Georgia State University
Digital Archive @ GSU

Philosophy Theses | Department of Philosophy

6-27-2007

Contextualist Responses to Skepticism

Luanne Gutherie

Follow this and additional works at: http://digitalarchive.gsu.edu/philosophy_theses

Recommended Citation
Gutherie, Luanne, "Contextualist Responses to Skepticism" (2007). *Philosophy Theses.* Paper 22.

This Thesis is brought to you for free and open access by the Department of Philosophy at Digital Archive @ GSU. It has been accepted for inclusion in Philosophy Theses by an authorized administrator of Digital Archive @ GSU. For more information, please contact digitalarchive@gsu.edu.

CONTEXTUALIST RESPONSES TO SKEPTICISM

by

LUANNE GUTHERIE

Under the Direction of Stephen Jacobson

ABSTRACT

External world skeptics argue that we have no knowledge of the external world. Contextualist theories of knowledge attempt to address the skeptical problem by maintaining that arguments for skepticism are effective only in certain contexts in which the standards for knowledge are so high that we cannot reach them. In ordinary contexts, however, the standards for knowledge fall back down to reachable levels and we again are able to have knowledge of the external world. In order to address the objection that contextualists confuse the standards for knowledge with the standards for warranted assertion, Keith DeRose appeals to the knowledge account of warranted assertion to argue that if one is warranted in asserting p, one also knows p. A skeptic, however, can maintain a context-invariant view of the knowledge account of assertion, in which case such an account would not provide my help to contextualism.

INDEX WORDS: contextualism, radical skepticism, theories of assertion, knowledge, relevant alternatives, external world

CONTEXTUALIST RESPONSES TO SKEPTICISM

by

LUANNE GUTHERIE

Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2007

Copyright by
Luanne Su-Lou Gutherie
2007

CONTEXTUALIST RESPONSES TO SKEPTICISM

by

LUANNE GUTHERIE

Major Professor: Stephen Jacobson
Committee: George Rainbolt
Andrea Scarantino

Electronic Version Approved:

Office of Graduate Studies
College of Arts and Sciences
Georgie State University
August 2007

TABLE OF CONTENTS

Introduction **1**
I. The Skeptical Problem Concerning the External World **2**
II. G.E. Moore's Proof of an External World and Contextualist Intuitions **6**
III. Relevant Alternatives and Contextualists **11**
IV. Some Objections **17**
V. A Larger Problem **30**
Works Cited **40**

Introduction

Contextualist theories of knowledge are frequently formulated with the intent of responding to skepticism about the external world. Such theories attempt to establish that the context in which the skeptic makes her argument imposes different requirements for knowledge than that required by the context of our everyday knowledge claims. Thus, the skeptical conclusion cannot be generalized to cover our ordinary knowledge claims, leaving much of our knowledge about the external world unscathed. In what follows, I will argue that this strategy, while insightful in pointing out the role that context plays in determining when we can appropriately say S knows that p, nonetheless fails to show that context determines when we can *truly* say S knows that p.

Following a brief description of what skepticism about the external world amounts to, I will explore several of the more popular contextualist arguments against skepticism and the intuitions upon which such theories are based. Having done so, I will then be in a position to spotlight one of contextualism's greatest strengths: it seems to resemble our normal justificatory practices more so than what is called for by the skeptic. The skeptic can answer this, however, by drawing a distinction between appropriate assertions of knowledge and true assertions of knowledge. Any contextualist theory must ultimately link these two concepts together in order to be successful and thus far, they have failed to do so. But first we must get a handle on the problem that contextualism professes to solve.

I. The Skeptical Problem Concerning the External World

There is a philosophical problem concerning our knowledge of the world around us, or as philosophers like to say, the external world. The problem is that, contrary to popular belief, we might have no such knowledge. We commonly believe ourselves to know many things, such as "I am currently sitting in a chair," or "My husband is watching television outside." The skeptic (the philosopher who argues we have no external world knowledge) will then raise the possibility that we might be dreaming or we might be systematically deceived by an evil demon. If such were the case, then we would have no basis for thinking that our beliefs about the external world are true. I feel the chair under me and see that I am sitting in it. But if I were in a dream, I would have these experiences regardless of whether I am actually sitting in the chair. Thus, such sensory experience cannot serve as evidence one way or another. And if sensory experience does not count as evidence, it is not clear what else is left. Without any other way of distinguishing between a dream world or the real one, we are not justified in any of our beliefs about the external world.

This type of skepticism has often been attributed to Rene Descartes. Sitting by the fire one night, he begins evaluating the many beliefs he had acquired since childhood in an attempt to determine if any of them could count as knowledge. Realizing that most of his beliefs were either acquired through or were in some way dependent on the senses, Descartes attempts to ascertain whether the senses are a reliable source of knowledge. He concludes that they are not:

> "…there are many…beliefs about which doubt is quite impossible…for example, that I am here, sitting by the fire, wearing a winter dressing-

> gown, holding this piece of paper in my hands, and so on.…How often, asleep at night, am I convinced of just such familiar events—that I am here in my dressing-gown, sitting by the fire—when in fact I am lying undressed in bed! Yet at the moment my eyes are certainly wide awake when I look at this piece of paper; I shake my head and it is not asleep; as I stretch out and feel my hand I do so deliberately, and I know what I am doing. All this would not happen with such distinctness to someone asleep. Indeed! As if I did not remember other occasions when I have been tricked by exactly similar thoughts while asleep! As I think about this more carefully, I see plainly that there are never any sure signs by means of which being awake can be distinguished from being asleep."[1]

Descartes evaluates the belief that his circumstances are indeed as they appear to be—that he is in reality sitting by the fire in a dressing-gown, holding a piece of paper in his hand. The possibility that he might be dreaming, however, immediately presses itself upon him. If he is only dreaming that he is sitting by the fire in his dressing-gown, then he is not really sitting by the fire. Despite appearances to the contrary, he is lying in bed. Descartes' belief that he is sitting by the fire would thus be false.

Let us look at this scenario a little more closely. The knowledge claim that Descartes is evaluating is that he is sitting by the fire in his dressing-gown. He justifies this claim by referring to his senses. It certainly appears to him as though he is sitting by the fire. He probably feels the chair under him, sees the fire before him, feels its warmth. He casts doubt on this basis for his claim, however, by introducing an alternative hypothesis that equally accounts for the appearances to his senses. If he were dreaming that he was sitting by the fire instead of actually sitting by the fire, he would experience the same perceptions. Moreover, we cannot discount this alternative hypothesis. Descartes attempts to do so at first by pointing out that his eyes are open, he does not

[1] Cottingham, Stoothoff, and Murdoch (1984), 12-13.

seem to be asleep, and everything looks very distinct. But each of these objections themselves rely on sense experience, which the dream hypothesis has already called into question. Thus, Descartes cannot use sense experience to support the claim that he is sitting by the fire.

This argument can then be generalized to all beliefs about the external world. All beliefs that presuppose the existence of objects would thus lack justification since we support such beliefs by, at the minimum, relying on sense experience. I know there is a tree outside. *How do you know?* I've seen it. *But couldn't you be dreaming right now?* I know that matter consists of atoms. *How do you know?* I learned it in school. *Couldn't the school have been a dream? Couldn't the subject matter in those textbooks have been a figment of your imagination? Couldn't those scientists have been dreaming the experiments?* Descartes' dream hypothesis disallows a huge swath of beliefs from counting as knowledge. Furthermore, the alternative hypotheses are not constrained to dream scenarios. As mentioned above, I could be systematically deceived by an evil demon, or a bodiless brain envatted in a tub of nutrients connected to electrodes that cause my sense perceptions (hereafter referred to as the BIV scenario). Any coherent scenario, regardless of how fantastical, would work as an alternative hypothesis as long as it accounts for our sense experience.

At this point, it is worth mentioning a distinction between two different interpretations of external world skepticism—high-standards and radical.[2] Some contextualists interpret skepticism about the external world as a demand for the

[2] Williams (2001), 73-77.

fulfillment of very high standards of knowledge. We might be justified somewhat for our beliefs since the alternative hypotheses are unlikely, but without the ability to eliminate these scenarios, we are not justified enough for such beliefs to count as knowledge. Radical skeptics, however, contend that we are not justified at all in our beliefs about the external world. They argue that, given all the evidence we will ever have (sense experience), skeptical scenarios are just as likely as our ordinary beliefs. Beliefs about the external world must be justified with appropriate evidence. Such evidence consists of sense experience. But there is no logical connection between sense experience and the objective state of the world, as skeptical scenarios have shown us. Nor can we infer inductively from sense experience to the external world without some access to that world independent of sense experience. Our evidence underdetermines what it is justifiable to believe.[3]

This distinction makes a difference in how successfully certain contextualist theories answer the skeptic. In answering high-standards skepticism, contextualists will tend to maintain that standards shift depending on context. Descartes' investigation takes place in a context where all practical concerns have been placed on hold. In such a context, one must be able to eliminate all alternative hypotheses, regardless of how unlikely they seem. But as explained above, one cannot. So in the context of contemplating the nature of knowledge before bedtime, Descartes indeed did not know whether he was sitting before the fire. As we return to our everyday lives, however, with our more pragmatic concerns, we do generally know what we claim to know. But this

[3] Michael Williams credits this analysis of skepticism to Ayer (1956).

leaves the radical skeptic unscathed, since the radical skeptic does not accept that the alternative hypotheses are unlikely in the first place. Regardless of how low a certain context might render the standards of knowledge, the radical skeptic would still maintain that we are never justified in our beliefs about the external world. I will explain this more fully in IV, but suffice it to say that a contextualist theory that would preserve our knowledge of the external world, even if only in ordinary contexts, would have to address radical, in addition to high-standards skepticism.

Despite these obstacles, contextualist theories do have one great strength. They recognize that, in most contexts, we do have a strong tendency to believe that we know things about the external world. To this we now turn.

II. G.E. Moore's Proof of an External World and Contextualist Intuitions

Skepticism is strangely evanescent. While utterly convincing when one is reflecting in an armchair or in a philosophy classroom, the problem of skepticism readily recedes from the mind in the living of everyday life. Indeed, in ordinary circumstances, the problem of skepticism seems largely irrelevant. In assessing whether Iraq had weapons of mass destruction prior to the war, I doubt the various intelligence officials considered it necessary to first eliminate the possibility they were being systematically deceived by an evil demon. As G.E. Moore notes, I know that I have hands because I am waving them in front of my face.[4] Skeptical possibilities or not, this is sufficient proof for my claim that I have hands. It is as sufficient as proofs get.

[4] Moore (1939), 602-605.

Moore's "Proof of an External World" is a fascinating insistence on the legitimacy of our ordinary justificatory practices. Moore's general strategy is to insist that the way we ordinarily prove things like "I know I have hands" is correct. If the basis for which I can claim to know I have hands is correct, then I must know also that there are external objects (my hands). Thus, skepticism is false. Moore begins by offering his proof of an external world: He waves first one hand, saying "Here is one hand." He waves the other, saying "and here is another." He then concludes that two human hands, and thus external things, exist. To the objection that this was not sufficiently rigorous to effectively counter skepticism, Moore lays out the criteria for a rigorous proof:

(I) The premise must differ from the conclusion.
(II) I must know the premise to be true.
(III) The conclusion must follow from the premise.

He then argues that his proof meets each of these three criteria. (I) The premises of his proof was the waving of his hands and his saying "Here is one hand" and "Here is another." The conclusion differed from the premises by being more general: "Two human hands exist at this moment." (II) Moore insists he knew the premises to be true:

> "I *knew* that there was one hand in the place indicated by combining a certain gesture with my first utterance of 'here' and that there was another in the different place indicated by combining a certain gesture with my second utterance of 'here.' How absurd it would be to suggest that I did not know it, but only believed it, and that perhaps it was not the case! You might as well suggest that I do not know that I am now standing up and talking—that perhaps after all I'm not, and that it's not quite certain that I am!

(III) If Moore's premises are true, then the conclusion that "Two human hands exist at this moment" must follow. Moore reminds us that we customarily do take proofs of this sort as conclusive. For instance, if a question arises as to whether there are three

misprints on a page, the best way to prove that is the case is to turn to the pages where the misprints are and point them out. In a court of law or scientific laboratory, a proof like this would conclusively settle the issue.

A skeptic would insist, however, that Moore does not know his premise—he does not know that he has hands. Moore's response is to insist that he does not have to offer a proof of his premises in order for him to know them to be true. In a separate paper,[5] Moore maintains that he is more certain of *his* premise than he is of the strongest premise of any skeptical argument. But since Moore refuses to supply an argument for his premise, the skeptic will simply say that Moore is being dogmatic.

Barry Stroud gives us an analogy that goes a long way in clarifying what Moore takes himself to be doing.[6] During a weekend party in a country house, a young duke is found stabbed on the far side of a large table in the hall. The butler was with him the entire time of the party except for a few seconds when he left to answer the telephone in the foyer where there were a lot of other people. Trying to solve the crime are an experienced detective and his eager young assistant. The assistant surmises that the murderer must have dashed into the room when the butler answered the phone and dashed out again before anyone noticed. The experienced detective disagrees. He says, "No, we know this table is here and is so large that no one could have come through that door and got around to this side of the table and stabbed the victim and got back out again before the butler returned." Note that the experienced detective does not need to prove to his apprentice that the table was still in existence during the murder for this to be a

[5] Moore, G.E. "On Certainty."
[6] Stroud (1984), 102-105.

successful refutation of the apprentice's theory. If anyone tried to prove the apprentice's point by denying the existence of the table, the detective would not be dogmatic in insisting that he is certain that the table did not disappear for a few seconds at the time of the murder. By dismissing without further justification an explanation that conflicts with what is already known, the detective is following normal justificatory procedures. In this circumstance, no argument for the table's continued existence is necessary. This is what Moore takes himself to be doing. Moore's proof reminds us that it is normal justificatory procedure to evaluate putative knowledge claims against what one already knows.

Here is another way of putting it. Let H designate any of the skeptical hypotheses mentioned above (such as, "I am dreaming" or "I am a BIV (brain-in-a-vat)") and O designate an hypothesis to which we might ordinarily appeal in an explanation of our sense perceptions (such as, "I have hands"). Moore's argument can be recast as the following:

(A) I know that O.
(B) <u>If I know that O, then I know that not-H.</u>
(C) I know that not-H.

The skeptic, on the other hand, argues:

(~C) I don't know that not-H.
(B) <u>If I don't know that not-H, then I don't know that O.</u>
(~A) I don't know that O.

Both arguments, Moore's and the skeptic's, are equally valid. The question then becomes which premise (A or ~C) should one hold onto. Rather than accepting the skeptic's premise (~C), Moore cleaves to what he believes is already known (A). Just as the detective dismisses the apprentice's conclusion because it conflicts with the presence

of the table, Moore dismisses the skeptic's conclusion because it conflicts with the presence of his hands.

And yet, we still feel that Moore's proof does not sufficiently answer the skeptic. This is because, as Stroud observes, Moore's strategy is effective only if the piece of knowledge used to deny a theory is not itself called into question by that theory. He gives another analogy. Following the detective's rebuke, the apprentice decides to be thorough and systematically eliminate anyone at the party who could not have committed the murder. He obtains a list of all party guests from the duke's secretary and shows conclusively that the only person on that list who could have committed the murder is the butler. He then announces to the detective that he knows the butler did it. "No," the detective replies, "that list was simply given to you by the secretary; it could be that someone whose name is not on the list was in the house at the time and committed the murder. We still don't know who did it."[7] In this circumstance, it would be absurd if the apprentice replied, "No, I know the butler did it. Thus, the list must be complete." Such a reply would not be effective against the detective's objection because it relies on a piece of knowledge that has already been called into question by the detective's objection. The detective's objection points out a flaw in the procedure the apprentice used to come to his conclusion. He raises a possibility (that someone could have snuck into the party without the secretary knowing) that is compatible with the apprentice's evidence. Unless the apprentice can eliminate that possibility, he does not know who

[7] Stroud (1984), 108.

committed the murder. He cannot simply insist that he is more certain the butler did it than that there was someone at the party whose name is not on the list.

Similarly, Moore's response to the skeptic relies on knowledge the skeptic has already called into question. Moore cannot refute the skeptic by simply pointing to a piece of knowledge he has about the external world when the skeptic has raised issues concerning the means by which we acquire all such knowledge. The skeptic has taken a position external to that of Moore's, so to speak. Moore takes for granted that his knowledge of the external world is legitimate, as we all commonly do. The skeptic, however, takes a step back and questions whether we are right in doing so. She determines that we are not. It is this that prevents Moore from successfully relying on his tactic of dismissing the skeptic's claim. Even though it disagrees with the rest of his knowledge about the external world, all that knowledge has been called into question as well.

III. Relevant Alternatives and Contextualists

While Moore's proof is ultimately insufficient, he does remind us that, ordinarily, it is legitimate to assess knowledge claims based on an already accepted body of knowledge (such as when the experienced detective refuses to consider that the table might have ceased to exist for a few moments). But we also see that, when presented with skeptical scenarios, the situation becomes like that of the apprentice who claimed to know the butler did the crime but failed to check the completeness of the list. There seems to be something wrong with the procedures with which we acquire knowledge about the external world and so we are prevented from relying on any such knowledge to

counter the skeptic. In an attempt to resolve this issue, certain philosophers—contextualists—contend that arguments for skepticism are effective only in certain circumstances, certain contexts. While conducting philosophy in an armchair by the fire, for instance, skeptical scenarios will defeat all knowledge claims about the external world. But in ordinary circumstances—in the living of our everyday lives—skeptical hypotheses do not function as effective defeaters. In other words, some contexts of external world knowledge claiming might be analogous to the detective saying, "I know the table was there," despite the threat of skeptical scenarios. We really are able to know of the existence of external objects because it agrees with the rest of our knowledge about the external world. Other contexts, however, are more like the apprentice's saying, "I know the butler did it." Skeptical defeaters prevent us from having knowledge of the external world because we cannot rely on any previous knowledge we might have about the external world. But it depends on the context. What this means is that while one may temporarily lose one's knowledge in the context of, say, philosophy, one regains it again in more mundane circumstances.

How is this accomplished? According to some contextualists, the skeptic is able to take away our knowledge of the external world by manipulating the semantic standards for knowledge. The very act of bringing up a skeptical defeater raises the standards for knowledge so high that we are no longer able to meet it. As the standards rise, so too must the strength of one's epistemic position when making a knowledge claim. Under the high standards imposed by the skeptic, one's epistemic position is never strong enough and so our external world knowledge claims are false. In ordinary circumstances,

however, the standards for knowledge fall back down to their normal levels and we again know things about the external world.

An early form of contextualism, the "relevant alternatives" view, maintained that claims to know are made within a context of possible defeaters. But not all defeaters must be eliminated in order to truly claim to know—just the relevant ones, and which count as relevant change from context to context. As the relevant defeaters change, so too does the meaning of the knowledge attribution. Gail Stine makes use of an example by Fred Dretske.[8] Suppose that I am viewing zebras at a zoo. As these creatures have all the stripes and other various fashionings of a zebra, as they are being housed in a cage labeled "Zebra," I believe them to be zebras and can truly say, "I know these animal are zebras." That I cannot eliminate the alternative hypothesis that these animals are mules painted to look like zebras does not make my knowledge claim false. That alternative is irrelevant in this case. As such, when I say, "I know these animals are zebras," what I really *mean* is, "I know these animals are zebras rather than horses or antelopes or some other animal commonly found at a zoo." However, if the context changes such that the painted mule alternative does become relevant, then I cannot truly say, "I know these animals are zebras." In such a context, "I know these animals are zebras" would really mean, "I know these animals are zebras rather than cleverly painted mules," a claim that is clearly false. In this way, relevant alternatives theorists argue that context changes the very meaning of S knows p:

> "So what one *means* when one says that John knows the animal is a zebra, is that he knows it is a zebra, as opposed to a gazelle, an antelope, or other

[8] Dretske (1970).

> 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."[9] (Italics added.)

This is how we can say "S knows p" is true in one context but false in another without contradiction.

For the sake of clarity, some comments about the notion of meaning must be made here. To employ Keith DeRose's terms, "know" can be used with the same "character" but with different "content." While the character of "know" remains the same across contexts, the content can change. In all contexts, the character of "S knows that p" remains roughly, "S has a true belief that p and is in a *good enough* epistemic position with respect to p." What counts as good enough is part of the content of the attribution and depends on certain factors of the context. When an alternative hypothesis is raised and thus made relevant, the content of "know" changes so that the knowledge attributor is now required to be in a stronger epistemic position, one that can discount such hypotheses, in order to count as knowing.

Yet, not all relevant alternative theories are contextualist. We must also pay attention to *how* features in the context affect the range of relevant alternatives.[10] DeRose draws a distinction between "subject factors" and "attributor factors." Subject factors are features of the context having to do with the putative knower's situation. Using Alvin Goldman's barn façade case as an example, a subject in ordinary circumstances can be truly said to know what he is seeing is a barn even if he cannot rule

---

[9] Stine (1976), reprinted in DeRose & Warfield (1999), p. 150.
[10] DeRose (1992).

out the possibility that it is a barn façade. But if the subject factors change such that the subject is now in an area where barn facades are a regular occurrence, then that what he is seeing is a barn façade becomes a relevant alternative. Attributor factors, on the other hand, have to do with the knowledge attributor's situation, with her linguistic and psychological context. If the attributor factors are such that she is in a philosophy class where Descartes' evil demon is being discussed, then this becomes a relevant alternative where ordinarily it would not be considered. Sometimes, the putative knower and the knowledge attributor are the same person, as in the case where I claim, "I know I have hands." When we attribute knowledge to someone else, however, or when one person attributes knowledge to another, then the subject of knowledge and the attributor of knowledge diverge.

This distinction is important in defining contextualist theories because only a contextualist will allow that attributor factors affect the range of relevant alternatives. An invariantist, someone who does not think the standards for knowledge change with the context, can allow that subject factors affect the range of relevant alternatives but not attributor factors. Take Dretske's zebra case, for example. Let's say Bob is looking at the zebras. In ordinary circumstances, I can truly say, "Bob knows those animals are zebras." If, unbeknownst to me, the zoo has a habit of replacing zebras with cleverly painted mules thereby changing the subject factors of the context, then the same attribution of knowledge is false, but not because the content of my knowledge attribution has changed. Indeed, it is difficult to see how I can mean, "Bob knows those animals are zebras rather than antelopes" in one context and "Bob knows those animals are zebras

rather than cleverly painted mules" in a different context when I remain unaware of the change in circumstance. Rather than raising the standards for knowledge, the change in subject factors have weakened my epistemic position. Subject factors change the truth value of knowledge attributions without changing the truth conditions (or meaning) of those attributions. Attributor factors, however, do affect the truth conditions of knowledge attributions. As DeRose notes:

> Attributor factors set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true: They affect *how good an epistemic position the putative knower must be in to count as knowing.* They thereby affect the truth conditions and the content or meaning of the attribution. Subject factors, on the other hand, determine whether or not the putative subject lives up to the standards that have been set, and thereby can affect the truth value of the attribution *without* affecting its content: They affect *how good an epistemic position the putative knower actually is in.*[11]

A contextualist can thus account for the intuitions behind both the skeptic's argument and Moore's argument. Under ordinary circumstances, I can truly say "I know I have hands" because the standards for knowledge are low enough that my being able to wave my hands around puts me in a good enough epistemic position to count as knowing. However, once a skeptical hypothesis has been raised and my range of relevant alternatives changes, my current epistemic position is no longer good enough. I can no longer truly say, "I know I have hands."

[11] DeRose (1992) p. 497.

IV. Some Objections

The strength of contextualism rests in its intuitive appeal. The skeptic's argument and Moore's appear convincing when viewed separately but incompatible when put together. Consider the following four propositions:

(A) I know that O.
(E) I know that O entails not-H.
(D) If I know that p, and I know that p entails q, then I know that q.
(~C) I do not know that not-H.

(A) is the first premise in Moore's argument. (~C) is the denial of the conclusion in Moore's argument and the first premise in the skeptic's argument. (D) is the principle of deductive closure. According to deductive closure, if (A) I know that O, and (E) I know that O entails not-H, then (C) I must also know that not-H. So, if (A), (E), and (D) are true, then (~C) must be false. If (~C), (D), and (E) are true, then (A) must be false. If (A), (E), and (~C) are true, then (D) must be false. (A), (E), (D), and (~C) are jointly inconsistent: we cannot accept all four at the same time. But which do we deny, as all four are also very plausible? This question is what Ram Neta calls the skeptical puzzle.[12] The appeal of contextualism rests on its ability to solve this puzzle in the course salvaging our ordinary claims to knowledge. We can avoid flatly denying either (A), (E), (D), or (~C). While there have been various objections leveled against contextualism, they seem to be answerable in a way that does not threaten this basic contextualist approach. In what follows, I will highlight some of these problems and their corresponding answers, culminating in a contextualist account that seems to adequately address all these problems. This will serve as a prelude to my discussion of an objection

[12] Neta (2003) p. 5.

that does threaten the basic contextualist approach, which I will take up in the next section.

Gail Stine's relevant alternatives theory has been alluded to briefly above. To summarize, Stine maintains S knows that p only if S can rule out all the alternatives to p that are relevant in S's context. Thus, in ordinary circumstances, I do not need to rule out the hypothesis that the zebras at the zoo are really cleverly painted mules because that alternative is not relevant in my current context. However, if a serious zebra-skeptic happens to raise the possibility that what I am seeing may be a cleverly painted mule, then that alternative becomes relevant and I would not know that I am looking at zebras without being able to rule out that alternative.

Ignoring, for the moment, DeRose's distinction between subject factors and attributor factors, the application of Stine's theory to the skeptical puzzle should be clear. In ordinary circumstances, skeptical hypotheses (H's) are not relevant alternatives. Thus, I do not need to rule out the notion that I might be a BIV in order to know that I have hands. And, if I know that I have hands, I also know that I am not a BIV. This is why, in ordinary circumstances, skeptical hypotheses fade into the background. The context is such that they are simply irrelevant. When conducting philosophy, however, the BIV hypothesis does become relevant. Because I cannot rule it out, in that context, I do not know that I have hands. In the philosophical context, skeptical scenarios regain their cogency. In effect, Stine argues that in ordinary contexts, we agree with Moore:

(A) I know that O.
(B) <u>If know that O, then I know that not-H.</u>
(C) I know that not-H.

But in extraordinary contexts when various H's become relevant, we agree with the skeptic:

(~C) I don't know that not-H.
(B) <u>If I don't know that not-H, then I don't know that O.</u>
(~A) I don't know that O.

Because the skeptic's argument applies only in very special circumstances that do not normally apply in our ordinary contexts, the skeptic's conclusion does not generalize to cover the contexts in which we ordinarily make our knowledge claims. My knowledge of external objects is preserved for the most part. Moreover, Stine can account for the plausibility of each of the propositions in the skeptical puzzle by maintaining they are all true but each only in certain contexts.

Stine's account, however, has a problem. She has to maintain not only (~C) I don't know that not-H, in circumstances where H is relevant, but also (C) I do know that not-H, in circumstances where H is not relevant. This is because Stine's relevant alternatives theory is constructed to preserve the deductive closure (D) that creates the entailment in (B). So, in ordinary circumstances, I know that O. To know this, however, I would need sufficient evidence to support my belief that O. But if I had sufficient evidence to know that O, and O entails not-H, I would have sufficient evidence to know that not-H, regardless of circumstance. As argued by Stewart Cohen[13] and further explicated by Ram Neta,[14] Stine is committed to the denial of either:

(D1) If I have sufficient evidence to know p and I know that p entails not-q, then I have sufficient evidence to know not-q.

[13] Cohen (1988).
[14] Neta (2003).

Or

(D2) My knowledge that O requires that my belief that O be based on some sufficient evidence for the proposition that O.

But (D1) and (D2) are just as plausible as (A), (D) or (~C). The denial of either (D1) or (D2) would counteract the intuitive appeal of not flatly denying (A), (D), and (~C).

In trying to resolve these difficulties, Cohen, like DeRose, appeals to the context of the knowledge attributor in addition to that of the putative knower. As discussed in the previous section, Stine's account pits knowledge as dependent on subject factors. The putative knower's context determines whether alternative H is relevant. According to Cohen, the relevance of H is also dependent on attributor factors. More specifically, let us say Jim believes he is looking at zebras because the creatures he is looking at have black and white stripes. That the zebras are really cleverly painted mules would be a relevant alternative if Jim's circumstances (his reasons for believing and other contextual factors) are such that the cleverly painted mule scenario is sufficiently probable, and what counts as sufficient depends on the context *I* am in, since I am the one claiming Jim knows something. In addition, if Jim does not have sufficient evidence to discount the cleverly painted mule scenario, the level of sufficiency again being dependent on my context, it would make the alternative relevant. Since the relevance of an alternative determines whether it needs to be ruled out in order for Jim to know that he sees zebras, and the relevance of an alternative depends on the context of the knowledge attributor, it is possible for one person to truthfully say that Jim knows he sees zebras and another to truthfully say Jim does not, without contradicting the first.

As a solution to the skeptical problem, this means that relative to some contexts, I have sufficient evidence to know that O and thus not-H. Relative to other contexts, I do not have sufficient evidence to know that not-H and so do not know that O. In all contexts, however, I can still maintain (D1) and (D2) in addition to the closure principle (D).

Neta points out, however, that if I want to maintain (D2), I would also need to have sufficient evidence that not-H. If my contention that O requires sufficient evidence, then there does not seem to be any reason why my contention that not-H would not similarly require sufficient evidence. Yet what would count as sufficient evidence for the proposition that not-H, regardless of how low my context sets the standard of sufficiency? Cohen attempts to resolve this problem by claiming that my belief that not-H is "intrinsically rational" and so does not require evidence. If this is the case, then what makes not-H more intrinsically rational than O, such that O requires evidence but not-H does not? This remains unclear.

David Lewis' contextualist account relies on what he calls the Rule of Attention. According to Lewis, S knows that p if and only if p holds in every possibility uneliminated by S's evidence (except for those possibilities we are properly ignoring). Proper ignoring, in turn, depends on attributor factors. While Lewis covers numerous contextual factors that affect the truth-value of knowledge attributions, it is the Rule of Attention that he relies on to resolve the skeptical puzzle. In keeping with the Rule of Attention, a possibility not ignored at all cannot be properly ignored. Thus, if the possibility that H comes up, the context becomes such that I cannot properly ignore it.

Since I have no evidence to discount this possibility, (~C) becomes true. I thus cannot truthfully say that I know O. However, in other contexts where the possibility that H does not come up, I can properly ignore it and so do not need to discount it in order to truthfully claim (A). Thus, I would know also that not-H by (D).

Unfortunately, Lewis' account does not successfully overcome an objection Keith DeRose levied against an earlier version of his argument.[15] According to DeRose, a Rule of Accommodation (which would presumably include the Rule of Attention) cannot adequately explain the persuasiveness of any skeptical argument that relies on a skeptical hypothesis. Imagine two skeptics, he says, one of which claims (~C) I do not know I am not a BIV. The second only dogmatically insists that (~A) I do not know that I have hands, and offers no reason to believe her claim. DeRose believes that a solution to the skeptical puzzle should at the very least be able to explain why the first skeptic (~C) is more convincing than the second (~A). The Rule of Attention is unable to do this. According to this rule, arguments based on the BIV scenario are persuasive because bringing up the BIV possibility creates a context where the BIV hypothesis must be discounted in order for there to be knowledge. But this means that the argument of the second dogmatic skeptic should be just as persuasive as the BIV skeptic. The dogmatic skeptic brings up the possibility that I do not have hands. The Rule of Attention states that, since this possibility has been brought up, I cannot ignore it and so must discount it. I cannot discount it because I cannot discount the possibility that I am not a BIV. Yet

[15] DeRose (1995).

(~A) is not near as convincing as (~C). The Rule of Attention, because it applies equally to both skeptics, cannot account for this.

DeRose offers his own contextualist theory to account for the persuasiveness of effective skeptical hypotheses. Building off of Nozick's account of knowledge, DeRose terms his account the Subjunctive Conditional Account (SCA). According to this account, the belief that I am not a BIV is problematic because I would have this belief regardless of whether it is true. My belief that I am not a BIV does not "track the truth," so to speak. SCA posits that we are strongly inclined to think we do not know p if we think p is a belief we would hold even if it were false. DeRose refers to this as an insensitive belief. In order for p to count as knowledge, p must be sensitive through a sufficient range of possible worlds, called "epistemically relevant worlds," centered around the actual world. What counts as a sufficient range depends on context. In contexts where the standards for knowledge are set relatively low, the range of epistemically relevant worlds is small. In contexts where the standards for knowledge are high, the range increases.

To apply SCA to the skeptical puzzle, DeRose introduces the "Rule of Sensitivity": When it is asserted that S knows (or doesn't know) p, then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which p is false. This is the mechanism by which context raises the standards for knowledge. When the skeptic asserts that I do not know I am not a BIV, she creates a context where the sphere of epistemically relevant worlds enlarges so that it includes at least the closest worlds in which I am a BIV. But since my belief that I am not a BIV is

insensitive, the truth does not track across the epistemically relevant worlds where I am a BIV. Thus, in this context, I do not know that I am not a BIV. Moreover, my epistemic position with respect to the proposition that I am not a BIV is at least as strong as my epistemic position with respect to the proposition that I have hands. According to DeRose, S's epistemic position with respect to p is at least as strong as S's epistemic position with respect to q just in case the following is true across all contexts of attribution: if, relative to that context, S knows that q, then relative to that same context, S knows that p. In other words, DeRose preserves the deductive closure first mentioned above (D). So, in this context, since I do not now that I am not a BIV, I do not know that I have hands. If, however, the subject of my being a BIV does not come up when I assert that I have hands, the context is such that the range of epistemically relevant worlds extend only to the closest worlds where I do not have hands (such as worlds where I am involved in a nasty car accident or some such thing). In this context, the belief that I have hands does track the truth through the range of epistemically possible worlds. And because deductive closure is retained, I can conclude that I am not a BIV.

Now, recall that DeRose fashioned SCA so that it could explain why certain skeptical hypotheses persuade and others do not. Consider the dogmatic skeptic's claim that I don't have hands. Bringing up this hypothesis fails to change the range of epistemically relevant worlds through which my belief must track. The epistemically relevant worlds for this claim are the same as those generated by the proposition that I do have hands, since those worlds already included the closest ones where I do not have

hands. Since my belief tracks through all the epistemically relevant worlds, I am not inclined to be convinced that I don't have hands.

Neta points out a flaw in DeRose's account.[16] To address unconvincing skeptical hypotheses that are nonetheless in accordance with the Rule of Sensitivity, such as I falsely believe that I have hands, DeRose adds that an effective skeptical hypothesis must also explain why I would have such a belief. Neta, however, offers the following counterexamples:

(H1) I falsely believe that I have hands because I've been hit on the head in such a way as to cause me to have that false belief, but not to have any of the sensory experiences and memories that would provide me with reasons for believing that I have hands.

(H2) I falsely believe that I have hands because I've taken a drug that causes me to have that false belief (but not to have any sensory experiences or memories that seem to support that false belief).

(H3) I falsely believe that I have hands because I want to believe it, and my wishes cause me to have that false belief (but not to have any sensory experiences or memories that seem to support that false belief).

Each of these hypotheses offer beliefs that are insensitive across the range of epistemically relevant worlds as well as explanations for why I falsely believe I have hands. Yet each of these three hypotheses remain unconvincing. Knowing what my beliefs should be based on (perceptual evidence, etc.), I would contend that in these situations, I know I don't have hands. We find that the explanations must be of a certain type. The hypotheses above offer merely causal explanations, whereas the explanation for the belief in an effective skeptical hypothesis must appeal to my evidence for that

[16] Neta (2003).

belief. Indeed, we find that effective skeptical hypotheses work by showing that the evidence to which we normally appeal is misleading.

This gives us a way of addressing one other concern about whether contextualist theories really do provide an answer to skepticism. The contextualist strategies above contend that the skeptic achieves her goals by illicitly raising the standards of knowledge. Michael Williams argues that such strategies do not address skepticism of the external world.[17] Rather, they address what he calls high-standards skepticism. High-standards skeptics insist that knowledge requires certainty. They do not deny that justified true belief is possible, but maintain rather that our ordinary standards of justification do not reach those necessary for knowledge. Indeed, DeRose charges the skeptic with holding the standards of knowledge so high that she ultimately changes the meaning of "know" from that of its ordinary use. It is analogous to insisting that a "doctor" is someone who has never had a patient die. If this is what one means by "doctor" or "know", it is difficult to see why one should care about problems concerning a concept no one employs in real life.

Skepticism about the external world, on the other hand, is what Williams calls a radical skepticism. Radical skepticism does not merely say that most of our beliefs are not justified enough to meet some standard. Rather, it denies the very possibility of justified belief. The radical skeptic insists that given all our available evidence, skeptical scenarios are no less likely than what we ordinarily believe. The only basis we have for

[17] Williams (2001), 74-75.

declaring skeptical hypotheses remote is based on the very perceptual evidence that the skeptic has called into question.

We can see how this plays out in DeRose's SCA. Recall that the "Rule of Sensitivity" demands that whenever it is asserted that S knows (or doesn't know) p, we enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which p is false. In making the claim that I know I have hands, it is automatically assumed that the closest epistemically relevant worlds where p is false are those where my hands have been hacked off by an automobile accident or some such relatively plausible scenario. It is certainly not assumed that the closest such world is the BIV world or the evil demon world. But what do I rely on to determine which worlds are the closest ones? If closeness depends on likelihood, then I have no basis for saying that the car accident world is closer than the BIV world. To determine which worlds are more likely, I would have to rely at least partially on the perceptual experiences the skeptic has called into question—perceptual experience from the past, say. The same applies if closeness depends on something like plausibility, or even similarity. I would thus have no basis for saying that my belief that I have hands tracks across all epistemically relevant words, since I would have no basis for saying that a BIV world would not be epistemically relevant.

According to Neta, introducing a skeptical hypothesis does not raise the standards of knowledge so high that we cannot meet them. Rather, as we saw in our discussion of what makes an effective skeptical hypothesis, it disqualifies what we can count as evidence for O. This goes back to our previous discussion of Stroud's objection to

Moore. Recall that Moore's strategy of appealing to our already existing body of knowledge is effective only if that body of knowledge has not been called into question by the skeptic. Otherwise, such an appeal would be no better than the apprentice dogmatically insisting that the list is complete when the master detective points out that the completeness of the list has not yet been confirmed. Skeptical hypotheses are similarly effective only when they call into question the procedures by which we come to the conclusion that O. They do so by disqualifying that which we would normally appeal to justify O—our evidence for O.

Neta maintains that we can recast the skeptical puzzle in the form of an evidentiary puzzle to take this into account:

(A*) I have evidence for O.
(D*) If I have evidence for p, and I know that p entails q, I have evidence for q.
(~C*) I have no evidence for not-H.

Like the skeptical puzzle, each of these three propositions are individually plausible but jointly inconsistent. Unless we wish to deny (D2), (A*) follows from (A). (D*) follows from (D) and (D2). And as for (~C*), what could count as evidence for not-H?

Accordingly, a contextualist theory that addresses radical skepticism would have to provide a solution to the evidentiary puzzle as well as the skeptical puzzle. Neta attempts this by contextualizing what counts as evidence. According to Neta, by raising the possibility that H, the skeptic shifts the conversational context in line with the following rule:

(Neta) When one raises an hypothesis H that is an uneliminated counterpossibility with respect to S's knowing that p at t, one restricts what counts in one's context of appraisal as S's body of evidence at t to

> just those mental states that S has, and would have, at t whether or not H is true.

A hypothesis H is "an uneliminated counterpossibility" with respect to S's knowing that p at t just in case (i) H implies that S doesn't know that p at t and (ii) H and "S knows that p at t" are introspectively indistinguishable for S. Two hypotheses h1 and h2 are "introspectively indistinguishable for S" just in case: (i) If h1 were true, then S would be in mental states M1; (ii) if h2 were true, then S would be in mental states M2; and (iii) there is no difference between M1 and M2 for S that is available to S's introspection. "Context of appraisal" refers to the knowledge attributor's context.

There are several advantages to this view. Like DeRose's contextualism, Neta's account explains why the BIV hypothesis is convincing but "I don't have hands" is not. The BIV hypothesis is an uneliminated counterpossibility to O whereas "I don't have hands" is not. Further, (Neta) recognizes that what counts as evidence is dependent on the conversational context of the knowledge attributor. So, for S to have evidence for p is for S to have evidence that favors p over some alternative(s) that are relevant in the attributor's context. For example, suppose that Richard is drinking Coke. Both Bob and Suzanne ask if Richard has evidence that he is drinking Coke. Bob asks because Richard frequently expresses a preference for Coke over Pepsi but Richard doubts that Bob can taste the difference. Suzanne asks because she does not think Richard is experienced enough in carbonated beverages to be able to taste the difference between Coke and Sprite. It would then be possible for Richard to have evidence that he is drinking Coke in Suzanne's context but not in Bob's. Since drinking Sprite is introspectively distinguishable from drinking Coke for Richard, his drinking Coke counts as evidence

that he is drinking Coke. That he might be drinking Pepsi is not relevant to Suzanne's context. Bob, however, does raise drinking Pepsi as an uneliminated counterpossibility. Because drinking Coke and drinking Pepsi is introspectively indistinguishable for Richard, he does not have evidence that he is drinking Coke in Bob's context.

Furthermore, there is reason to believe that the evidentiary puzzle and the skeptical puzzle are generated by (Neta). Once the skeptic raises the BIV hypothesis, she shifts the context such that the only evidence I can appeal to is just those mental states that are introspectively available to me regardless of whether or not I am a BIV. Thus, what I had originally considered to be my knowledge of the external world no longer counts as evidence. In such a context, (~C*) and (~C) are true. When no such hypothesis has been raised, however, (A) and (A*) are true. I know that I have hands because my being able to wave them around counts as evidence. That I am handed is introspectively distinguishable from non-handedness. (D) and (D*) are true regardless of context.

V. A Larger Problem

When looking at how "knows" is utilized in non-philosophical conversation, it does appear that the evidence necessary to justify a knowledge claim may change from one context to another. Such seeming variability in our attributions of knowledge is part of what makes contextualism so enticing. Unfortunately, it is less clear whether this variability in what appears to be appropriate attributions of knowledge reflect variability in the truth conditions of knowledge.

Here is the problem. In appealing to contextualist intuitions, a contextualist will frequently present two cases of the same individual making or denying a knowledge claim. In Case A, the individual will claim to know p and one intuits that this claim is true. In Case B, the conversational context will have shifted and the same individual will now claim to not know p. And one intuits that this claim is also true. The contextualist will then explain the conflicting intuitions by positing that the change in conversational context in Case B raised the standards of knowledge. Thus, one can both know p in Case A and not know p in Case B because the standards for knowledge have been raised in Case B. Here is an example from DeRose:[18]

Bank Case A. My spouse and I are driving home on a Friday afternoon and we are discussing whether to stop at the bank to deposit our paychecks. Passing the bank, we see that the lines are quite long so I suggest we drive straight home and go on Saturday instead. My spouse says, "Maybe the bank will be closed tomorrow. A lot of banks close on Saturdays." I reply, "No, I know it'll be open. I was there two weeks ago on Saturday. It's open until noon."

Bank Case B. Imagine the same situation as in Case A except we have just written a very large and important check that requires the deposit of our paychecks by Saturday for our account to reach sufficient funds. Again, the lines are long and I suggest we go back on Saturday. My spouse reminds me that banks are closed on Sunday, so the paycheck must be deposited either today or Saturday. And, "Banks do change their hours. Do you know the bank will be open tomorrow?" Though I am just as confident as

[18] DeRose (1992).

in Case A that the bank will be open on Saturday, I reply, "Well, no. I'd better go and make sure."

There is a strong sense here that I know the bank will open on Saturday in Case A but not in Case B. As mentioned above, contextualism provides one way of explaining these conflicting intuitions. An invariantist, however, can object by drawing a distinction between true assertions of knowledge and warranted assertions of knowledge. My claim to know the bank will open on Saturday is false in Case A but warranted, given the low cost of being wrong and time-saving benefits of being right. In that case, it is not knowledge that context affects but warranted assertability.

In "Assertion, Knowledge, and Context," Keith DeRose proposes a solution to this problem.[19] He suggests that we marry contextualism with a theory of warranted assertion called the knowledge account of assertion, thus providing a needed link between knowledge and warranted assertion. It is generally agreed that in order to appropriately assert that p, one must be well-enough positioned with respect to p. The controversy revolves around what counts as "well-enough positioned." According to the knowledge account, one is well-enough positioned only if one knows that p. One can see how this maneuver is meant to work. If there is general agreement that context affects warranted assertibility and warranted assertibility is linked directly to knowledge, then there can be general agreement that context affects knowledge.

DeRose first argues that the skeptic's objection to contextualism cannot be just the bare claim that the contextualist is mistaking the conditions for warranted assertions

[19] DeRose (2002).

with the truth-conditions of knowledge. As explained above, the warranted assertability objector will have to maintain that, despite intuitions to the contrary, S either knows p in both contexts or does not know p in both contexts. One of the contextualist's assessments is wrong. If we have intuitions that S knows that P in one of those contexts, it is only because S has fulfilled the conditions for asserting p appropriately, which, as I argued above, can be different from fulfilling the conditions for actually knowing p. So, in that context, "S knows that p" would be false while "S appropriately asserted p" would be true. But, objects DeRose, this type of maneuver could be used to make all sorts of ridiculous arguments by allowing the proponent of a view to explain away any counterexample. Imagine someone who tries to argue that "S knows that p" has the same truth conditions as "S believes that p." Then, to address the numerous counterexamples that appear to show that the truth of p might also be required, the proponent of this view could say that the truth of p only appears to be a truth-condition of "S knows that p" because it is an assertability condition.

What is needed for a plausible employment of this maneuver is a general conversational rule that the assertability condition can fall under. For instance, if someone asked me for my mother's maiden name and I knew it to be Lee, it would be somehow inappropriate, though true, for me to answer, "It's possible that it's Lee." Such an answer would be inappropriate because an assertability condition of "It's possible that p" might be that you are not in a position to know that P is true. And this, in turn, falls under a general conversational rule that says: When you're in a position to assert either of two things, then, other things being equal, assert the stronger. I am in a position to

assert either the stronger, "It's Lee," or the weaker, "It's possible that it's Lee." If I assert, "It's possible it's Lee," I generate a false implicature that I do not know my mother's maiden name is Lee. The plausibility of this view (as compared to the "S knows that p" = "S believes that p" view above) stems from the fact that the "Assert the stronger" conversational rule applies to propositions of any content, not just the term in question.

In order for the warranted assertability objection against contextualism to be plausible, there needs to be a general conversational rule for the assertability conditions of "S knows that p" to fall under. And indeed there is. There is a general conversational rule that says: One should assert something only if one is well-enough positioned with respect to that proposition to properly assert it. This is a conversational rule that applies to "S knows that p" and how well positioned one must be in to assert "S knows that p" varies based on context. So, one can be well-enough positioned with respect to "S knows that p" to assert it even though one is not well-enough positioned to actually know it.

It is this objection that the knowledge account of assertion is brought in to address in the second stage of DeRose's argument. According to the knowledge account of assertion, one must know that p in order to be well-enough positioned with respect to p to assert it. Hence, by asserting p, you represent yourself as knowing p, though you are not asserting that you know p. Warranted assertions, then, depend on knowledge and exhibit context variability because, as the contextualist will maintain, knowledge exhibits context variability.

Although DeRose's focus is less on arguing for the knowledge account of assertion and more on establishing that the knowledge account, if true, supports contextualism, it is worth addressing some of the more serious objections to this theory of assertion. As I will argue further below, the way in which the knowledge account deals with these objections actually renders its benefits to contextualism moot.

The first counterexample concerns cases in which a person reasonably makes an assertion that is false. Take Goldman's barn façade case mentioned in the previous section. Given how unusual it is to find oneself in barn façade country, the subject can hardly be blamed for asserting that he knows what he is seeing is a barn when it is in fact a barn façade. This certainly seems to be a case in which one does not know p but can make a warranted assertion that p.

The second counterexample concerns cases in which a person reasonably makes an assertion that is true, but that still does not qualify as knowledge. Consider the following example, adapted from Stroud. During wartime, airplane spotters are given a quick course on how to spot different types of airplanes. They learn from their manuals that an airplane with features *x*, *y*, and *w* is an E, and an airplane with features *x*, *y*, and *z* is an F. The completely trained airplane spotter will then announce the presence of an F only if he sees all three features *x*, *y*, and *z*. If he sees only *x* and *y* but cannot determine the third feature, he does not know whether the plane is an E or an F.

Suppose also that, unbeknownst to the airplane spotter, there is another airplane, G, that also has features *x*, *y*, and *z*. The manual never mentioned them because it was unnecessary for the airplane spotters to be able to identify them. While they are as

common as F's, they are not very dangerous. With the exigencies of war, it was more expedient to train the airplane spotters to identify E's and F's but ignore the G's. In such a situation, the airplane spotters do not know an airplane is an F if they see features *x*, *y*, and *z*, even if they are really looking at an F. Just as the spotters do not know an airplane is an F if they see only features *x* and *y* because it might be an E, they do not know an airplane is an F if they see features *x*, *y*, and *z* because it might be a G. But given the exigencies of war, it certainly seems as if the airplane spotter is warranted in asserting he knows an airplane is an F when he sees features *x, y,* and *z.*

In defending the knowledge account of assertion, Timothy Williamson addresses such counterexamples in part by insisting that an assertion made by an asserter who does not know is technically unwarranted.[20] The intuition that the subject can not be held blameworthy for making a reasonable but false assertion or a reasonable true assertion that does not count as knowledge can be explained by invoking rules or conventions of behavior other than those governing solely assertions. In the airplane spotter case, for example, the warranted assertion intuition might derive from moral or behavioral rules that come into play when action is required to save lives. So while the airplane spotter technically did make an unwarranted assertion by falsely representing himself as knowing when he really did not, this is counterbalanced by the positive consequences of his making the unwarranted assertion. The asserter might have been blameless in his act of asserting, but the assertion itself remains unwarranted.

[20] Williamson (2000), Sections 11.4-11.5.

This removes any benefit that the knowledge account of assertion might provide to contextualism. DeRose would like to answer the invariantist by maintaining that contextual factors affect whether an assertion of knowledge is warranted because only knowledge warrants assertion, and contextual factors affect the truth conditions of knowledge. Thus, recalling the bank cases, if I am warranted in my assertion of knowledge in Bank Case A, it is because my assertion of knowledge is true in Bank Case A. The knowledge account of assertion, however, allows the invariantist to explain the intuitions behind the bank cases without agreeing with the antecedent of the preceding conditional. The invariantist can simply say that I am technically not warranted in my assertion of knowledge in Bank Case A. Rather, it only seems so because of extraneous behavioral rules that are not constitutive of solely assertions. In fact, I falsely represented myself as knowing that the bank will be open on Saturday, but am warranted in doing so because of more pragmatic concerns. In this way, the invariantist can maintain that it is neither knowledge nor warranted assertion that is context dependent, but rather behavior in general.

Indeed, it is often the case that we make knowledge claims not for the sake of transmitting knowledge alone but because we think doing so will accomplish some more desired goal like saving time or saving lives. Contextual factors may determine how important it is to actually succeed in transmitting knowledge in order to accomplish that goal. Perhaps this is what underlies the intuitions upon which contextualism is based.

To establish the link between context and knowledge, contextualists need to explain the relationship between attributor factors and the fact that p, such that attributor

factors should play a role in determining whether one is justified in believing p. From the airplane spotter's point of view, any airplane with features *x*, *y*, and *z* are F's. He has no reason to doubt this as this is what he learned in training and this is what it says in the manual. From our point of view, however, we are able to see that there are also G planes that the manual does not cover. From our epistemic position outside of the airplane spotter's situation, we can see that the manual is an inadequate basis for the plane spotter's knowledge claim. The airplane spotter is not himself at fault for thinking he knows a plane is an F. And his ability to identify F's is perfectly adequate for the purposes of the war effort. Nonetheless, the fact that an airplane might be an F or a G is independent from whether the airplane spotter's training was perfectly adequate for the war. It is this independence between the facts of the matter concerning the identity of the airplanes and the purposes for the plane spotter's knowledge claims that proves problematic for contextualism.

Now, I do not believe a contextualist would claim the airplane spotter knows an airplane with features *x*, *y*, and *z* is an F. That it might be a G is a relevant alternative due to subject factors. Just as one would not know what one is seeing is a barn in barn-façade country, one does not know an F from a G when there are also G's flying around. The problem is, rather, that there seems to be a disconnect between the attributor factors in the airplane spotter's context and the facts of the matter. A contextualist would maintain that the exigencies of war (the attributor factors in this case) play a part in determining how good an epistemic position the airplane spotter must be in to know that an airplane is an F. But the exigencies of war have nothing to do with whether the airplane is an F or a G.

The skeptic believes herself to be in a position similar to our position relative to the airplane spotter. The exigencies of war make it such that the airplane spotter thinks he knows a plane is an F when he knows no such thing. Standing outside the airplane spotter's context and removed from the exigencies of war, we can see that the airplane might very well be a G. The skeptic similarly places herself outside any context in which practical needs might drive claims to know. From there, she can see that what we take to be the external world might just be a dream or some other skeptical hypothesis. Just as we can see that the exigencies of war play no part in whether a plane is in fact an F or a G, the skeptic sees contextual factors as playing no part in whether what we sense around us is in fact the external world or merely a dream. As long as there is this independence between: (1) the contextual factors that drive knowledge claims, and (2) the facts around us; contextualism will remain implausible.

Do I believe that the skeptic is right, then, in that we really have no knowledge of the external world? No, I do admit that the intuition behind (A) I know that O is quite strong. And while contextualism might not be the answer to the skeptic, it is nonetheless possible that something else might step into the breach. I am simply arguing that contextualism is not the answer.

Works Cited

Cohen, Stewart (1988), "How to be a Fallibilist", *Philosophical Perspectives*, 2, pp. 91-123.

Cottingham, J., Stoothoff, R., and Murdoch, D. (1984) (eds.), *The Philosophical Works of Descartes*, vol. 2 (Cambridge: Cambridge University Press).

DeRose, Keith (1992), "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research,* 52, pp. 913-29. Repr. in Heumer (2002), 491-506.

_____ (1995), "Solving the Skeptical Problem", *Philosophical Review*, 104, pp. 1-52. Repr. in DeRose and Warfield (1999), 183-219.

_____ (2002), "Assertion, Knowledge, and Context", *Philosophical Review*,

_____ and Warfield, Ted A. (1999), *Skepticism: A Contemporary Reader* (New York: Oxford University Press).

Dretske, Fred (1970), "Epistemic Operators", *Journal of Philosophy*, 67, pp. 1007-1023. Repr. in DeRose and Warfield (1999), 131-144.

Heumer, Michael (2002), *Epistemology: Contemporary Readings* (New York: Routledge).

Jacobson, Stephen (2001), "Contextualism and Global Doubts about the World", *Synthese*, 129, pp. 381-404.

Lewis, David (1996), "Elusive Knowledge", *Australasian Journal of Philosophy*, 74, pp. 549-567. Repr. in DeRose and Warfield (1999), 220-239.

Moore, G.E. (1939), "Proof of an External World", *Proceedings of the British Academy*, 25. Repr. in Heumer (2002), 602-605.

Neta, Ram (2003), "Contextualism and the Problem of the External World", *Philosophy and Phenomenological Research*, 56, pp. 1-31.

Stine, Gail (1976), "Skepticism, Relevant Alternatives, and Deductive Closure", *Philosophical Studies*, 29, pp. 249-261. Repr. in DeRose and Warfield (1999), 145-155.

Stroud, Barry (1984), *The Significance of Philosophical Scepticism* (New York: Oxford University Press).

Williams, Michael (2001), *Problems of Knowledge: A Critical Introduction to Epistemology* (New York: Oxford University Press).

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