demonstrative sometimes requires a demonstration to aid communication. After the demonstration has been given, it should be clear what the reference is. In the Intercom example, the second speaker actually gives a demonstration. The second speaker, in a moment of exasperation, exclaims “I can see him!” with the word “him” referring to Frank. When the second speaker makes this exclamation, it demonstrates the room to which the speaker is referring. Fluent and competent speakers are often successful at figuring out indexical

¹⁹⁷ Rysiew 2001, 484.

¹⁹⁸ Rysiew 2001, 484.

¹⁹⁹ This is according to Kaplan (1989).

references.²⁰⁰ This points to the conclusion that speakers are not semantically blind to indexicals; rather, they indeed *see* the context-sensitivity of sentences containing indexicals.²⁰¹

There are a number of important disanalogies between the kind of dispute in DeRose's Intercom example and the kind of dispute between the ascriber and sceptic who are disputing whether someone knows, although this latter kind of dispute is the one that DeRose is ultimately trying to explain. Unlike the dispute in the Intercom example, these disputes over whether someone knows are commonplace and do not strike listeners as comical, but rather often are taken quite seriously. It is further difficult to imagine a conciliator being able to smoothly resolve disputes of this nature. The contextualist tries to play this role in the dispute between the ascriber and sceptic, but after the contextualist has explained the context-sensitivity of the word "knows" to the parties involved, it is doubtful that it would have the same effects as in the dispute between the speakers in the Intercom example. That is to say, it is doubtful that the conversational participants disputing whether someone knows would suddenly recover from their purported semantic blindness, see the allegedly different semantic standards in place, and feel that there was no genuine disagreement between them.²⁰²

At this point the contextualist may want to reiterate the claim that speakers remain semantically blind to the context-sensitivity of epistemic vocabulary. But repeating this claim now involves a major risk. For if I am right that speakers are not semantically blind to what pure indexicals and true demonstratives refer to and often recognize a kind of context-sensitivity for sentences containing indexicals, then the contextualist is taking a major risk by treating the alleged indexical "knows" as an exception, as the sole indexical to which speakers remain blind. Instead of being well-theoretically motivated, if the contextualist were to make this move, then it would appear be *ad hoc*. That is to say, making the word "knows" an exception appears to be *ad hoc*, because this move is motivated in such a way as to preserve the crucial explanatory roles of speakers' being both semantically blind to epistemic vocabulary and to the indexicality of the word "knows" play in adequately solving the sceptical "paradox." The contextualist needs to

²⁰⁰ See: Rysiew 2001, 477-514.

²⁰¹ By this statement I do not mean to suggest that ordinary speakers recognize the indexicals contained in the sentences as such. I mean for this to be a theoretical observation. Again, Rysiew (2001) discusses this matter in some detail.

²⁰² Feldman (2001) contends that the disagreement between speakers in disputes about whether someone knows is a genuine one, and that even if speakers waver in their intuitions about whether someone knows this does not in itself reveal the context-sensitivity of what they are disputing (Feldman 2001, 73-74).

present cases of similar semantic blindness for indexicals in order for the contextualist's explanation of the word "knows" as an indexical to be well-theoretically motivated. DeRose argues that his Intercom example is one such case, but it is an odd case – to say the least. The speakers in his example rather foolishly misunderstand one another, even after a demonstration to aid communication has been given by the second speaker. The demand for the contextualist to cite genuine cases of speakers' being semantically blind to indexicals, cases which need to be similar to cases of speakers' being semantically blind to epistemic vocabulary, is a demand that remains unmet. This makes the contextualist's contention that there are genuine cases of speakers' being semantically blind to epistemic vocabulary appear exceptional and *ad hoc*.

Another significant disanalogy between a semantic account of indexicals and the contextualist account, which tries to include the word "knows" as a fellow of familiar indexicals, is that the meanings of pure indexicals or true demonstratives do not simply change with changes in *conversational* contextual factors. As it is used by contextualists, the term "context" always is intended to mean the conversational context of speakers or listeners where "knowledge"-ascriptions or "knowledge"-denials take place. Whenever contextualists speak of "attributor or contextual factors," these are features of the *conversational* context, such as the mentioning or considering of error-possibilities by speakers.²⁰³

It is worth noting that these error-possibilities do not have to be anchored in the world, but can rather be of outlandish radical sceptical error-possibilities, such as the possibility of being tricked by a Cartesian evil demon at this very moment. The contextualist claims that the speaker raising these error-possibilities does not need any evidence for the existence of such a demon or even a reason to plausibly believe that such a being could exist. This makes the error-possibilities rather free-floating, since whether or not these error-possibilities are taken seriously, according to the contextualist, depends simply on what the conversation is like, without regard to what the world is like (or often taken to be like).

According to contextualists, however, these contextual or attributor factors set the semantic standards of "knowledge"-ascriptions. It was mostly on this basis that DeRose (1992) replied to Yourgrau's objection to contextualism. So if the contextualist is to make the analogy between the word "knows" and indexicals work, then there should be semantic standards for

²⁰³ DeRose 1992, 915.

indexicals as well. A natural way to interpret the contextualist on this point is to recognize that indexicals gain their content from the context in which they are used. In this way, *contextual factors*, such as person, place, and time, determine the content of the indexical and the meaning of the particular sentence containing the indexical. But notice that these *contextual factors* that determine the content of indexicals are different from the contextual factors of which the contextualist speaks, since these factors which determine the content of indexicals are *extra-conversational*. To see this, consider, for instance, that when a speaker sincerely utters “I am here,” the content of the indexical “I” will be the particular speaker who uses it and the content of the indexical “here” will be the location of that speaker. This kind of contextual specification for indexicals is of no help to the contextualist who wants variable semantic standards for the word “knows,” since the contextual factors that create the index for indexicals remain outside of the conversational context. But the conversational contextual factors to which the contextualist appeals all remain inside of the conversational context. This is an important asymmetry that the contextualist should not overlook, although too often the contextualist glosses over this important difference by simply talking about “contextual factors.”

2.6 Cohen’s Recent Account

The contextualist account of the word “knows” as an indexical, which may have looked promising as an explanation for the context-sensitivity of epistemic vocabulary, is riddled with so many problems as to make it unsalvageable. To avoid these problems, Cohen (1999; 2000; 2008) offers a new account of indexicality which maintains that the conversational context of the ascriber sets the standards and determines the content of the ascription. However, unlike on Cohen’s (1988) and DeRose’s (1992) old accounts, Cohen’s recent account holds that “...ascriptions of knowledge involve an indexical reference to standards.”²⁰⁴ So whenever a true “knowledge”-ascription is made, there will be a corresponding indexical reference to standards, although from the point of view of ascribers, this indexical reference remains implicit and ascribers may fail to recognize the standards in effect at the time of the ascription.²⁰⁵ Cohen also thinks that the set of conversational contextual factors, or what David Lewis (1979) calls the conversational score, *alone* determines which standard will be in effect for both “knowledge”-

²⁰⁴ Cohen 1999, 61.

²⁰⁵ *Implicit* indexical reference is redundant, but useful stylistically in emphasizing semantic blindness, see: Brian McLaughlin 1991, x-xvii.

ascriptions and their analogues. If so, then conversational contextual factors exclusively set the semantic standards. Moreover, Cohen claims that semantic blindness is a much more pervasive phenomenon than critics of contextualism assume it to be, which, if this were true, would avoid my criticism that the contextualist account of semantic blindness appears to be *ad hoc*.

In the last section, I held that fluent and competent speakers tend to see, with a reasonably high degree of clarity, the context-sensitivity of sentences containing pure indexicals or true demonstratives. Cohen (2008) does not contest this claim, but instead argues that speakers are semantically blind to other terms that are not pure indexicals or true demonstratives, but are terms which nonetheless can crucially provide analogues to “knowledge”-ascriptions and “knowledge”-denials. On this point, Cohen states:

Some philosophers find this semantic blindness thesis implausible. If ‘know’ is context sensitive, competent speakers should be aware of it. What can the contextualist say in defense of the semantic blindness thesis? First of all, we can note that semantic blindness exists for other terms in a natural language.²⁰⁶

In this statement, the terms that Cohen has in mind include: “flat,” “bald,” “rich,” and “happy.”²⁰⁷ Cohen claims that these terms are context-sensitive in the sense that whether a predicate containing one of these terms is true depends on which particular standard is in effect at the time of the ascription, a standard which speakers can supposedly remain blind to altogether, or at the beginning of a conversation, or can lose sight of in the course of a conversation.²⁰⁸

Cohen’s favourite example of a context-sensitive term that may provide an analogue to the term “knows” is the term “flat.” Viewing the term “flat” as a useful term with which to compare the term “knows” is not without precedent. A number of prominent epistemologists, including Unger (1975; 1984), Lewis (1979), and Dretske (1981; 1991), have also done so; although these philosophers have come to different conclusions from their analyses. The leading

²⁰⁶ Cohen 2008, 428.

²⁰⁷ Cohen 2008, 428; Also see: Cohen 1999, 60.

²⁰⁸ With the exception of Cohen (1988), Cohen rarely qualifies what he means by “standards.” It remains clear that meeting a standard makes the sentence “S knows that p” true. Recall that Cohen (1988) spoke about “standards of relevance,” since he then espoused a contextualist version of the relevant alternatives theory. Yet, Cohen (1999, 2000, 2008) abandons the relevant alternatives view, and instead stresses justification or evidence. So, standards could be cast in terms of whatever sets the required strength of evidence or justification and which needs to be met in order for the sentence “S knows that p” to be true. This would be done by evaluating the epistemic subject’s evidence to see whether it is strong enough to face error-possibilities raised in the conversational context of the ascriber. James Pryor (2001) classifies Cohen’s (1999) contextualism as espousing a *strength-of-evidence* view.

similarity between the terms “flat” and “knows,” which has been the focus of much philosophical attention, is that both terms seem to be absolute terms. The term “flat” is absolute in the sense that it is inconsistent to claim that surface Y is flatter than surface X, if surface X is said to be flat. Likewise, the term “knows” is absolute in the sense that it is inconsistent to claim, if Jones knows that p, that Smith can know that p better than Jones knows that p.²⁰⁹

Cohen notes that the predicate “is flat,” when it enters into descriptions of surfaces, can be satisfied *simpliciter* for, say, the true ascription “the table is flat,” or according to varying standards, as it can be for, say, the true ascription in the form of a comparative “the table is flatter than the pavement.” But how do we determine which standard satisfies the predicate *simpliciter* or satisfies it to varying degrees?

2.7 Unger’s Invariantism

Cohen thinks that one option, taken by Unger (1975), is to espouse invariantism, which views the term “flat” as an absolute term with extremely high, determinate and fixed standards. These standards are so high, in fact, that if something is said to be flat then nothing could be flatter than it.²¹⁰ Unger notes that the absolute term “flat” literally *means* that a surface is devoid of *any* bumps or curves. So the term “flat” correctly applies only to those surfaces which are without any bumps.²¹¹ From this, Unger claims that we speak falsely whenever we call roads, tables, or fields, “flat,” because each of these types of surface admit of some bumps, however minute or microscopic. Presumably, the predicate “is flat” is only satisfied *simpliciter* or *per se* for mathematical objects, such as Euclidean planes, because these objects are devoid of bumps and nothing could be flatter than them. Unfortunately, given that a natural language contains a number of familiar absolute terms, such as “flat,” “empty,” and “certain,” which speakers usually do not recognize as such and apply liberally to objects and states of the world, it follows from Unger’s observations and claims that there will be a pandemic of semantic blindness. So whenever speakers use these terms, they will mostly speak falsely. Unger thinks that this is the price we pay for having a language that contains absolute terms. Like any good epistemologist,

²⁰⁹ Cohen (1999) has controversially argued for the gradability of “knowledge”-ascriptions in the sense that “knowledge”-ascriptions can be weaker and stronger. Stanley (2005) criticizes Cohen’s claim for the gradability of “knowledge”-ascriptions by contending that saying that S can “better” know that P than another subject (note the clumsiness of this construction) is undermined by linguistic data, which strongly suggest that “knowledge” cannot be a matter of degrees. See: Stanley 2005, 35-46.

²¹⁰ Unger 1975, 45-91. DeRose (1992) calls Unger’s (1975; 1984) invariantism *sceptical invariantism*.

²¹¹ Unger 1975, 54.

Unger aims to say things which are true, or at least not false. Accordingly, he embraces the high invariant standards that are required for using absolute terms correctly. As a consequence, he accepts the scepticism that follows from this: speakers are massively ignorant and often speak falsely whenever they use absolute terms.²¹² Most epistemologists, including Cohen, find the sceptical consequences of Unger's invariantism unpalatable and ultimately unacceptable.

Cohen cannot accept Unger's claims about the meaning of the absolute term "flat," because accepting them would ensure that the predicate "is flat" cannot be sensitive to conversational contextual factors. On Unger's invariantist view, whether a surface is flat or not depends entirely on whether that surface meets the description of what it *means* for a surface to be flat; namely, for that surface to be without any bumps or curves whatsoever. For Unger, a surface X either is flat or is non-flat and bumpy. To find a model for the contextualist treatment of the word "knows," however, Cohen needs the predicate "is flat" to vary with features of the conversational context of the speaker so that there could be cases in which one speaker could say of surface X that "surface X is flat," while another speaker could say of the same surface X that "surface X is not flat," with both the "flatness"-ascription and "flatness"-denial being mutually compatible and true at the same time. Unger's view does not allow for this, since the meaning of the term "flat" is invariant and absolute. Accordingly, the first speaker would flatly contradict the second speaker in the above case – pun intended.

2.8 The Conversational Score

Instead of taking this first option and endorsing Unger's invariantism, Cohen contends that the better option is to adopt Lewis's (1979) account of context-sensitive terms, an account which treats the term "flat" as context-sensitive and avoids the inevitable scepticism that results from invariantism. In considering this option, Cohen asks "How precisely do the standards for these predicates [predicates containing context-sensitive terms] get determined in a particular context of ascription?" and he answers:

This is a very difficult question to answer. But we can say this much. The standards are determined by some complicated function of speaker intentions, listener expectations,

²¹² In order to avoid confusion, Unger is only considering the word "flat" as it applies to surfaces and is not considering figurative uses of the word, as in "My beer went 'flat'" or "I feel 'flat'".

presuppositions of the conversation, salience relations, etc.—by what David Lewis calls the conversational score.²¹³

But what exactly does Lewis mean by the conversational score?

According to Lewis (1979), what happens in a conversational context or particular language game is like – to use his leading analogy – keeping score in a baseball game.²¹⁴ Lewis considers a language game to be like a baseball game in the following respects: components of the score are abstract entities; what play is correct depends on score; score evolves in a rule-governed way; efforts by participants are often cooperative; and the history of a particular game influences score.²¹⁵ In a baseball game, the score is determined by the play on the field along with constitutive rules about what it means to score in the game, constitutive rules which define what counts for, say, scoring a homerun, a walk, an out, and so on.²¹⁶ Analogously, Lewis claims that conversational elements create a context in which speakers understand why the conversation is taking place, are able to keep track of what has happened in the course of the conversation, and can anticipate what might happen next. Lewis claims that speakers can figure out what counts for what; essentially, speakers keep track of what the conversational score is.

To make the idea of conversational score clearer, consider a main example that Lewis uses to demonstrate how *salience relations* can change the conversational score. Lewis presents an example in which he asks you – the reader – to imagine that you are in a room with both him and Bruce, Lewis’s pet cat that keeps dashing about wildly. Bruce is the *only* cat in the room when Lewis says the following to you:

²¹³ Cohen 1999, 61.

²¹⁴ Lewis 1979, 342.

²¹⁵ Lewis 1979, 344-346.

²¹⁶ A constitutive rule is a rule that creates or defines a new form of behaviour and often is expressed by saying what counts for what, such as what counts for a checkmate in chess. A constitutive rule is to be distinguished from a *regulative* rule, which regulates an already existing form of behaviour and often is expressed in the form of an imperative or command, such as “Get your elbows off of the table!” For a discussion of the differences between constitutive and regulative rules, see: Searle 1969, 33-42. For Lewis’s brief discussion of constitutive rules, see Lewis 1979, 343-344. I want to briefly note that Lewis also considers an alternative account to constitutive rules for scorekeeping, which involves taking a legal realist or operationalist view of scorekeeping, which holds that the score is whatever the scoreboard reads. Lewis immediately raises some difficulties for this alternative, however, such as having to answer the question: what is supposed to count as the scoreboard? Is it the board filled with bright lights for all to see, the sweat-stained score sheets in the hands of coaches and spectators, or the score in the umpire’s head? All of these cannot count, because there could be discrepancies between them, which may need to be resolved. Nevertheless, constitutive rules are needed to say what play counts for what.

The cat is in the carton. *The cat* will never meet our other cat, because our other cat lives in New Zealand. Our New Zealand cat lives with the Cresswells. And there he'll stay, because Miriam would be sad if *the cat* went away.²¹⁷

According to Lewis, the conversational participants of this particular conversation – you, the listener, and Lewis, the speaker – are able to figure out which cat has become salient, despite the different cats being referred to by what appear to be identical instances of the definite description “the cat.” Conversational participants usually understand that these seemingly identical instances of this description of “the cat” have different denotations. In the above quotation, the first and second instances of “the cat” denote Bruce, whereas, after Lewis speaks of his other cat in New Zealand, the New Zealand cat becomes more salient than Bruce does to the conversation, so that the third instance of “the cat” denotes the New Zealand cat. When Lewis starts to talk about his New Zealand cat, the conversational score is changing so that listeners will count the next instance of “the cat” as the New Zealand cat. The salience relations have changed. In this way, conversational participants are able to keep score in the conversation and know what counts for what.

But, aside from salience relations, what conversational rules have allowed Lewis to so smoothly move from talking about Bruce to talking about his New Zealand cat? Lewis's answer is that conversations exhibit *rules of accommodation*, which allow for swift conversational transitions.²¹⁸ That rules of accommodation can factor into the conversational score is what Lewis considers the greatest difference between scorekeeping in a conversation and scorekeeping in a baseball game. A conversation is loosely structured so as to permit an array of different moves and ways to score, whereas a baseball game is tightly structured in such a way as to enable its participants to recognize, with a much greater degree of precision than in a conversation, what definitely counts as correct play and a score. As Lewis sees it, “...conversational score does tend to evolve in such a way as is required in order to make whatever occurs count as correct play. Granted, that is not invariable but only a tendency.”²¹⁹ As for correct play in the above conversation about Lewis's pet cats, introducing and describing the New Zealand cat as “the cat,” after that description had already been used to denote Bruce, was easily accommodated by the conversational participants, because of salience relations. According

²¹⁷ Lewis 1979, 348 [my emphasis].

²¹⁸ Lewis 1979, 346-350.

²¹⁹ Lewis 1979, 347.

to Lewis, this made the move correct play. Correct play in a baseball game, however, is much more well-defined. For instance, if I am the batter and the count is three balls and no strikes and I suddenly decide to waltz over to first base, I will not thereby score a walk. This move breaks the rules of baseball and cannot be normally accommodated by the participants – by the players, coaches, umpires, and perhaps also by the spectators – of the game. Unlike a baseball score, the conversational score often adjusts to the conversational moves of the participants, which is because of the rules of accommodation.

But how does Lewis's idea of the conversational score serve Cohen's purposes? It may do so in the following way: Lewis claims that what he calls *standards of precision*, which permit an utterance to be treated as *true enough* or acceptable for the sake of conversation, are also included in the rules of accommodation and thereby in the conversational score as well.²²⁰ Lewis argues that for some context-sensitive and vague terms, such as the terms "bald" or "rich," conversational participants will treat utterances containing these terms as true enough according to the standards of precision in force at the time of the utterance.²²¹ The standards of precision, however, are also vague; but Lewis does not consider this a problem. He thinks that speakers will figure out the strictness or laxity of the standards of precision by testing whether listeners will accept what has been uttered as true enough. For instance, imagine speaker D says "France is hexagonal." If listener E does not object to what speaker D said, then listener E will probably also treat as true enough the utterance of "Italy is boot-shaped."²²² Speaker D could do further tests. But if there is no protest from listener E after a few tries, then the standards of precision are likely to be quite lax and quotidian. Accordingly, the conversational participants are able to find out the conversational score from these tests. But, as sometimes happens, a conversational move may cause the standards of precision to soar to extreme heights, so that not only is the utterance of "France is hexagonal" vehemently protested, but, in "Ungerian" style, the utterance of "there are material objects that are hexagonal" could also be vehemently protested. In this case, if the speaker making these utterances accepts the ultra-precise standards favoured by the other conversational participants, then speakers and listeners may accommodate these standards.

²²⁰ Lewis 1979, 352.

²²¹ Lewis 1979, 352. Notice that, for these utterances (sometimes both these senses are being used) I have dropped the "of a sentence" part in "utterance of a sentence," since I have some worries about Lewis's semantic considerations or lack thereof for the conversational score. I have done so in order not to mislead my reader.

²²² Lewis 1979, 352.

Interestingly, Lewis thinks that Unger's observation of the meaning of the term "flat," as an absolute term, is the correct one: "Some might dispute Unger's premise that "flat" is an absolute term; but on that score it seems to me that Unger is right."²²³ In parallel fashion to the example about the approximate shape of different countries, Lewis holds that Unger is correct in the context of a conversation, provided that ultra-precise standards are in force at the time when the speaker is denying that a surface is "flat." The conversational score sometimes allows for extreme "flatness"-denials. According to Lewis, this will not, however, affect the truth of previous utterances containing the term "flat" made under less precise standards. These previous utterances remain true enough. As Lewis states:

The right response to Unger, I suggest, is that he is changing the score on you. When he says that the desk is flatter than the pavement, what he says is acceptable only under raised standards of precision. Under the original standards the bumps on the pavement were too small to be relevant either to the question whether the pavement is flat or to the question whether the pavement is flatter than the desk. Since what he says requires raised standards, the standards accommodatingly rise. Then it is no longer true enough that the pavement is flat. That does not alter the fact that it *was* true enough *in its original context*.²²⁴

Lewis concludes that the ultimate "arbiter" of correctness is the standard of precision in force at the time of the ascription or denial. This remains the case even after the standards of precision have changed, and thereby the conversational score as well has changed.

2.9 Problems with Cohen's Recent Account

It is unclear whether Cohen wants to import into his recent account of indexicality all of the elements that Lewis makes a part of the conversational score. Cohen often emphasizes the importance that salience relations play in determining standards. So much is clear. As Cohen (1999) states: "In the case of knowledge ascriptions, salience relations play a central role in determining the standards. In particular, when the chance of error is salient, it can lead knowledge ascribers to intend, expect, presuppose, etc., stricter standards."²²⁵ As we have seen, conversational participants can figure out the conversational score by paying attention to salience relations, as in the example of Lewis's pet cats. Cohen remains silent, however, on whether he wants to incorporate the other parts of the conversational score, such as the rules of

²²³ Lewis 1979, 353.

²²⁴ Lewis 1979, 353 [original emphasis].

²²⁵ Cohen 1999, 61.

accommodation generally, into his account. A warranted interpretation from what Cohen says is that, along with salience relations, he also wants to take Lewis's description of the dynamics of the standards of precision, of how the standards tend to become raised. Cohen's account leans quite heavily on Lewis's explanation of the conversational score. Cohen is so eager to explain semantic blindness, however, that he rather optimistically overlooks some of the features of Lewis's explanation of the conversational score that will present theoretical difficulties for his own account. In what immediately follows, I will reveal how these features in Lewis's explanation can cause problems for Cohen's recent account of indexicality.

Lewis's description of how the standards of precision shift in a conversation may at first appear to help Cohen to explain how ascribers and deniers can sometimes remain semantically blind to the standards in force at the time of the ascription or denial. As Lewis claims, "...a player of language games who is so inclined may *get away with it* if he tries to raise the standards of precision as high as possible..."²²⁶ As I understand Lewis on this point, the uncooperative or Sneaky Player – if I may call him that – trying to “get away with it” is trying to raise the standards so that the other players will not see this change, and hence these other players will not be able to protest the Sneaky Player's move for the simple reason that they never saw it. This may not be cheating or breaking the rules of language games, but it is not exactly an innocent move either. That aside, often the conversational participants will be aware of the rules of accommodation, because they tacitly, if not explicitly, agree to them, as is evident from speakers conducting tests to see which standards are in effect and figure out what the conversational score is. But perhaps the other players will sometimes unwittingly let the Sneaky Player “get away with it,” and the standards will rise up accordingly. So much looks promising for Cohen's recent account.

It is doubtful that Cohen, however, can develop out of Lewis's explanation of the conversational score a suitable account of semantic blindness. One problem is that, although conversational participants can lose sight of the standards of precision during the course of a conversation, on Lewis's explanation of the conversational score, there may remain a real possibility of genuine disagreement between these participants. An ascriber could genuinely disagree with a denier if there is *the* (one and only) conversational score, and it reads in favour of

²²⁶ Lewis 1979, 353 [my emphasis].

one of the conversational participants at the time of the ascription or denial.²²⁷ The account that Cohen needs to provide, however, must hold that the “knowledge”-ascription does not contradict the “knowledge”-denial at *the same time*.²²⁸ One piece of textual evidence that appears to challenge my reading is that Lewis clearly concedes to Unger that the ultra-precise standards at the time of the “flatness”-denial can be correct, but maintains that these standards do *not* conflict with the correctness of the less strict, rather imprecise, standards that were in force for the previous “flatness”-ascriptions in the original context. This seems to allow for a lack of contradiction. This is of no real help to Cohen, however. For notice that Lewis is giving a *diachronic* explanation, in the sense that he is explaining how the standards have mutated over time, whereas what Cohen needs for his recent account to work is a *synchronic* explanation, since contextualism states that both the ascriber and the denier can be correct at the *same time*, without contradiction.²²⁹ Therefore, Cohen cannot simply use Lewis’s diachronic explanation, because Cohen needs a synchronic explanation to match his recent account of the indexicality of standards for the word “knows.”

To elaborate on the possibility of a genuine dispute between ascribers and deniers, it appears that the Sneaky Player who “got away with it” *wins* if the scoreboard reads in his favour. Presumably, in similar circumstances, the uncooperative sceptic, as a conversational participant, would also win if the conversation were about “knowledge.” But a principal attractive feature of contextualism, if not a theoretical necessity for its viability, is its *even-handedness* and ability to

²²⁷Lewis is not always consistent about labelling, and switches between “*the* conversational score” and “conversational score” without the word “the.” Lewis talks about “*the* conversational score” in his section titled “conversational score”(Lewis 1979, 345). Note that Cohen talks about Lewis’s notion of “*the* conversational score” in the quotation I provided (Cohen 1999, 61). There is, however, a small literature of Lewisian scholarship that asks questions such as whether a *multiple* scoreboard view could be given, or a *gap*-view with a single scoreboard (DeRose 2004) or another single scoreboard view with a conversational dynamic such that the conversational participants are not genuinely disagreeing. Unfortunately, I do not have the space to discuss these scholarly issues, see: DeRose 2004, 1-21; and Feldman 2004, 23-33 for his comments on DeRose 2004. As this relates to Cohen (1999), I will say this much; the question of what the conversational score is and whether there can be a single scoreboard view is a vexed question. So Cohen cannot remain silent on this question and these issues. He must clarify his position. Cohen cannot simply import whatever this feature of Lewis’s view is, whether it is the one and only scoreboard or many scoreboards, since it is unclear what that very feature is that Cohen is importing.

²²⁸ I realize that this formulation may look awkward, but recall that the “knowledge”-ascription does not contradict the “knowledge”-denial, so long as the word “knowledge” remains in quotations, since this signifies that the formulation is *meta-linguistic* or that “knowledge” is supposed to be indexical.

²²⁹ Recall Lewis’s emphasis in the last quotation, which reveals that he is giving a diachronic explanation: “That does not alter the fact that it *was* true enough *in its original context*.” Recall the above quotation from Cohen (1988) in which he advertised the attractiveness of treating the word “knowledge” as an indexical as allowing for “knowledge”-ascription and “knowledge”-denials to apply to the same subject and be made simultaneously, without contradiction. Cohen retains these features of indexicality, but modifies the type of indexicality.

explain away disagreement, which is supposed to reveal that there is no winner or loser, and ultimately no real opposition between the sceptic and ascriber. Both participants can supposedly be simultaneously making “knowledge”-ascriptions and “knowledge”-denials according to non-conflicting standards.²³⁰ The indexicality of “knowledge”-ascriptions and “knowledge”-denials – or the indexicality of “flatness”-ascriptions and “flatness”-denials, for that matter – is supposed to be what shows that speakers are really at “cross-purposes” or “talking past one another.” That the conversational participants can lose sight of the conversational score may be a form of semantic blindness – although something more akin to “conversational blindness” may be a better term for this – but it is not an explanation that will explain away the disagreement between ascribers and deniers. It is not, therefore, an explanation that will serve Cohen’s purposes.

What I have said so far presents theoretical difficulties for Cohen’s recent account. But there may be a deeper, more fundamental, theoretical problem that Cohen must confront. Much of what I have said in criticizing Cohen hints at a toxic ambiguity that is lurking in the background, which, without Cohen removing it, threatens to poison his recent account. What is threatening is that the term “context” is ambiguous. As Bach (2005a) notes, the term “context” could either refer to those parameters whose values delimit the semantic values of context-sensitive terms with variable contents or to the shared set of salient beliefs that speakers presume or presuppose in a particular conversation.²³¹ Which sense of the term is Cohen using? This is a problem for Cohen, because he does not remove the ambiguity. This problem moves to the forefront, since it appears that Cohen is using the first sense of “context” when he discusses the semantics and indexicality of terms such as “knowledge” and “flat,” but Lewis seems to be using the second sense. Conflating the two senses, however, really confuses the two, because they have distinct features; roughly, the first sense being more akin to semantics, whereas the second is more akin to pragmatics. But without getting into thorny issues of classification or the specific

²³⁰ See Crispin Wright’s (2005) excellent essay for a discussion of how contextualists view even-handedness as one of the main virtues of their position. Wright is a critic of contextualism and ultimately exposes how this even-handedness quickly evaporates, because the contextualist must accept the *factivity* of knowledge; that is to say, that knowledge implies truth or a *fact* about the world. But Wright is keenly aware of the contextualist’s move to semantic ascent and meta-linguistic claims. Wright thinks that these meta-linguistic formulations of “knowledge,” however, *implicate* (conversational) that something factual is presupposed, which is hardly an even-handed way of dealing with the sceptic who wants *everything* to hang at the level of possibility, that it is possible that I am, say, being tricked by an all-powerful evil demon at this moment. But the factivity of “knowledge” suggests more than possibility; it suggests something *factual* about the world. The sceptic is bound to charge the contextualist with being unfair or not being *completely* even-handed, and it looks like the sceptic has good grounds for this charge. (Wright 2005, 236-262).

²³¹ Bach 2005a, 21.

issue of whether indexicality is to be located in semantics or pragmatics or whether it lies nearer to the border – if there is a well-defined, non-fuzzy border or if there is any border at all – we may thankfully avoid these issues by noting that Cohen takes himself to be considering the *semantics* of indexicality and epistemic vocabulary.²³² This would strongly suggest that the sense of “context” that Cohen has in mind is closer to the first sense, especially when he tries to find a suitable account of semantic blindness in Lewis’s explanation of the conversational score. Yet Cohen cannot develop such an account out of Lewis’s explanation of the conversational score, if Lewis is using “context” in the second sense of the set of shared salient beliefs presumed and presupposed by conversational participants – a description of the second sense of the term “context” which looks very similar to Lewis’s own description of the conversational score.²³³

Cohen cannot take the second option of elucidating the standards for “flatness”-ascriptions or “knowledge”-ascriptions by using Lewis’s idea of the conversational score. It is further difficult to see how Lewis’s conversational participants could be semantically blind, although perhaps they could temporarily suffer from something more akin to “conversational blindness.” But even if they were to suffer from semantic blindness, there would still need to be a lack of genuine disagreement between the ascriber and denier and a form of evenhandedness exemplified by the contextualist. Yet, given that Lewis provides a diachronic explanation of a lack of contradiction, it may be that conversational participants can genuinely disagree with one another if the conversational score reads in their favour. Many of the same problems with the old account of the word “knows” as an indexical return and remain unresolved in Cohen’s recent account of indexicality.

²³² Cohen 1999, 60-61.

²³³ It may appear that I have been too hasty in my final judgment, because it may look like I have overlooked what is sometimes called in pragmatics *speaker-meaning* from perhaps playing a significant role in determining the meaning of “knowledge”-ascriptions. It may appear that the speaker who raises certain error-possibilities in the conversational context (perhaps in the second sense) which affect the semantic standards of “knowledge”-ascriptions could also do so with speaker-meaning. Although speaker-meaning usually is understood as speakers trying to *inform* listeners, I suppose it could also be used to *misinform* in the way that, say, the Sneaky Player may have done to the other conversational participants – although this is a stretch of Lewis’s description of that player, since he tries to “get away with it” without notice. Still, I think we can safely assume that the other *non-sneaky* players aim for knowledge and the truth, and also to inform and not be devious. This makes it especially hard to make sense of how the sceptic could have the conversational “wool” pulled over his eyes, so to speak. But it must make sense for Cohen’s account to explain semantic blindness. I do not really see this as a viable option, especially given that Cohen has “semantical considerations.” And, of course, if we are going to talk strictly about *speaker-meaning*, then the speaker cannot be semantically blind to such meaning since such meaning depends on what she means, is transparent to her, and ultimately depends on her intentions. Accordingly, there cannot be semantic blindness to such meaning, since this would make the very notion of speaker-meaning incoherent.

2.10 In the Grip of “Flatness”-Scepticism

But where does this negative conclusion leave us? Does it mean that we should accept Unger’s observation that the meaning of the term “flat” is simply absolute, after all? But if so, then we would have to also accept the rather unpalatable sceptical conclusion that follows from this. I think that Cohen is right to reject Unger’s conclusion, but my reasons for being opposed to Unger are not ascriber contextualist reasons. The main point of my raising these questions, however, is to acknowledge that Cohen may have some of the facts or data right about how speakers behave when these speakers consider whether or not a surface is flat. Yet it is my contention that even if Cohen is right about some of these facts or data, he provides the wrong theoretical explanation of them.

The following is Cohen’s (2008) characterization of how a speaker can go from being a confident ascriber of “flatness”-ascriptions to quickly being overtaken by Ungerian considerations and engulfed by scepticism:

But when we are in the grip of ‘flatness skepticism,’ we are inclined to deny the truth of previous utterances of the form ‘X is flat.’ That is, we mistakenly think that our utterance ‘Nothing is flat’ conflicts with our (and others’) previous utterances of the form ‘X is flat.’²³⁴

Cohen continues by providing a theoretical explanation to cover both phenomena of “flatness”-scepticism and “knowledge”-scepticism:

This is exactly the same kind of mistake the contextualist attributes to us regarding ‘knows.’ So it appears that competent speakers can be blind to the context sensitivity of a term in their language. This lends credibility to the contextualist’s claim that competent speakers are semantically blind with respect to ‘knows.’²³⁵

I have supplied reasons to challenge Cohen’s main claim that cases of speakers’ being semantically blind to the indexicality of “flatness”-ascriptions are analogous to cases of speakers’ being semantically blind to terms in the epistemic vocabulary. To his credit, Cohen immediately acknowledges a relevant disanalogy between these two supposed forms of semantic blindness: “There is, however, a relevant disanalogy between ‘know’ and ‘flat.’ Contextualist theories of flatness ascriptions gain easy and widespread acceptance among people, but

²³⁴ Cohen 2008, 428.

²³⁵ Cohen 2008, 428.

contextualist theories of knowledge do not.”²³⁶ I am not so sure about Cohen’s empirical claim that contextualist theories of “flatness”-ascriptions “gain easy and widespread acceptance among people.”

In response, I will use Dretkse’s (1981; 1991) relevant alternatives theory to provide an alternative explanation of Cohen’s claim about the widespread acceptance of contextualist theories of “flatness”-ascriptions in the next section. For now note that there is a deeper sense in which these alleged forms of semantic blindness are disanalogous, even by the contextualist’s own lights, but which Cohen fails to fully report in his reporting of the above disanalogy: in considerations of “knowledge,” the supposed “intuition” that generates the so-called sceptical “paradox” is that speakers rather freely admit to not knowing the falsity of sceptical error-possibilities. In considerations of “flatness,” however, speakers only grudgingly accept “flatness”-scepticism – after all, they have to be in its grip – and find such sceptical error-possibilities counterintuitive and Unger’s general sceptical conclusion outlandish. Cases of “flatness”-scepticism, however, are supposed to be a very similar to cases of “knowledge”-scepticism, since Cohen ultimately wants to use the former cases to explain the latter cases.

Unlike the contextualist, however, I really do not wish to have this discussion to be in terms of intuitions. As Bach (2005b) claims, talk of intuitions is a risky business for philosophers to engage in for a number of reasons, not the least of which is that we may be simply wrong about our own alleged intuitions or the alleged intuitions of others.²³⁷ Furthermore, *philosophical* intuitions are thought to *constrain* theories and used to *judge* theories as plausible or implausible, whereas facts or data, some of which I am taking Cohen to be presenting in his purported cases of “flatness”-scepticism, are to be *explained by and incorporated into theories*.²³⁸

I will take Cohen’s description of the behavior of speakers in the so-called grip of “flatness”-scepticism as containing some data that should be explained theoretically, although I have qualms and reservations about Cohen’s description, since I think that something more akin to speaker confusion or a loss of confidence is happening. Unlike Cohen, I do not think that speakers in the so-called grip of “flatness”-scepticism are inclined to *deny* their previous claims,

²³⁶ Cohen 2008, 429.

²³⁷ Bach 2005b, 62-63. For a challenge to the value of intuitions in philosophy generally, see: Nichols, Stich, and Weinberg (2003).

²³⁸ Bach 2005b, 63.

but rather, in these situations, speakers will simply retract previous claims because of a loss of confidence. To be more specific, Cohen views ascribers who suddenly find themselves in the so-called grip of “flatness”-scepticism to be mistakenly thinking that there is a conflict between their previous “flatness”-ascription and their current “flatness”-denial. Doubtless, Cohen needs to say this to explain how the ascription and denial have different meanings and are really non-contradictory. In the case of knowledge, however, I think it would be better to view the speaker as failing to meet the doxastic or belief condition, which is because of a lack of confidence. So, the speaker may feel some conflict, yet not of the kind that comes from the speaker having a belief and subsequent disbelief about the same proposition, but rather a belief and then a loss of that belief. A similar assessment should be made for Cohen’s description of the speaker going from making confident “flatness”-ascriptions to then supposedly making “flatness”-denials (which should be viewed as *retractions* of the “flatness”-ascriptions, in my opinion). In what follows, I will present Dretske’s (1981) relevant alternatives theory and his clarification of it (1991) as a non-contextualist theory that provides a way around Unger’s general sceptical conclusion. I also intend for Dretske’s theory to provide an alternative explanation of Cohen’s description of being in the grip of “flatness”-scepticism.²³⁹

2.11 Dretske’s Theory and Relationally Absolute Terms

Using his version of the relevant alternatives theory, Dretske (1981) claims that “flat” is a *relationally* absolute term – along with the terms “empty” and “knows.”²⁴⁰ To understand what he means by a relationally absolute term, consider first the mundane fact that the term “flat” is liberally applied by speakers to different *types* of surfaces, such as roads, fields, tables, and so on. According to Dretske, whether the application of the term “flat” is correct or incorrect crucially depends on the *type* of surface being considered at the time of the application. Dretske argues that what counts as a *relevant* bump on a table, for instance, may not count as a *relevant* bump on a road.²⁴¹ From this observation, Dretske claims that there are distinct thresholds for each type of surface. In applying the term “flat” to a type of surface, Dretske holds that, after the distinct threshold has been passed, there will be no more relevant alternatives in the form of, say, bumps, to rule out for that type of surface. This means that if X is a flat road, then road Y cannot

²³⁹ Cohen likely would want to continue talking in terms of intuitions, but I have given some of my reasons for why I oppose such talk.

²⁴⁰ Dretske 1981, 367.

²⁴¹ Dretske 1981, 366.

be any flatter than road X, since all of the relevant bumps for road X *as* a road, as that type of surface, have already been ruled out. Dretske's theory, then, expresses the absolute character of the term "flat," since another road cannot be flatter than a flat road – both roads *as* tokens of that type of surface – while the term also exhibits a relational character, since whether the road is flat is *in relation* to the type of surface that it is.²⁴²

Unger (1984) proposes another somewhat complimentary interpretation of Dretske's idea of a relationally absolute term, which is to view it as revealing a theory of *sortalism*: the theory that the identity conditions of a thing depend on the *sort* of thing that it is.²⁴³ Although this provides a different angle from which to understand Dretske's relevant alternatives theory and his idea of a relationally absolute term, Unger's interpretation of Dretske as a sortalist should receive cautious acceptance at best. This is because Dretske's interest is with regard to the identity of a thing as having the relevant property in question. That is to say, Dretske is not interested in identifying, say, roads as a sort of thing *per se*, but rather is interested in identifying roads which have the relevant property of being flat. Flat roads may be a sort of road, but to say so is to make a qualification on sortalism – to make the identity conditions depend on both the sort of thing that the thing is and the sort of thing having the relevant property that those sorts of things have – which is my reason for saying that Unger's interpretation should receive cautious acceptance, if it receives any acceptance at all.

In the face of Unger's (1975) threatening sceptical conclusion, Dretske's (1981) theory holds the possibility of relief and compromise. Dretske agrees with Unger that the term "flat" has an absolute character.²⁴⁴ Yet he disagrees with Unger that *any* bump, visible or microscopic, on a particular surface, simply disqualifies that surface from being flat, since what counts as a bump will depend on the type of surface under consideration. Dretske's treatment of the term "flat" as a relationally absolute term captures the distinctive semantic character of the term, on which Unger focuses, while it allows the term to also exhibit the "contextual relativity," which is manifested in the way that ordinary speakers liberally, but not incorrectly, make claims about

²⁴² Dretske 1981, 365.

²⁴³ Unger 1984, 30-33. Although Unger does not say so, it may be of historical interest to note that the famous British empiricist John Locke is often considered the first sortalist – although Locke does not use this name to identify his view – since he proposes what are considered sortal identity conditions for terms such as "animal," "man," and most famously "person." See: Chapter 27 of his *An Essay Concerning Human Understanding*; for an excellent book on the topic of sortalism, see: E.J. Lowe (1989).

²⁴⁴ Dretske 1981, 365.

objects and surfaces in the world. Dretske's treatment of relationally absolute terms harmonizes and explains semantic considerations, pragmatic interests, and empirical facts about the way in which speakers behave. It also has the theoretical virtue of avoiding Unger's sceptical conclusion that speakers are massively ignorant, for, according to Dretske, speakers can accept the term "flat" as an absolute term and still apply it to surfaces which have *irrelevant* bumps on them.²⁴⁵

More recently, in response to Cohen's (1991) insistence that particular temporary conversational interests of speakers set the standards for the correct application of context-sensitive terms, Dretske (1991) clarifies his (1981) theory by contending that group interests can play a significant role in setting the threshold for the relevant alternatives that need to be ruled out.²⁴⁶ Dretske's (1981) previous statement of the relevant alternatives theory did not discuss how group interests could play this role, and the correct application of the term under consideration seemed to depend simply on the type of thing under consideration. To appreciate Dretske's (1991) reason for making this clarification, consider, for instance, that a teacher may walk into a classroom and remark that "it is empty," because there are no students or other people present, while a custodian may enter that same classroom at the same time and remark that "it is not empty," because there are many chairs, desks, and other furniture in the classroom. According to Dretske's clarified theory (1991), the teacher and custodian, as *representatives* of different groups with different interests, are not contradicting one another in stating that the classroom is empty and not empty at the same time, because their statements differ with respect to the correct application of the terms and the relevant alternatives that need to be ruled out to correctly say of a space, in the form of the classroom, that it is empty. The thresholds are set differently in this case, which are in relation to the interests of the representatives. Although he does not say so, Dretske's (1991) clarification of his (1981) theory makes the "relational" aspect of relationally absolute terms more prominent, because a surface is flat *in relation* to both its type and group interests.

²⁴⁵ Dretske uses the term "contextual relativity" to mean the environment of *epistemic subjects*; he is not using it to mean the conversational context of *ascribers* (Dretske 1981, 365). Dretske (1981) clarifies this point in a reply to Cohen (1991), see: Dretske 1991, 191-196. Also, to satisfy Cohen's (1991) curiosities, Dretske declares that he did not have *ascriber* contextualism in mind when wrote his (1981) classic essay, although he admits to not being so alert to the differences between ascriber contextualism and his relevant alternatives theory (Dretske 1991, 192). I feel that this exchange between two distinguished philosophers on important issues concerning ascriber contextualism and the relevant alternatives theory is worthy of scholarly attention.

²⁴⁶ Dretske 1991, 194.

Contextualists seek to have the particular, if not peculiar and perhaps sometimes even idiosyncratic, temporary conversational interests of ascribers and deniers play a decisive role in setting the standards which determine the correct application of relationally absolute terms.²⁴⁷ For the contextualist whether the application of the term is correct or incorrect will be in relation to the conversational context of the ascriber or denier.²⁴⁸ The contextual relativity with which Dretske is concerned, however, is beyond particular temporary conversational interests. For Dretske is concerned with longer-lasting, more permanent, group interests. Group interests often will anchor the error-possibilities in what the speakers *qua* representatives assume the world to be like. Yet, despite the diversity of group interests, this does not threaten to set things adrift. For with a high degree of regularity, these group interests will coincide with one another about what the world is *generally* like, with different members agreeing that the world has some flat stretches of road, flat slates on pool tables, flat tops of desks, and so on – even if, say, the speaker *qua* pool player judges the tops of desks generally to be too bumpy to cover them up with slate and felt and turn them into pool tables. This is in stark contrast to what contextualism sometimes permits, since, to give an explanation of “knowledge”-scepticism or “flatness”-scepticism, the contextualist must honour and take seriously even those radical sceptical error-possibilities raised by the sceptic in sometimes free-floating, radical sceptical, conversational contexts. On this score, it is telling that the contextualist claims to have a clear understanding of how the standards rise, but admits that it remains a mystery as to how those standards drop down again: “But for some reason raising of standards goes more smoothly than lowering.”²⁴⁹ Notice that this is not a difficulty for Dretske, however, because thresholds shift with and track the changes of the speaker *qua* representative’s interests. Speakers can be different representatives on different occasions. For example, a speaker *qua* detective announces that “the room is full (of fingerprints)” on one occasion, but later that speaker, now as speaker *qua* teacher, can announce

²⁴⁷ For an intriguing discussion of this, see: Unger 1984, 30-33; especially where he makes the following statement with Dretske (1981) in mind: “Sortalism differs from contextualism: According to sortalism, there is no very direct connection with a user’s temporary interests and the truth-value of what he says with those expressions in the relevant range. Unlike the contextualist, the sortalist recognizes as having relevant semantic import only those enduring and widespread interests of ours that are already reflected in the categorical words of our common language” (Unger 1984, 31). It is worth noting that Unger (1984) no longer supports invariantism, as he (1975) once did, and he (1984) is not here supporting contextualism; rather, he is a semantic relativist who considers the arguments for invariantism and contextualism to be on par. It is worth noting too that he is dismissive of sortalism as a position of compromise. Also, for a discussion about how particular temporal conversational interests determine standards and his criticism of Dretske (1981) on this matter, see: Cohen 1991, 27.

²⁴⁸ For a clear statement of this, see: Cohen 1991, 17-37.

²⁴⁹ Lewis 1979, 352.

of the same classroom that “the room is not full (of people),” naively ignoring any forensic evidence at the scene. In this case, the larger set of relevant alternatives decreases with a shift in the speaker’s role, which is similar to a shift from higher to lower standards. Unless the contextualist can solve the mystery of how the standards can shift back down after they have gone up, then the particular temporary conversational interests of the denier raising radical sceptical possibilities poses a serious theoretical difficulty for the contextualist. Dretske’s theory can meet this criticism. Better yet, the contextualist criticisms of Dretske on this matter are at best “doubled-edged” – with the rather sharper edge, if I may say so, facing the contextualist.

But what explains “flatness”-scepticism? How can speakers go from being extremely confident in making judgments about flat surfaces to subsequently doubting whether those same particular surfaces are flat? Jonathan Vogel (1990) cites as evidence of such a phenomenon that when ordinary speakers attend to error-possibilities the very act of attending to them makes ordinary speakers treat even rather improbable error-possibilities as much more likely.²⁵⁰ In these circumstances, speakers tend to have strong doubts about whether they really know what they originally claimed to know, and then often retract claims to knowledge. The evidence that Vogel cites may also explain cases of “flatness”-scepticism. The speaker claiming that, say, the table is flat, may retract that claim if someone else calls attention to an unrealized error-possibility, such as that there may be unseen bumps on its surface. The claim may be retracted for the sole reason that the person who attends to the error-possibility may treat that error-possibility as having a much higher probability than it actually has. Even if that speaker retracts the claim, but really should have continued to claim that the table is flat, because she had indeed ruled out all of the relevant alternatives, this would not tell against Dretske’s theory. In this case, the speaker *qua* representative may be simply confused about which standards are in effect. This would be especially possible if the conversational participants are different representatives using different standards for the application of the same term. But could there be any place for the Ungerian sceptic on Dretske’s view? I would suggest that there could be, and that that sceptic exemplifies the standards of a geometer, although such a geometer would have to be rather eccentric to go out of her way to deny the flatness-claims of others. Yet even these standards remain legitimate. Moreover, unlike contextualist standards, these standards often remain public – for all eyes to see.

²⁵⁰ Vogel 1990, 19.

CHAPTER 3

THE “SKEPTICAL PROBLEM,” TRACKING, AND CONTEXTUALISM

I intend for this final chapter to support my claim that distinctively epistemological commitments to either internalism or externalism make the difference for whether the sceptical “paradox” arises. The focus will be on DeRose’s (1995) externalist contextualist “solution” to what he calls the “Skeptical Problem,” which he admits only differs from Cohen’s sceptical “paradox” with respect to its name and details.²⁵¹ DeRose relies on Nozick’s (1981) philosophical explanation of radical scepticism, which comes out of Nozick’s subjunctive conditionals account of knowledge. Nozick admits that agents cannot know the denials of radical sceptical hypotheses, but maintains that this has no tendency to prevent agents from continuing to know ordinary empirical propositions. A familiar pattern emerges: Nozick rejects the closure principle in the face of radical sceptical hypotheses, but in the process mistakenly compromises his externalist commitments by making an internalist assumption about knowledge when radical sceptical hypotheses are under examination. To preserve the closure principle and avoid what he calls “abominable conjunctions,” DeRose contextualizes Nozick’s subjunctive conditionals account of knowledge. Since DeRose is an externalist about knowledge, if he were aware of Nozick’s mistake, then DeRose would realize that the application of ascriber contextualism is simply unnecessary. The “Skeptical Problem” is not a problem for the externalist who remains steadfast in his commitments to externalism.

3.1 Solving the “Skeptical Problem” and the Argument from Ignorance

DeRose (1992; 1995) observes that, “Contextual theories of knowledge attributions have almost invariably been developed with an eye towards providing some kind of answer to philosophical skepticism.”²⁵² DeRose (1995) continues to propose intriguing ways of answering “philosophical skepticism” in a distinctively contextualist manner. More than just answering the sceptic, however, DeRose claims to have solved the “Skeptical Problem.”²⁵³ But what is the “Skeptical Problem”? The “Skeptical Problem” is to figure out exactly what is wrong with the Argument from Ignorance. So we must understand this argument to see what the problem is.

²⁵¹ See: DeRose 1995, 3 Note 5. As we shall see, the main difference between these contextualists is their *epistemological* orientations, with Cohen being an internalist and DeRose an externalist.

²⁵² DeRose 1992, 917. DeRose repeats the claim in: DeRose 1995, 5.

²⁵³ The full title of his paper is “Solving the Skeptical Problem,” see: DeRose 1995, 1-52.

Where “O” stands for an ordinary proposition (for example, the proposition “I have hands”) and “H” a sceptical hypothesis (for example, the hypothesis “I am a bodiless brain-in-a-vat [henceforth, BIV] who has been electrochemically stimulated to have extremely similar experiences to those experiences I would have or have had outside of the vat”), DeRose presents the *Argument from Ignorance* (henceforth, the Argument) as follows:

- (1) I don’t know that not-H
- (2) If I don’t know that not-H, then I don’t know that O. So,
- (3) I don’t know that O.²⁵⁴

DeRose’s Argument has the same structure as the *argument for closure-based scepticism*. DeRose’s Argument also receives the same kind of contextualist treatment that Cohen (1988, 1999) gives to the argument for closure-based scepticism, which is to quarantine the Argument to the sceptical context and keep the quotidian context pristine and free of scepticism. But unlike Cohen’s (1988, 1999) internalist contextualist treatment of scepticism, DeRose’s (1995) contextualist treatment of scepticism is externalist.²⁵⁵ Moreover, DeRose claims that his contextualist externalist “solution” to the Argument better explains how it is that speakers could become ensnared by having conflicting “intuitions.” When the Argument is made using a radical sceptical scenario, such as the BIV scenario, DeRose holds that ordinary speakers sometimes find premises (1) and (2) intuitive, but have a strong intuition against the conclusion (3). As in the argument for closure-based scepticism, the problem, however, is that if premises (1) and (2) are true and the Argument is valid, then this deductively necessitates the conclusion (3). Yet no one but the sceptic claims that the Argument is sound without qualification. The contextualist claims that the Argument is sometimes sound, but quickly makes a number of qualifications on this claim; first, that the Argument is sound only in the sceptical context, in which premise (1) and the conclusion (3) are true; and second that premise (2), as a partial expression of the closure principle, holds relative to the sceptical context, so that the set of entailments of the main proposition in the sceptical context have the property of being closed. DeRose’s externalist

²⁵⁴ DeRose 1995, 17. In DeRose’s view, solving the “Skeptical Problem” requires explaining what is wrong with the Argument and how it sometimes “traps” people.

²⁵⁵ Cohen, for instance, characterizes the relationship between his account and DeRose’s account as follows: “My internalist account construes the strength of one’s epistemic position toward a proposition as to a large part determined by the strength of one’s reasons or evidence. DeRose’s view differs by providing an externalist interpretation of the strength of one’s epistemic position”(Cohen 1999, 70).

“solution” to the Argument has an advantage over Cohen’s internalist “solution,” given that Cohen’s characterization of intrinsically rational beliefs appears to be *ad hoc*. This makes DeRose’s “solution” worthy of closer examination.

After DeRose presents the Argument, he defends the view that the Argument indeed poses a real epistemological problem. He defends this view from both an uncritical dismissive response and what he calls “the Moorean reaction,” both of which are responses that try to block the Argument by rejecting premise (1) that I do not know that sceptical hypotheses are false. To those who respond uncritically and dismissively to the possibility of being a BIV because of the sheer improbability of this, DeRose replies, “For however improbable or even bizarre it may seem to suppose that I am a BIV, it also seems that I don’t know that I’m not one.”²⁵⁶ But why do I fail to know that I am not a BIV? Notice that DeRose casts the Argument in terms of what I can know from within a first-person perspective. DeRose believes that the most effective radical sceptical hypotheses employ this perspective, especially when he observes and then asks: “Upon hearing the hypothesis, typically one can’t help but projecting oneself into it. How would things seem to me if that situation obtained?”²⁵⁷ This first-person perspective is a prominent feature of the Argument, which Duncan Pritchard (2005) focuses on in his characterization of how radical sceptical hypotheses pose problems:

The problem that a sceptical hypothesis such as the BIV hypothesis poses is that it is subjectively indistinguishable to the agent whether her experiences indicate that the world is as she takes it to be as opposed to being completely unlike how she takes it to be, which it would be if she were the victim of the sceptical hypothesis.²⁵⁸

The BIV in DeRose’s paradigmatic example of a radical sceptical hypothesis is supposedly having experiences sufficiently similar to those that I am having right now, but the BIV’s experiences are caused by it being hooked up to a super-computer which generates its experiences, unlike mine which are supposedly being caused by my being a human and standing in a certain relation to the (non-vat) world – at least I hope so. If I were the BIV being electrochemically stimulated, however, it is imaginable that I would still have experiences like those experiences that I am having right now. Presumably, if I were a BIV, I would have no idea of it, for my experiences, from my point of view, would be *subjectively indistinguishable* to me,

²⁵⁶ DeRose 1995, 3.

²⁵⁷ DeRose 1995, 18.

²⁵⁸ Pritchard 2005, 24.

either as caused by the BIV or being a human in the world. The apparent problem posed by the BIV radical sceptical hypothesis, then, is for me to know that I am not a BIV and figure out how I can know this. DeRose's point about the effectiveness of the BIV hypothesis is that I cannot rule it out on the basis that it *seems* to me that I am not a BIV, for that is how it would *seem* to me if I were one. But if we concede premises (1) and (2), then the sceptic appears to be an inference away from carrying the day.

From a theoretical point of view, I do not think that the Argument is a problem for the externalist about knowledge, and that it, therefore, should not bother DeRose either, despite him staunchly defending its status as a genuine problem for all epistemologists. The "Skeptical Problem" is rather a problem for the internalist about knowledge. Accordingly, as an externalist about knowledge, DeRose proposes a contextualist "solution" to the "Skeptical Problem" which is ultimately poorly motivated and unnecessary. My reasons for thinking so will become clearer as these sections develop.

To deal with the Moorean reaction that counters with common sense and denies that the Argument exposes any such ignorance on the part of putative knowers, DeRose contends that the Moorean reaction fails to explain how anyone could ever be ensnared by the Argument: "In seeking a solution to this puzzle, we should seek an explanation of how we fell into the skeptical trap in the first place..."²⁵⁹ Thus, just like Cohen (1988, 1999), DeRose argues that the Moorean reaction fails to meet a *desideratum* for a satisfying resolution of the Argument.

DeRose further holds that an adequate resolution of the Argument has to account for, what Pritchard (2001) calls, the "phenomenology of scepticism," which requires that we explain why ascribers often refrain from making "knowledge"-ascriptions when sceptical hypotheses are the topic of conversation and how conversational factors can be sufficient to cause this form of behaviour in ascribers.²⁶⁰ To account for the "the phenomenology of scepticism," DeRose attempts to refine and contextualize Nozick's philosophical explanation of radical scepticism, which uses an externalist subjunctive conditionals account of knowledge.²⁶¹ DeRose's main supposed refinements include: avoiding abominable conjunctions that arise from Nozick's

²⁵⁹ DeRose 1995, 4.

²⁶⁰ Pritchard 2001, 331.

²⁶¹ Nozick 1981, 167-288.

rejection of the closure principle; and removing the threat of scepticism, while preserving the closure principle, by endorsing ascriber contextualism. I need to present the most critical parts of Nozick's subjunctive conditionals account of knowledge to put us in a position to appreciate and criticize both DeRose's account of the "phenomenology of scepticism" and his attempt to solve the "Skeptical Problem."

3.2 Nozick's Subjunctive Conditionals Account of Knowledge

Nozick (1981) defines knowledge as follows: "knowledge is a real relationship to the facts, subjunctively contoured."²⁶² In seeking an externalist account of knowledge, Nozick *tentatively* proposes, as a first approximation, four individually necessary and jointly sufficient conditions on knowledge.²⁶³ Where "p" is some ordinary contingent proposition and "S" is a subject, Nozick claims that the following four conditions must hold in order for S to know that p:

- (a) Truth Condition: p is true
- (b) Belief Condition: S believes that p
- (c) Variation Condition: If p were not true, S would not believe that p
- (d) Adherence Condition: If p were true, S would believe that p.²⁶⁴

When conditions (a) and (b) are met, these conditions hold in the actual world as a matter of fact, whereas conditions (c) and (d) are tracking conditions containing subjunctive conditionals, which are met by a possible-worlds-analysis, in which the agent's belief tracks the truth (of p) through different possible worlds. The tracking conditions (c) and (d) are intended to ensure that the factual conditions (a) and (b) are non-accidentally related to one another, with the variation condition (c) being the one on which Nozick primarily focuses.²⁶⁵ On Nozick's view, the tracking conditions ensure that that agent's true belief is responsive to counterfactual circumstances, which makes the true belief non-accidental in the actual world and an instance of

²⁶² Nozick 1981, 288.

²⁶³ Strictly speaking, only after making the appropriate modifications to conditions (b), (c), and (d) by making the belief-clauses of each condition relative to a method, are the conditions, on Nozick's *considered* account, jointly sufficient for knowledge.

²⁶⁴ Nozick 1981, 172-178.

²⁶⁵ A number of critics and commentators, including DeRose (1995), call the "variation condition" the "sensitivity condition" or principle. It seems to me that very little turns on this and the difference is probably one of mere words. For a discussion of the terminology, and for this claim that the differences in terminology are insubstantial, see: Pritchard 2008, 438. I will continue to call this condition the "variation condition," as Nozick did. Nevertheless, in discussions of whether a belief tracks the truth across possible worlds, the variation condition is almost always the condition on which Nozick is focused, with the meeting of adherence condition often being simply assumed.

knowledge.²⁶⁶ Nozick's account of knowledge is admittedly externalist, because the agent cannot be cognitively aware of the non-accidental connection between the factual conditions (a) and (b), since possible worlds and counterfactuals are a matter of analysis, which remain beyond that agent's first-person perspective.²⁶⁷ A consideration of Nozick's concern over conditions (a) and (b) standing in a non-accidental relationship to one another, together with his denial of the need for any intermediate justification condition (as traditionally understood) on knowledge, reveals the kinship between Nozick's account and Armstrong's (1973) externalism; although Armstrong stresses the naturalistic character of the crucial relationship between belief and truth in the *actual* world, whereas Nozick finds subjunctive contours on that relationship.²⁶⁸

To explain how the tracking conditions (c) and (d) are met, Nozick provides a possible-worlds-analysis.²⁶⁹ Understood as an analytic tool, a possible world is a world which is supposed to be causally isolated from the actual world and other possible worlds, and which differs from the actual world in a certain relevant respect that makes it useful to an analysis of subjunctive conditionals.²⁷⁰ Possible worlds are ordered in terms of their similarity to the actual world, with

²⁶⁶ There are a number of famous epistemological cases that Nozick's account solves, including: Gettier cases, the barn facades case, and stopped clock cases. Duncan Pritchard (2005) recently has claimed that a central intuition of epistemology is an anti-luck intuition. Nozick's account rules out luckiness of true belief, although Pritchard (2005) and Sosa (1999) find Nozick's characterization of the variation condition (or sensitivity condition) problematic and opt for the safety principle, which is akin to the adherence condition in Nozick's account.

²⁶⁷ Nozick admits that his subjunctive conditionals account of knowledge is externalist in orientation (Nozick 1981, 280-283).

²⁶⁸ Armstrong's externalism is naturalistic because he talks about "reliable indicators" of the truth, which are found in the actual world. According to Armstrong (1973), beliefs formed on the basis of "reliable indicators" typically lead to knowledge.

²⁶⁹ Nozick 1981, 173.

²⁷⁰ Pritchard 2009, 25-26. In defining a "possible world," I am relying on Pritchard's (2009) way of talking about possible worlds as an analytic tool, which he ascribes to Nozick. I also believe that Nozick employs the notion of "possible worlds" for epistemological *purposes*. The metaphysical issues surrounding the possible-worlds-analysis are extremely controversial and thorny, given that a possible-worlds-analysis often commits one to a controversial metaphysical view that traces its history to the seventeenth-century German philosopher Gottfried Leibniz, whose theology spoke of this world as the best of all possible worlds, and to David Lewis's (1968, 1973, 1986) *modal realism*, which boldly asserts that possible worlds are as real as the world that we call our own. So Leibniz and Lewis do not claim that a possible world is to be used merely as an analytic tool, but claim that possible worlds are indeed real. I do not have the space to discuss these metaphysical issues. I will say that, *strictly speaking*, the agent in the possible world cannot be me, for the very reason that the relation of transitivity does not hold across possible worlds, and thereby there is no identity relation between me and the agent in that possible world. Accordingly, Lewis calls the agent in a possible world, the agent who is in many ways like me but not me, my "counterpart." See his "counterpart" theory: Lewis 1968. For a discussion of particular issues of identity in Lewis's theory, see: Grayling 2005, 58-65. I shall talk in the way that Nozick does, without identifying the agent in the possible world as my "counterpart"; that is to say, without distinguishing the agents – me and my counterpart – and will talk simply about the agent.

nearby possible worlds being much more similar to the actual world than are far away ones.²⁷¹ With this possible-worlds-analysis in mind, Nozick's aim is to figure out whether the agent, if resituated in a possible world where what was believed in the actual world would now be false because at least one fact about the actual world is different in the nearby possible world, would in that possible world now (mistakenly) believe what that agent did in the actual world. If so, then the agent's actual world belief is merely accidental.²⁷² Essentially, Nozick is trying to find out whether the agent's belief in the actual world would track the truth across possible worlds or if the agent in the possible world would simply lose sight of these tracks that lead to the truth. To get a sense of how this analysis is supposed to work, consider, for instance, the most widely cited example in the literature: the proposition "I have hands." To know this proposition, we need to see whether my belief tracks the truth. In the actual world, assume it is true both that I have hands and that I believe I have two. Consequently, conditions (a) and (b) are met. For my true belief to be an instance of knowledge, however, we need to test the strength of this connection between the factual conditions (a) and (b) by using subjunctive conditionals and the possible-worlds-analysis. One nearby possible world, slightly different from ours, in which the proposition "I have hands" is false is the possible world where I have been the victim of a grisly accident and have lost my hands. One far away possible world, very different from ours, in which the proposition "I have hands" is false is the possible world where I am a handless BIV hooked up to a super-computer on Alpha Centauri. This is a far away possible world, since the actual world is radically dissimilar to this possible world, with this far away possible world having super-technology that we do not currently have. Most important to this analysis, however, is to notice that were I in that nearby possible world where I have been the victim of that grisly accident, my belief would indeed track the truth, because I would see that my hands are missing and would thereby believe that I am handless.²⁷³ Yet this is not so in the far away possible world

²⁷¹ Nozick 1981, 173. See: Pritchard 2009, 25-26.

²⁷² A paradigmatic case would be to imagine the agent in the actual world having formed a belief about what time it is from having seen a (broken) *stopped* clock. In the case, it happens that the agent looks at the clock at just the right time and forms a true belief about what time it is. But in a nearby possible world where the clock remains stopped and the agent happens to look at the stopped clock a minute before or a minute after she did in the actual world, then she will form a false belief about what time it is, which will reveal the accidentality of the relevant belief under examination. (The possible world where everything else is the same but the time in that possible world is always a minute behind ours is indeed a *possible* world.)

²⁷³ Although few critics and commentators make them, there would need to be *ceteris paribus* clauses made for this example, since I may not track, in this possible world, this belief that I am missing hands, if I were, for instance, the recent victim of the accident in a hospital bed doped up on painkillers or perhaps acutely suffering from phantom

where I am a BIV on Alpha Centauri, since I would continue to (falsely) believe that I have hands, given that the super-computer feeds me (as a BIV) a simulated experience of hands, although I am really a handless BIV. Therefore, my belief that I have hands would not track the truth in this far away possible world or in far away possible worlds generally.

From this general conclusion, that I do not know the denials of radical sceptical hypotheses because my beliefs do not track the truth in far away possible worlds, Nozick thinks that he must reject the closure principle. For given that the proposition “I have hands” entails the falsity of “I am a handless-BIV,” Nozick thinks that it must be the case that if I can know that “I have hands” (in nearby possible worlds), can know that this proposition entails the falsity of the entailment that “I am a handless-BIV,” but yet always fail to know this entailment that “I am not a handless-BIV” (in far away possible worlds), then the closure principle does not hold. It is worth noting that Nozick, unlike Dretske (1970), is not explicitly giving a relevant alternatives account of knowledge to motivate his rejection of the closure principle, although Nozick’s account holds that in order for the subject to have knowledge that subject’s belief must track the truth in *all* of the *nearby* possible worlds, a subset of the possible worlds, which makes some theoretical commitment to salience or relevance, however implicit.²⁷⁴ The major problem with Nozick’s rejection of the closure principle, however, is that he ends up compromising his externalist commitments by making an internalist assumption about the method by which the agent gains a belief.

In rejecting the closure principle, Nozick makes a mistake remarkably similar to Dretske’s mistake, insofar as both externalists compromise their externalist commitments by claiming that subjects do not know the denials of sceptical alternatives, because these subjects cannot, from a first-person perspective, have such knowledge. To see this mistake, we need to appreciate Nozick’s peculiar characterization of the term “method.” On Nozick’s *considered* account of knowledge, factual condition (b) and counterfactual conditions (c) and (d) need to hold relative to the method used by the subject in coming to a belief. Nozick demonstrates this need using his famous example of the grandmother:

limb syndrome. In these variations of the example, I may believe that I have hands, despite having none, which would mean that my belief would not track the truth in these circumstances.

²⁷⁴ This commitment should not be surprising, however, given that possible worlds are ordered according to how similar they are to the actual world.

A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. Clearly, we must restate our conditions to take explicit account of the ways and methods of arriving at the belief.²⁷⁵

Nozick clearly needs to distinguish which method is being used to come to the relevant belief when testing knowledge, since, in the grandmother example, the grandmother's belief that her grandson is well, gained from seeing him, indeed tracks the truth in nearby possible worlds; whereas, this same belief, gained from the testimony of her family, would not track the truth in nearby possible worlds. On the one hand, in a nearby possible where her grandson would be sick in a hospital bed, were the grandmother to base her belief on sight, then the grandmother would *not* believe that he is well, since she would not see him as being even ambulatory. On the other hand, in a nearby possible world where her grandson is sick in a hospital bed, were she to base her belief on the testimony of her family, then the grandmother would now (falsely) believe that her grandson is well. The grandmother's belief tracks the truth via the method of sight, but not via the method of testimony, which makes it necessary to identify the method that the agent used to gain the belief when evaluating whether a true belief counts as knowledge. Although Nozick clearly distinguishes the method of sight from that of testimony in the grandmother example, what should we say about the previous example of having hands? Should we also distinguish between methods, making sure to distinguish between my seeing I have hands from my being a BIV having a simulated experience of having hands when in fact I have none?

One would think that the answer to this question is that these are definitely distinguishable methods being used in coming to the belief that I have hands; namely, that basing my belief that I have hands on sight is distinguishable in terms of method from basing my belief that I have hands on receiving stimulation from the super-computer to which I am, as the BIV, connected. Nozick claims that an agent can use a method "...without knowledge or awareness of what method he is using," as any thoroughgoing externalist should.²⁷⁶ But Nozick goes on to characterize "method" using a distinctively internalist criterion:

²⁷⁵ Nozick 1981, 179.

²⁷⁶ Nozick 1981, 184.

...any method experientially the same, the same ‘from the inside’, will count as the same method. Basing our beliefs on experiences, you and I and the person floating in the tank are using, for these purposes, the same method.²⁷⁷

To make this equivalence of method based on experiential or phenomenological sameness is, for the externalist, to make an egregious mistake. The reason it is egregious is that externalists claim that knowledge depends on external relations; that is to say, on whether my beliefs are *in fact* appropriately connected to the world, and so these beliefs need not depend on what I know from within a first-person perspective, even in the face of radical sceptical alternatives.²⁷⁸ In response to Nozick, Michael Williams (1991) asks the following critical question: “Whether my beliefs track the truth is not determined by ‘internal’ factors, so why should the method I am using be individuated in such terms?”²⁷⁹ I submit that Nozick individuates methods in internalist terms for the very same reason that Dretske (1970; 2005) disallowed subjects from knowing the falsity of sceptical alternatives; namely, because Nozick too is trying to prevent bootstrapping in epistemology by denying the closure principle, and achieves this through his internalist characterization of methods. Let me briefly try to justify this claim.

In considering what a subject can know from within the first-person perspective when that subject is using the *method* of deductive inference, a method which centrally concerns the status of the closure principle, Nozick writes:

The problem is that the method of inferring q from know p is indistinguishable by the person from the method of inferring q from (believed) p. We have said that knowledge is a real connection of belief to the world, tracking, and though this view is external to the viewpoint of the knower, as compared to traditional treatments, it does treat the method he uses as identified from the inside, to the extent it is guided by internal cues and appearances.²⁸⁰

This is one of the main pieces of argumentation that Nozick uses to justify his claim that seemingly different methods which are experientially (internally) alike are really the same method. For instance, when I use the method of deductive inference, Nozick is arguing that I *cannot* know, from my knowing that I have hands, and my knowing that I have hands entails that I am not a BIV, that I am, therefore, not a BIV. This is a closure argument which I am attempting

²⁷⁷ Nozick 1981, 184-185.

²⁷⁸ Perhaps, especially so, since it looks like radical sceptical alternatives are testing an externalist’s commitment to externalism.

²⁷⁹ Williams 1991, 343.

²⁸⁰ Nozick 1981, 232.

to make, but one which I fail to make, according to Nozick, because I cannot know this entailment on Nozick's possible-worlds-analysis (provided, of course, that this peculiar characterization of methods is in place). Doubtless, if I have hands, and a BIV is a handless being, then this entails that I am not a BIV. There is nothing objectionable with this inferential process. But in blocking the inference about what I can know from a limited first-person perspective, from the method I am using, Nozick is making precisely the same kind of mistake that Dretske (1970) makes in not realizing that *my knowing*, from a first-person perspective, that I am not a BIV, is not an entailment of my knowing that I have hands. To be clear, "my knowing" is simply not found among the *logical* consequences of known propositions. Knowledge from a first-person perspective lies outside of the scope of the closure principle, so that it would not be my knowing from within the first-person perspective the falsity of a sceptical hypothesis. This is why Nozick's thinking otherwise undermines his motivation for rejecting the closure principle. For an externalist, the way I can know that I am not a BIV is not from within a first-person perspective, but rather is by the way in which my belief is *in fact* appropriately related to the world, which would include factual conditions and may also include counterfactual conditions, all of which can remain outside of my perspective and external to me. Nozick should have kept saying so. He should not have compromised his externalist commitments by making an internalist assumption about methods in the face of the BIV radical sceptical hypothesis.

3.3 DeRose's Rule of Sensitivity and the Phenomenology of Scepticism

DeRose's externalist contextualist account of non-propositional "knowledge" relies on Nozick's idea of tracking, especially Nozick's variation condition, to account for the "phenomenology of scepticism."²⁸¹ DeRose is trying to explain how conversational factors in the conversational context, in which the ascription is made, could raise the standards that determine the proper application of "knowledge"-ascriptions to epistemic subjects. Recall that DeRose (1992) was concerned with how attributor or contextual factors, such as the mentioning or considering of error-possibilities in conversational contexts, could cause variation in the truth-

²⁸¹ Recall that, for the contextualist, the word "knowledge" should be put in scare quotes to mark that contextualism is a metalinguistic thesis, involving semantic ascent, which is about the *meaning* of the word "knowledge," and not about propositional knowledge. I have tried to do this throughout, and this is a perfect example of why it is important to do so. In DeRose's words: "Unlike Nozick, I'm not presenting an analysis of propositional knowledge"(DeRose 1995, 21). This is because DeRose, as a contextualist, is concerned with the meaning of the word "knowledge" and how that word is context-sensitive.

conditions or semantic standards that determine the meaning of “knowledge”-ascriptions.²⁸² In building upon this idea, DeRose (1995) tries to discover the conversational mechanism that shifts contexts and explains how contextual factors can cause variation in the meaning of “knowledge”-ascriptions. It is crucial for DeRose to do so, since he claims that the “knowledge”-ascription in the quotidian context *means* something different from the “knowledge”-ascription in the sceptical context.

By leaning on Nozick’s understanding of tracking, DeRose claims that the context-shifting conversational mechanism is the following *Rule of Sensitivity*:

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.²⁸³

Richard Feldman (2001) provides a nice gloss on what DeRose means by “sensitivity”: “S’s belief that p is sensitive if it is not the case that if p were false, S would believe p.”²⁸⁴ DeRose’s Rule of Sensitivity, then, is really Nozick’s variation condition applied to non-propositional “knowledge” in a conversational context. Feldman also succinctly captures DeRose’s notion of *strength of epistemic position* as it is applied to DeRose’s “solution” to the “Skeptical Problem” as follows: “The strength of one’s epistemic position is measured by the range of worlds through which one accurately tracks the truth value of a proposition. The further out from reality that one tracks a proposition, the stronger one’s epistemic position.”²⁸⁵ However, as Feldman also correctly points out, *context* plays a key role in determining what DeRose calls “the sphere of epistemically relevant worlds,” in which the subject’s belief needs to remain sensitive in this set of possible worlds for that subject to have “knowledge.”²⁸⁶ Consequently, if the BIV sceptical hypothesis, for instance, is raised in a conversation, then when the ascriber tries to ascribe “knowledge” to the subject, that subject must be in a strong enough epistemic position to rule out the BIV hypothesis. This would mean that that subject’s belief must remain sensitive in the far

²⁸² DeRose 1992, 915.

²⁸³ DeRose 1995, 37. DeRose introduces the Rule in the following way: “Let’s call the conversational rule this new account posits as the mechanism by which the skeptic raises the standards for knowledge the ‘Rule of Sensitivity’”(DeRose 1995, 37).

²⁸⁴ Feldman 2001, 70.

²⁸⁵ Feldman 2001, 70.

²⁸⁶ Feldman 2001, 70-71; DeRose 1995, 37.

away possible world where the subject is the BIV, provided that the “knowledge”-attribution is made within the (sceptical) context, in which that far away possible world is within the sphere of epistemically relevant worlds. Although, in principle, it is supposedly possible for a subject to rule out the BIV hypothesis, it is practically impossible to do so. Thus, the Rule of Sensitivity raises the standards that determine how good an epistemic position the subject needs to be in, which, in turn, determine the proper application of the “knowledge”-ascription.

DeRose diverges from Nozick, however, in denying the importance of the subject’s belief having to be evaluated relative to the method used to come to that belief. According to DeRose, it is unnecessary for the contextualist to follow Nozick on this point:

One possibility here is to follow Nozick very closely by modifying that generalization [S knows that P if S’s belief tracks the truth] so that it refers to Nozick’s modified, rather than his original, third condition, and thus, like Nozick, explicitly relativizing our account to the method by which S believes that P. Often, though, context takes care of this for us.²⁸⁷

Essentially, context takes care of this because ascribers or deniers will be aware of, say, in the grandmother example, the family’s false testimony to the grandmother *as* an error-possibility raised in the conversational context, rather than this error-possibility – at the first-order level of propositional knowledge – as one which would simply prevent the grandmother’s belief from tracking the truth across nearby possible worlds.

Nozick’s pivotal move to ensure the possibility of knowledge was to claim that knowledge could be gained so long as the subject’s belief tracks the truth in all of the nearby possible worlds. To achieve this, Nozick denied the closure principle and did so by making a dubious internalist assumption about methods, specifically by contending that my belief does not track the truth in far away possible worlds if we assume that the method by which I come to believe that I have hands from seeing them is the same as the one by which I come to this same belief by being a BIV. According to DeRose, it follows from Nozick’s analysis that I can know that I have hands *and* cannot know that I am not a handless BIV, since my belief that I have hands tracks the truth in nearby possible worlds but fails to do so in far away ones. DeRose calls the conjoining of these two facts of analysis, which yield this “intuitively bizarre result,” an

²⁸⁷ DeRose 1995, 22.

“abominable conjunction.”²⁸⁸ Abominable conjunctions are the result of Nozick’s rejection of the closure principle.²⁸⁹ In response to Nozick’s rejection of the closure principle, DeRose claims that “...it would be no wiser to bet one’s immortal soul on O’s [that I have hands] being true than to bet on not-H’s [that I am not a handless BIV] being true.”²⁹⁰ To preserve the closure principle and thereby avoid abominable conjunctions, DeRose contextualizes Nozick’s account using the Rule of Sensitivity. Specifically, DeRose combines the Rule of Sensitivity with what he calls the *Basic Strategy*, which states:

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.²⁹¹

DeRose contends, then, that the closure principle holds relative to the context in which the “knowledge”-attribution is made, so that speaker A’s attribution that “S knows that P” is true if S is in a strong enough epistemic position with respect to the standards in force in speaker A’s context at the time of the “knowledge”-attribution. By talking about strength of epistemic position in terms of possible worlds and contextual standards in terms of speaker A’s conversational context, DeRose employs ascriber contextualism to keep closed the set of propositions and entailments in the particular context in which the “knowledge”-attribution or “knowledge”-denial was issued. Essentially, contextualism preserves the closure principle and is the means of evaluating arguments containing allegedly context-sensitive terms from the epistemic vocabulary.

DeRose’s Basic Strategy enables the ascriber to ascribe “knowledge” to agents in the quotidian context with lax standards and the sceptic to deny “knowledge” to agents in the sceptical context with strict standards. Although the Rule of Sensitivity tends to raise the standards, DeRose clearly wants to ensure that “knowledge” can be ascribed in the quotidian context. On this point, he writes:

²⁸⁸ DeRose 1995, 28.

²⁸⁹ Note that very rarely would anyone utter an abominable conjunction, because of its conversational impropriety. Moreover, Nozick (1981) and Dretske (1970) are aware that rejections of the closure principle are counterintuitive.

²⁹⁰ DeRose 1995, 32.

²⁹¹ DeRose 1995, 5.

For the fact that the skeptic can install very high standards that we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates.²⁹²

The reason why DeRose focuses so much attention on the very high standards of scepticism is that he is trying to account for the “phenomenology of scepticism,” which he believes the Rule of Sensitivity reveals, insofar as the sceptic raises radical sceptical error-possibilities in the conversational context that the other conversational participants recognize as the subject being unable to rule out; namely, because that subject's belief would be insensitive and would not track the truth across possible worlds if these error-possibilities were to obtain.

Things are not this straightforward, however. DeRose's most glaring oversight is overlooking Nozick's mistaken internalist assumption about methods. In the quotation above, in which DeRose discusses Nozick's modification to his subjunctive conditionals account of knowledge by making conditions (b), (c), and (d) relative to a method, DeRose clearly misses Nozick's mistake and instead enthusiastically declares that Nozick's modification is unnecessary for the ascriber contextualist, since context will take care of this. Had he realized Nozick's mistake, DeRose likely would not have applied ascriber contextualism to Nozick's subjunctive conditionals account of knowledge in order to try to solve the “Skeptical Problem.” But DeRose's account was really doomed from the beginning, seeing that he believed that the “Skeptical Problem” was indeed a genuine problem for all epistemologists, because agents seemingly cannot know the falsity of radical sceptical hypotheses. Recall that his initial explanation of why the “Skeptical Problem” poses a problem is that I, as an agent, cannot know that I am not a BIV from it *seeming* to me that I am not BIV, because that is how it would *seem* to me if I were a BIV. Of course, I cannot know this *if* we are making an internalist assumption about knowledge. Nozick makes this assumption and ends up denying the closure principle. DeRose makes this same assumption, but applies ascriber contextualism to Nozick's subjunctive conditionals account to preserve the closure principle and to attempt to solve the “Skeptical Problem.” DeRose sees this application as an improvement on Nozick's externalist epistemology. But the needed improvement would be to correct Nozick's mistaken internalist assumption about methods. The “Skeptical Problem” would disappear had DeRose not followed

²⁹² DeRose 1995, 39.

Nozick and made internalist assumptions. Ultimately, DeRose's "solution" to the "Skeptical Problem" is unnecessary.

CONCLUSION

A main lesson I have drawn throughout has been that whether the sceptical “paradox” (or “Skeptical Problem”) arises depends on whether a commitment is made to internalism. A familiar pattern that has emerged in the course of this investigation has been that a commitment to externalism proves that there is no obvious reason to view radical sceptical hypotheses as threatening, and thereby renders the application of ascriber contextualism unnecessary. Both Stine (1976) and DeRose (1995) conformed to this pattern. After having made commitments to externalism, they should have continued to view radical sceptical hypotheses as non-threatening. Yet both of them continued to perceive radical sceptical hypotheses as threatening, and applied ascriber contextualism to their epistemologies to remove this supposed threat. Externalism alone, however, would have achieved this and would have prevented the sceptical “paradox” from arising.

Ascriber contextualists have tried to emend Dretske’s (1970) and Nozick’s (1981) epistemologies by contextualizing them in order to preserve the closure principle and explain scepticism. But ascriber contextualists have not realized that these epistemologies are fundamentally flawed, because of mistaken internalist assumptions. Correcting these assumptions would have been enough, but ascriber contextualists do not recognize these mistakes, and rather have inherited and continue to suffer from them. Cohen (1988, 1999) appeared to be the exception to this pattern, for he explicitly accepted internalism. A significant finding made in the first chapter was that Cohen’s “solution” to the sceptical “paradox” appeared *ad hoc* and was ultimately unacceptable. A related, more general, negative conclusion by Pritchard (2005) was that it is impossible to be internalistically justified in believing the denials of radical sceptical hypotheses. Consequently, there can be no real possibility of an internalist “solution” to the sceptical “paradox,” even if that “solution” combines internalism with ascriber contextualism.

Despite these negative conclusions, an anticipated response from ascriber contextualists is that they will contend that they alone have the only viable explanation for how the sceptical “paradox” arises in the first place. Of course, if one is an externalist, then there is no need to explain how it arises, since there is no “paradox.” The internalist, however, will indeed need such an explanation. To provide such an explanation, we have seen that ascriber contextualists

contend that the meaning of epistemic vocabulary is context-sensitive. Yet the findings of the second chapter proved that none of the accounts of the indexicality of the word “knows” do what they purport to show. This casts doubt on whether such an account can be given. But without such an account, ascriber contextualists cannot continue to claim that they have the only viable explanation of how the sceptical “paradox” arises in the first place.

One important issue, beyond the scope of this thesis, but in the near vicinity of the epistemological issues discussed, is how externalists can provide an adequate response to radical scepticism. This response will not be from within a first-person perspective, but will be developed from a distinctively externalist point of view. Critics of externalism find externalist responses to radical scepticism to be question-begging and to have given up on a traditional goal of epistemology that traces its roots to Descartes. Moreover, externalists tend to be less than sanguine in developing such responses. Evidence of this may come from the fact that many externalists make internalist assumptions in the face of radical sceptical hypotheses. Lately, however, there have been significant developments in such externalist responses. This should give externalists reasons to be optimistic and remain steadfast in their commitments. Michael Bergmann (2006), for instance, canvasses a number of different viable externalist responses to radical scepticism. Michael Williams (1991) too has developed an externalist response to radical scepticism, which is to diagnose the “doubts” of this form of scepticism as *unnatural*. The character of these “doubts” is markedly different from the character of what are normally considered doubts. For example, upon setting out on a long foreign travel, the doubt I experience in believing that I may not have bought enough insurance for the trip is quite different from the “doubt,” if any, I may experience in considering whether I am a BIV. The former sense of doubt is quite natural, whereas the latter sense seems unnatural and manufactured. The latter “doubts” are typical of radical scepticism. Their artificiality should make us suspicious of them, however. At the very least we should not exaggerate their significance, although ascriber contextualists have a bad habit of doing so.

WORKS CITED

- Armstrong, David. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press (1973).
- Bach, Kent. "Context *ex Machina*." In *Semantics versus Pragmatics*. Ed. Zoltan Gendler Szabo. Oxford: Oxford University Press (2005a): 15-45.
- _____. "The Emperor's New 'Knows.'" In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Eds. Gerhard Preyer and Georg Peter. Oxford: Oxford University Press (2005b): 51- 89.
- Bergmann, Michael. *Justification without Awareness: A Defense of Epistemic Externalism*. Oxford: Oxford University Press (2006).
- Blaauw, Martijn. "Challenging Contextualism." *Grazer Philosophische Studien* 69 (2005): 127-146.
- BonJour, Laurence, "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5 (1980): 53-73.
- Chisholm, Roderick. *The Problem of the Criterion*. Milwaukee, Wisconsin: Marquette University Press (1973).
- _____. "On the Nature of Empirical Evidence." In *Empirical Knowledge*. Eds. Roderick Chisholm and Robert Swartz. Englewood Cliffs, N.J.: Prentice Hall (1973b).
- Cohen, Stewart. "Knowledge and Context." *Journal of Philosophy* 83 (1986): 574-583.
- _____. "Knowledge, Context, and Social Standards." *Synthese* 73 (1987): 3-26.

- _____. "How to be a Fallibilist." *Philosophical Perspectives* 2 (1988): 91-123.
- _____. "Skepticism, Relevance, and Relativity." In *Dretske and His Critics*. Ed. Brian McLaughlin. Cambridge, MA: Blackwell (1991): 17-37.
- _____. "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives* 13 (1999): 57-90.
- _____. "Contextualism and Skepticism." *Philosophical Issues* 10 (2000): 94-107.
- _____. "Is Knowledge Contextual?" In *Contemporary Debates in Epistemology*. Eds. Matthias Steup and Ernest Sosa. Oxford: Blackwell (2005): 56-71.
- _____. "Ascriber Contextualism" In *The Oxford Handbook of Skepticism*. Ed. John Greco. Oxford: Oxford University Press (2008): 417-436.
- Corazza, Eros. *Reflecting the Mind: Indexicality and Quasi-Indexicality*. Oxford: Clarendon Press (2004).
- Dancy, Jonathan. *An Introduction to Contemporary Epistemology*. Oxford: Blackwell Publishing (1985).
- DeRose, Keith. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52 (1992): 913-929.
- _____. "Solving the Skeptical Problem." *Philosophical Review* 104 (1995): 1-52.
- _____. "Contextualism: An Explanation and Defence." In *Epistemology*. Eds. John Greco and Ernest Sosa. Oxford: Blackwell (1999): 187-205.
- _____. "Now You Know It, Now You Don't." *Proceedings of the Twentieth World Congress of Philosophy* (2000): 91-106.

- _____. "Single Scoreboard Semantics." *Philosophical Studies* 119 (2004): 1-21.
- _____. "The Ordinary Language Basis for Contextualism and the New Invariantism." *The Philosophical Quarterly* 55 (2005): 172-198.
- _____. "'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments Against Contextualism." *Philosophy and Phenomenological Research* 73 (2006): 316-338.
- _____. "Gradable Adjective: A Defence of Pluralism." *Australasian Journal of Philosophy* 86 (2008): 141-160.
- Dretske, Fred. "Epistemic Operators." *Journal of Philosophy* 67 (1970): 1007-1023.
- _____. "Conclusive Reasons." *Australasian Journal of Philosophy* 49 (1971): 1-22.
- _____. "The Pragmatic Dimension of Knowledge." *Philosophical Studies* 40 (1981): 363-378.
- _____. "Dretske's Replies." In *Dretske and His Critics*. Ed. Brian McLaughlin. Cambridge, MA: Blackwell (1991): 180-221.
- _____. "The Case against Closure." In *Contemporary Debates in Epistemology*. Eds. Matthias Steup and Ernest Sosa. Oxford: Blackwell (2005): 13-26.
- Feldman, Richard. "Skeptical Problems, Contextualist Solutions." *Philosophical Studies* 103 (2001): 61-85.
- _____. "Comments on DeRose's 'Single Scoreboard Semantics.'" *Philosophical Studies* 119 (2004): 23-33.
- Furmerton, Richard. *Knowledge*. Oxford: Blackwell Publishing (2006).
- Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis* 26 (1963): 144-146.
- Goldman, Alvin. "A Causal Theory of Knowing." *Journal of Philosophy* 64 (1967): 355-372.

_____. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73 (1976): 771-791.

_____. "What is Justified Belief?" In *Justification and Knowledge*. Ed. George Pappas.

Dordrecht: D. Reidel (1979): 1-23.

Grayling, A.C. *An Introduction to Philosophical Logic*. 3rd Edition. Oxford: Blackwell Publishing (2005).

Hawthorne, John. "Reply to Cohen." *Philosophical Issues* 10 (2000): 117-120.

_____. *Knowledge and Lotteries*. Oxford: Oxford University Press (2004).

_____. "The Case for Closure." In *Contemporary Debates in Epistemology*. Eds. Matthias Steup and Ernest Sosa. Oxford: Blackwell (2005): 26-43.

Heller, Mark. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy* 77 (1999): 91-114.

Kaplan, David. "Demonstratives." In *Themes from Kaplan*. Eds. J. Almog, J. Perry, H. Wettstein. Oxford: Oxford University Press (1989): 481-563.

Kornblith, Hilary. "Internalism and Externalism: A Brief Historical Introduction." In *Epistemology: Internalism and Externalism*. Ed. Hilary Kornblith. Oxford: Blackwell Publishing (2001): 1-9.

Lewis, David. "Counterpart Theory and Quantified Modal Logic." *Journal of Philosophy* 65 (1968): 113-126.

_____. *Counterfactuals*. Oxford: Blackwell Publishing (1973).

_____. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8 (1979): 339-359.

- _____. *On the Plurality of Worlds*. Oxford: Blackwell Publishing (1986).
- _____. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (1996): 549–567.
- Lowe, E.J. *Kinds of Being: A Study of Individuation, Identity and the Logic of Sortal Terms*. Oxford: Basil Blackwell (1989).
- Moore, G. E. "Certainty" in *Philosophical Papers*. London: Allen and Unwin (1959).
- Neta, Ram. "Contextualism and the Problem of the External World." *Philosophy and Phenomenological Research* 63 (2003): 1-31.
- Nichols, Shaun; Stich, Stephen; Weinberg, Jonathan. "Normativity and Epistemic Intuitions." In *Biological and Cultural Bases of Human Inference*. Eds. Riccardo Viale, Daniel Andler, Lawrence Hirschfeld. Mahwah, New Jersey: Lawrence Erlbaum Associates (2006): 191-222.
- Nozick, Robert. *Philosophical Explanations*. Oxford: Oxford University Press (1981).
- Predelli, Stefano. *Contexts: Meaning, Truth, and the Use of Language*. Oxford: Oxford University Press (2005).
- Pritchard, Duncan. "Contextualism, Scepticism, and the Problem of Epistemic Descent." *Dialectica* 55 (2001): 327-349.
- _____. "Resurrecting the Moorean Response to Scepticism." *International Journal of Philosophical Studies* 10 (2002): 283-307.
- _____. *Epistemic Luck*. Oxford: Oxford University Press (2005).
- _____. "Sensitivity, Safety, and Antiluck Epistemology." In *The Oxford Handbook of Skepticism*. Ed. John Greco. Oxford: Oxford University Press (2008): 437-455.

- _____. *Knowledge*. New York: Palgrave MacMillan (2009).
- Pryor, James. "Highlights of Recent Epistemology." *British Journal for the Philosophy of Science* 52 (2001): 95-124.
- Reiber, Steven. "Skepticism and Contrastive Explanation." *Nous* 32 (1998): 189-204.
- Russell, Bertrand. *The Analysis of the Mind*. New York: MacMillan (1921).
- Rysiew, Patrick. "The Context-Sensitivity of Knowledge Attributions." *Nous* 35 (2001): 477-514.
- _____. "Motivating the Relevant Alternatives Approach." *Canadian Journal of Philosophy* 36 (2006): 259-280.
- Searle, John. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press (1969).
- Schiffer, Stephen. "Contextualist Solutions to Scepticism." *Proceedings of the Aristotelian Society* 96 (1996): 317-333.
- Sosa, Ernest. "On Knowledge and Context." *Journal of Philosophy* 83 (1986): 584-585.
- _____. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13 (1999): 141-151.
- Stanley, Jason. *Knowledge and Practical Interests*. Oxford: Oxford University Press (2005).
- Steup, Matthias. *An Introduction to Contemporary Epistemology*. Upper Saddle River, N.J.: Prentice Hall (1996).
- Stine, Gail. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29 (1976): 249-261.
- Unger, Peter. "An analysis of Factual Knowledge." *Journal of Philosophy* 65 (1968): 157-170.

_____. *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press (1975).

_____. *Philosophical Relativity*. Oxford: Blackwell Publishing (1984).

Vogel, Jonathan. "Are There Counterexamples to the Closure Principle?" In *Doubting:*

Contemporary Perspectives on Scepticism. Eds. M.D. Roth and G. Ross. Dordrecht:
Kluwer (1990): 13-27.

_____. "The New Relevant Alternatives Theory." *Philosophical Perspectives* 13 (1999): 155-
180.

Williams, Michael. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*.
Oxford: Blackwell Publishing (1991).

_____. *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford
University Press (2001).

_____. "Scepticism and the Context of Philosophy." *Philosophical Issues* 14 (2004): 456-475.

Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press (2000).

_____. "Knowledge, Context, and the Agent's Point of View." In *Contextualism in Philosophy:
Knowledge, Meaning, and Truth*. Eds. Gerhard Preyer and Georg Peter. Oxford: Oxford
University Press (2005): 91-114.

Wright, Crispin. "Contextualism and Scepticism: Even-handedness, Factivity and Surreptitiously
Raising Standards." *The Philosophical Quarterly* 55 (2005): 232-262.

Yourgrau, Palle. "Knowledge and Relevant Alternatives." *Synthese* 55 (1983): 175-190.