

Contextual Knowledge Ascriptions & Non-Contextual Knowledge

Michael Stephen Zarella, Jr.

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in
partial fulfillment of the requirements for the degree of

Master of Arts

In

Philosophy

Joseph C. Pitt, Chair

Philip R. Olson

Walter R. Ott

December 4th 2013

Blacksburg, VA

Keywords: Knowledge, Knowledge Ascriptions, Contextualism, Skepticism

Contextual Knowledge Ascriptions & Non-Contextual Knowledge

Michael Stephen Zarella, Jr.

ABSTRACT

In this paper I pair a contextualist theory of knowledge ascriptions with a non-contextual definition of knowledge, specifically the principle that knowledge is infallible belief. This combination is unusual because proponents of contextualism, understood as a semantic thesis, either do not engage epistemology or suggest that the criteria for knowledge are also context sensitive. In order to sustain the pairing that I suggest, the truth conditions of a knowledge ascription must be distinct from the criteria for knowledge. I believe that this distinction is important and fruitful for two reasons: 1) the distinction allows us to preserve both the principle that knowledge is infallible belief and the conviction that we know a lot; 2) the distinction explains the paradox that certain skeptical arguments are not obviously unsound even though their conclusion does seem absurd. Since I uphold a definition of knowledge that is not context sensitive, my treatment of skepticism is unlike prominent contextualist treatments.

Table of Contents

I. Introduction.....	1
II. The Rock of Fallibilism & the Whirlpool of Skepticism.....	4
III. The Paradox of Skeptical Arguments.....	18
IV. References.....	24

I. Introduction

The relationship between the criteria for knowledge and ascriptions of knowledge is not straightforward. One issue of importance to both epistemology and the philosophy of language is how the criteria for knowledge relate to the conditions in which an ascription of knowledge is true. A reasonable supposition is that ascriptions of the form ‘S knows that P’¹ are true only if S knows that P; however, this is not necessarily the case. *Contextualism* offers a method for evaluating the truth of the ascription ‘S knows that P’ without determining whether S knows that P. Contextualism, as I use the label, refers to a family of semantic theses that contend that the meaning of the word ‘know(s)’ depends on the context of the ascription. According to contextualism, the truth of the ascription ‘S knows that P’ is a function of S’s epistemic position, i.e., S’s evidence or reasons for believing that P is the case, and the contextual meaning of ‘know(s).’² As a semantic thesis, contextualism does not postulate that the criteria for knowledge are context sensitive. A contextualist could combine or merge the semantic thesis with the epistemic thesis that the criteria for knowledge are context sensitive, as Cohen and DeRose propose.³ Yet, a contextualist could maintain that the conditions in which a knowledge ascription is true are context sensitive even though the criteria for knowledge are not context sensitive, so long as the semantic thesis works with the epistemic thesis. According to the latter pairing, the ascription ‘S knows that P’ could be true even if S’s belief that P does not constitute knowledge.

I believe that a contextualist account of knowledge ascriptions is correct. I also believe that the criteria for knowledge are not context sensitive; specifically, I maintain that knowledge is infallible belief. Combining a contextual theory of knowledge ascriptions with non-contextual

¹ The ascription ‘S knows that P’ can represent either a 1st person or 3rd person ascription. Furthermore, the ascription ‘S knows that P’ could be expressed in speech, writing or thought, either in its exact logical form or not.

² DeRose also insists that ‘S knows that P’ is true only if S has a true belief that P. This point is discussed on p. 17.

³ Cohen 1999 p. 60; DeRose 1995, 2009 p. 14

criteria for knowledge is the task at hand. This combination is unusual because proponents of contextualism either do not engage epistemology or suggest that the criteria for knowledge are also context sensitive. The combination I suggest will prove fruitful for two reasons: 1) We can preserve both the principle that knowledge is infallible belief and the conviction that we know a lot; 2) We can explain the paradox that certain skeptical arguments are not obviously unsound even though their conclusion does seem absurd. To begin this project I will further describe contextualism in contrast with invariantism.

Contextualism & Invariantism

Peter Unger introduced the terms *contextualism* and *invariantism* to signify competing positions regarding the nature of ascriptions involving certain terms such as ‘flat’ and ‘know(s).’⁴ The issue of contention is how to account for the variance in our willingness to ascribe flatness or knowledge. Invariantism regarding ‘know(s)’ is the position that the meaning of ‘know(s)’ does not depend on the context of ascription. Since the meaning of ‘know(s)’ is invariant, instances of the ascription ‘S knows that P’ express the same proposition regardless of the ascriber’s context. Accordingly, the truth-value of those ascriptions does not vary with the context of ascription. Even though the meaning of the word ‘know(s)’ is fixed, an invariantist can acknowledge that the *appropriateness* of our knowledge ascriptions may depend on the circumstances. This acknowledgement allows for the possibility that a knowledge ascription is false, but appropriate nonetheless. For instance, it may be appropriate for me to state that ‘S knows that P’ during a casual conversation with friends, even though *strictly speaking*, my ascription is false. One strategy for accounting for this phenomenon is the so-called Warranted

⁴ Unger 1984 p. 6

Assertability Maneuver. The details of such maneuvers vary but the general strategy is to appeal to pragmatic rather than semantic factors to explain the appropriateness of such assertions.

Conversely, contextualism contends that the proposition expressed by the ascription ‘S knows that P’ depends on the context of ascription because the meaning of the word ‘know(s)’ depends on the context of ascription. Factors such as the intentions, purposes, expectations and presuppositions *of the person ascribing knowledge* determine the content of the ascription. For example, the ascription “Jacqui knows that she is looking at a red barn” can express different propositions if thought or spoken by ascribers with different expectations or presuppositions regarding the presence of barn façades. It is important to note that the facts that comprise the subject’s circumstances, i.e., whether there is in fact a barn or merely a barn façade, or even the number of barn façades are in the area, etc., only influence the meaning of a knowledge ascription insofar as those facts contribute to the context of ascription—the intentions, purposes, expectations and presuppositions of the person ascribing knowledge.

Since the proposition expressed by ‘S knows that P’ is context sensitive, the truth of ‘S knows that P’ is context sensitive. One way to interpret the context sensitivity of ‘know(s)’ is to follow Cohen and DeRose in their suggestion that the meaning of ‘know(s)’ is determined by the relevant epistemic standard, which may be *higher* or *stricter* in some contexts than others.⁵ Another way to interpret this phenomenon is through the idea of *relevant alternatives*. According to this view, the meaning of ‘know(s)’ is determined by the alternatives to P that S has (and has not) ruled out, and may vary because only certain alternatives are relevant to certain contexts. Although a semantic analysis of any particular scenario would be complex in its details, contextualism offers a straightforward account of our conversational practices: the

⁵ Cohen 1999, DeRose 2009

appropriateness of our knowledge ascriptions varies because the *meaning* of a knowledge ascription is context sensitive and generally a knowledge ascription is true given its context.

Since Unger introduced the debate between contextualism and invariantism, nuanced positions have been developed in support of each general position.⁶ Although the debate between contextualism and invariantism regarding ‘know(s)’ is ongoing, I believe that a contextualist account is correct. I side with contextualism because it does seem that we *mean* something different by ‘know(s)’ depending on the context. Moreover, I believe that contextualism provides a good account for the supposition that our knowledge ascriptions are generally true. However, I am not directly concerned with fully explicating and defending contextualism against invariantism in this paper. Instead, I am interested in arguing for the combination of a version of contextualism with a non-contextual definition of knowledge, namely the principle that knowledge is infallible belief. Likewise, my aim is not to settle the debate as to whether knowledge can be fallible or must be infallible, but I do want to motivate the position that we should define knowledge as infallible belief.

II. The Rock of Fallibilism & the Whirlpool of Skepticism

David Lewis provocatively wrote that epistemologists “are caught between the rock of fallibilism and the whirlpool of skepticism. Both are mad!”⁷ Lewis explained that our predicament arises because it seems to us that knowledge is, by definition, infallible belief. Lewis is appealing to the idea that for S’s belief that P to be knowledge, S must have ruled out all the alternatives to P. Although, if a belief must be infallible to be knowledge, then the skeptic is correct to point out that few, if any, of our beliefs about the world constitute knowledge. Yet,

⁶ A notable contribution is Jason Stanley’s Interest-Relative Invariantism. *Knowledge and Practical Interests*, 2005

⁷ Lewis 1996 p. 550

this skeptical conclusion strikes us as mad because it seems to conflict with our conviction that we know a lot. To be sure, we take ourselves to know many things. We know that George Washington was the first president of the United States, we know that the earth is more than 10,000 years old, and we each know that we have two hands, but those beliefs are not infallible. The skeptic is adept at suggesting alternatives that (1) entail that our belief is false and (2) cannot be ruled out because the alternative is consistent with our reasons or evidence for holding the original belief. So, perhaps we should accept that a fallible belief can be knowledge; perhaps S can know that P even if S has not, or even cannot, rule out a certain alternative in which P is not the case? Yet, Lewis resists: “If you are a contented fallibilist, I implore you to be honest, hear it afresh. ‘He knows, yet he has not eliminated all possibilities of error.’ Even if you’ve numbed your ears, doesn’t this overt, explicit fallibilism still sound wrong?”⁸ Doesn’t it sound mad?

I recognize Lewis’s point concerning the tension between the principle that knowledge is infallible belief and the conviction that we know a lot. However, unlike Lewis, I maintain that an acceptable solution will be one that avoids Lewis’s precarious *Psst*. Lewis submitted that “S knows that *P* iff *S*’s evidence eliminates every possibility in which not-*P* – *Psst*! – except for those possibilities that we are properly ignoring.”⁹ According to this account of knowledge, S need not eliminate *every* possibility in which not-*P* since some possibilities are properly ignored. Which possibilities are properly ignored is determined by a set of rules that Lewis proposes. Lewis resorts to the *Psst* when demarcating the alternatives that S must rule out from the alternatives that S can properly ignore because of Lewis’s “Rule of Attention.” According to the “Rule of Attention,” an alternative to P is not properly ignored if it is not ignored; in other words, if an alternative to P is *attended to* it becomes a relevant alternative, no matter how far-fetched

⁸ Lewis 1996 p. 550

⁹ Lewis 1996 p. 554

the alternative is considered to be.¹⁰ On the other hand, it is a feature of Lewis's theory that *unattended to* alternatives to P can be properly ignored, so long as they do not fall within the domain of one of his other rules. Hence, the *Psst* is Lewis's way to suggest that some alternatives are properly ignored without attending to them. I do not support Lewis's qualification that some alternatives are properly ignored, at least not with respect to the question of whether S's belief that P constitutes knowledge. Lewis's account of knowledge only requires that it *seem* to S that he or she has ruled out the alternatives to P. Since *unattended to* alternatives are properly ignored, Lewis's account permits the possibility that S's belief is fallible. Before I propose my own solution for accommodating both the principle that knowledge is infallible belief and the conviction that we know a lot, I will try to motivate the view that knowledge is infallible belief.

Knowledge as Infallible Belief

Infallibilism and fallibilism alike insist that S's belief that P is knowledge only if P is true. The disagreement between infallibilism and fallibilism concerns S's epistemic position regarding P. Cohen provides a straightforward characterization of the fallibilist's position: *S can know P on the basis of R even if there is some alternative to P compatible with R.*¹¹ In Cohen's formulation, R stands for S's epistemic position regarding P (which for Cohen should be construed in terms of *reasons*). Although fallibilism insists that S's belief that P is knowledge only if P is true, fallibilism does not require that *S's epistemic position guarantee that P is true*—from S's perspective there may remain alternatives to P that are not ruled out. Hence, for fallibilism, the criterion that P is true is distinct from the “justification” criterion, broadly construed. According to fallibilism, S's belief that P is knowledge only if P is true *and* only if S's

¹⁰ Lewis 1996 p. 559

¹¹ Cohen 1999 p.58

epistemic position regarding P meets certain conditions or standards, but S need not believe that P *because* P is true. It is the gap between S's epistemic position regarding P and the truth of P that infallibilists consider unacceptable because *from S's perspective, P may be false!*

According to infallibilism, the truth of P is not just a necessary criterion for knowing that P. For knowledge, it is not sufficient that S believes that P *and* P is true; rather, S must believe that P *because* P is true. Importantly, this principle did not originate unmotivated in some philosopher's head; there is strong intuitive force behind it. Gettier's famous paper capitalized on precisely this intuition. Gettier describes scenarios in which S believes that P, S is in a strong epistemic position regarding P, *and* P is true, yet Gettier points out that S does not know that P and almost everyone agrees. Gettier's analysis explains that his subjects' beliefs do not constitute knowledge because the beliefs are not held *because* P is true.¹² In Gettier's particular scenarios and "Gettier" cases in general, S's epistemic position consists of reasons and/or evidence for P, but the reasons and/or evidence are not actually based on P, but rather an alternative to P. Gettier's paper is provocative precisely because it elicits our intuition that S knows that P only if S believes that P *because* P is true.

The insight that S's belief that P constitutes knowledge only if S believes that P *because* P is true can be stated formally as the requirement that S's belief must be *sensitive* to the truth of P: *S's belief that P constitutes knowledge only if S would not believe that P if P were false.* If S does in fact believe that P because P is true, then it is the case that S would not believe that P if P were false. Fred Dretske and Robert Nozick each developed a theory of knowledge around this

¹² Gettier analyzes the first of his two cases (he does not offer as much detail in analyzing the second case) by saying: "But it is equally clear that Smith does not know that (e) is true [that the man who will get the job has 10 coins in his pocket]; for (e) is true in virtue of the number of coins in Smith's pocket, while Smith does not know how many coins are in Smith's pocket, and bases his belief in (e) on a count of the coins in Jones's pocket, whom he falsely believes to be the man who will get the job." In other words, Smith has a true belief but does not hold the belief because it is true; rather, he holds the belief for other reasons. Gettier 1963 p. 122

basic idea; however, neither was motivated to insist that knowledge is infallible belief.¹³

Accordingly, neither required *complete* sensitivity. In order to capture the principle that knowledge is infallible belief through the notion of sensitive belief, the construal of sensitivity cannot admit of degrees; a belief is either sensitive or it is insensitive. Insofar as S's belief is sensitive to P, i.e., S would not believe that P if P were false, then S has ruled out all the alternatives in which P is false, thus satisfying the principle that knowledge is infallible belief: *S's belief that P is knowledge if and only if S has ruled out all the alternatives in which ~P*. The sensitivity criterion ensures that S has an infallible belief by bringing together the traditional criteria of knowledge: S's belief that P is knowledge only if S's epistemic position guarantees that P is true.

What begins as an insight that knowledge requires that S believe that P *because* P is true, when applied strictly, severely limits what we know. Consider my belief that I have hands. I might suppose that my belief is sensitive because if I did not have hands, i.e., P is false, *because my hands were amputated*, then I would not believe that I have hands. But that alternative is just one of many possible alternatives to my belief. Suppose that I do not have hands *because I am a bodiless brain in a vat (BIV)* that is stimulated to believe that I have a body with hands, etc. If that alternative is the case, then I would still believe that I have hands. Hence, the status of my belief is that I *would* believe that P even if P is false and so my belief is not sensitive to the truth of P. Instead, my belief is sensitive to my experience of having hands which, according to the BIV hypothesis, is consistent with me not having hands. The point of this example generalizes to the view that adhering to the sensitivity principle requires that S rule out *all* the alternatives in

¹³ Dretske 197; Nozick 1981

which P is false because establishing that a belief is sensitive to P requires establishing that S's belief is not instead sensitive to an alternative in which P is false.

If I do not know that I have hands because that belief is not sensitive, what do I know? The consequence of infallibilism is skepticism and Cohen acknowledges that this prospect leads many, perhaps most, epistemologists to accept fallibilism.¹⁴ According to Cohen, whose appeal to contextualism relates to his defense of fallibilism, "the motivation for fallibilism stems from the widely held view that what we seek in constructing a theory of knowledge is an account that squares with our strong intuition that we know many things."¹⁵ He explains that even though the view that knowledge is infallible belief is an attractive principle in the abstract, its appeal is not sufficient to withstand the overwhelming case against it from our intuition that we know a lot.¹⁶

I respect Cohen's and Lewis's concern that infallibilism is in tension with the conviction that we know a lot, but I disagree with their acceptance of fallibilism. Like Lewis, I propose a solution that attempts to accommodate both the principle that knowledge is infallible and the conviction that we know a lot. My argument is that our conviction that we know a lot is a judgment that our knowledge ascriptions are generally true, not a judgment that our beliefs constitute knowledge. If that is so, then the supposition that infallible-knowledge is in tension with our conviction that we know a lot presumes that the ascription 'S knows that P' is true only if S knows that P; however, this is not necessarily the case. A contextualist account of knowledge ascriptions can provide the tools for determining that the ascription 'S knows that P' is true even if S's belief that P does not constitute knowledge. Hence, we can account for the conviction that we know a lot while admitting that our beliefs do not constitute knowledge.

¹⁴ Cohen 1999 p.58

¹⁵ Cohen 1999 p.58

¹⁶ Cohen 1999 p.58

“We know a lot.”

We take ourselves to know a lot. We take ourselves to know that the earth is not flat, to know where we live and, dammit, we know that we have hands. Moreover, we frequently ascribe knowledge to others. But what does this conviction signify? The immediate answer is that our conviction signifies the judgment that when we take ourselves and others to know various Ps, it is true that we know those Ps, but that answer is ambiguous. The widespread conviction that we know a lot is *not* the conclusion that, after analysis of our beliefs with respect to the presumed correct definition of knowledge, most of our beliefs are indeed knowledge. Of everyone who is confident that they know a lot, relatively few have any experience affirming a definition of knowledge and then assessing which beliefs constitute knowledge. Therefore, the conviction that we know a lot does not seem to signify that type of judgment. Nonetheless, the conviction does seem to represent the judgment that our knowledge ascriptions are generally true—that when we claim to know that P, or we claim that someone else knows that P, etc., we are speaking the truth. Certainly our knowledge ascriptions can be mistaken, but we take ourselves to know a lot because we believe that generally when we ascribe knowledge to ourselves and to others our ascriptions are true.

Since our conviction that we know a lot signifies the judgment that our knowledge ascriptions are true, but does not signify an analysis that concludes that our beliefs constitute knowledge, we can account for this conviction with a contextualist account of knowledge ascriptions, which affirms that indeed our knowledge ascriptions are generally true. According to the core tenant of contextualism, the truth of a knowledge ascription is a function of S’s epistemic position regarding P and the contextual meaning of ‘know(s)’ – the question of whether S’s belief that P constitutes knowledge is a distinct question. Although the truth of a

knowledge ascription is distinct from the question of whether a belief constitutes knowledge, a knowledge ascription is nonetheless a report that S's belief constitutes knowledge. Therefore, the proposed account of knowledge ascriptions must accord with the proposed account of knowledge. Although I do not accept Lewis's Relevant Alternatives theory of knowledge, I do consider the notion of relevant alternatives to be central to understanding the relation between contextual knowledge ascriptions and knowledge.

Knowledge & Contextual Knowledge Ascriptions

Determining whether a belief is knowledge is an issue that is independent of contexts of ascription. Determining whether a belief is knowledge involves only an assessment of S's epistemic position regarding P and the definition of knowledge. Since knowledge requires that S rule out *all* the alternatives to P, the factors that comprise contexts of ascriptions, i.e., intentions, expectations and presuppositions, are irrelevant. Conversely, knowledge ascriptions are essentially context sensitive. An ascription of knowledge is a report that S knows that P *by a person in a particular context*. In other words, a knowledge ascription is a report that S's belief that P *counts as knowledge according to the ascriber*. In order to determine whether S knows that P simpliciter, i.e., not according to any particular person, we must abstract from the context of ascription. In doing so, the issue transforms from an issue regarding knowledge ascriptions to an issue about knowledge.

The common element between (α) contextual reports that a belief is knowledge and (β) determinations that a belief is knowledge (abstracted from any particular context) is the notion of sensitive belief. Knowledge ascriptions report that S's belief *is* a sensitive belief – though what the ascriber considers to be a sensitive belief is influenced by the context of ascription. Factors of

the ascriber's context, such as their intentions, purposes, expectations and presuppositions, influence what the ascriber means by 'knows' and in turn whether the ascriber will count S's belief as knowledge. The ascriber's context influences the meaning of 'knows' by limiting the extent that S's belief must be sensitive to P; specifically, by limiting the alternatives that S must rule out in order for the ascriber to consider S to know that P. To put the point another way, a knowledge ascriptions is a report that '(for all intents, purposes, presuppositions, expectations, etc.) S's belief is sensitive to P.' The set of factors comprising the context of ascription are parenthetical because *usually* they are implicit in the ascription, though, as I will discuss, not necessarily. Insofar as S's epistemic position satisfies the contextual meaning of 'know(s),' i.e., S's epistemic position rules out the requisite alternatives to P, the ascription 'S knows that P' is true. Nonetheless, if we abstract from the context of ascription, i.e., we do not excuse the factors that limit the alternatives that must be ruled out, it is likely that S's belief that P is not (strictly) sensitive, and hence is not knowledge.

Since knowledge ascriptions report knowledge with varying degrees of sensitivity, a fair question asks: "Why not define knowledge as a belief that is *sensitive enough*?" There are two related responses to this question. First, that approach would define knowledge based on a cursory description of how the word 'know(s)' is used. That is a legitimate position to take, especially if the proponent of that position argues that there is further analysis to be given. However, that is not my position. Certainly the way that the word 'know(s)' is used should inform its definition, but ultimately I believe that the definition of knowledge should be more than a description of its use. As discussed above, I submit that an analysis of purported instances of knowledge leads to defining knowledge as sensitive belief. Moreover, and this is the second point, the idea that knowledge is a (strictly) sensitive belief is present in our ascriptions of

knowledge. A person ascribing knowledge is claiming that S's belief that P *is* sensitive, though, the factors comprising the context of ascription adulterate the notion of sensitivity.

A potential worry is that contextualism introduces a hierarchal conception of contexts of ascription such that our knowledge ascriptions are never *really true* since there is a stricter standard which carries more authority in which the ascription is false. One way to imagine this view is to envision contexts of ascriptions as a set of concentric circles: the outermost circle is the most stringent context because all alternatives are relevant within that context. However, a hierarchical conception of contexts is not a necessary feature of contextualism. I suggest a model in which contexts of ascription are not hierarchical but different and independent. Michael Williams's *substantive contextualism* argues a similar point by rejecting the hierarchy of epistemological kinds;¹⁷ though, Williams's argument is a claim about the nature of knowledge and my point is a claim about knowledge ascriptions.

The model I am suggesting can be captured through the image of a pepperoni pizza. If we imagine the entire pie as containing all the alternatives to P, then each piece of pepperoni should be understood as a context within which only certain alternatives to P are relevant. Just as individual pepperoni are distinct, each context of ascription is distinct from every other. The meaning of 'know(s)' varies from one context to another but the various meanings of 'know(s)' are not relative to each other. In other words, it is not that the meaning of 'know(s)' involves *higher* or *stricter* standards in certain contexts compared to others; it's just that *different* alternatives are relevant in different contexts. Contexts may overlap, given that certain alternatives are relevant to both, but each context retains its autonomy since contexts are not *within* each other in the sense that concentric circles are within each other.

¹⁷ Williams, Michael. *Unnatural Doubts*, 1991. Particularly chapter 3 "Epistemological Realism"

Importantly, there is no *ultimate* context of ascription. If a person considers all possible alternatives to be relevant for determining whether S knows that P, then another way to describe that person's position is to say that he or she wants to determine whether S knows that P *independent* of any set of intentions, purposes, expectations and presuppositions that may limit the alternatives that are relevant. In other words, he or she wants to determine whether S knows that P independent of a context of ascription. A context of ascription is a set of intentions, purposes, expectations, presuppositions, etc. that comprise the perspective of the ascriber. If every alternative is on the table then the issue is no longer, *does S know that P from this perspective*, this issue has become *does S know that P (simpliciter)*. The latter issue is concerned with whether S can rule out all the alternatives to P, i.e., whether S's belief is knowledge, not whether a person in a particular context would ascribe knowledge.

An Illustrative Scenario

In order to illustrate my view and draw out some details it will be helpful to consider a scenario based upon the barn façade case popularized by Alvin Goldman. Suppose I am driving through the country side with my sister and niece and we are playing a game of sorts. My niece points at various objects and asks, "Do you know what *that* is?" I reply, "Yes, that is a red barn."¹⁸ I ascribe knowledge to myself because, without much reflection, I take myself to have a sensitive belief – to know – that she is pointing at a red barn. I can see that it is a red barn and I believe that if it was not a red barn, perhaps because the structure was as house or was blue, I would not believe that it is a red barn. I consider my belief to be sensitive but of course my belief is not strictly sensitive. My belief is based on my visual perception of a barn such that I would

¹⁸ Here is an instance in which a knowledge ascription has the logical form 'S knows that P' without explicitly being verbalized in that form.

believe that the structure is a barn even if it is just a barn façade. Nonetheless, *my ascription is true* because my epistemic position does satisfy the contextual meaning of ‘know.’ My purpose in claiming to know that my niece is pointing at a red barn is to identify a barn, to identify the color red and other similar facts I believe she might be interested in. Also, I presuppose that no one has erected a barn façade in the country-side that I am driving through. Importantly, my intention is *not* to give the impression that I can tell the difference between a barn and the façade of a barn from 100 yards while driving 40mph, especially since I have no reason to expect that it might be a barn façade. For these reasons among others, when I ascribe knowledge to myself I do not *mean* that I have ruled out the possibility that it is a barn façade. What I mean is relatively modest, something like: I have good vision and an unobstructed perception of a red barn. According to *this* meaning of ‘know’ my ascription is true even though my belief does not qualify as knowledge since my belief is not sensitive.

Meanwhile, suppose that unbeknownst to me my sister did realize that I was driving through an area with many barn façades amongst many actual barns. As Goldman’s scenario suggests, our intuition is that my sister should not ascribe knowledge to me. Why not? My sister should refrain from ascribing knowledge to me because my epistemic position cannot satisfy her meaning of ‘know.’¹⁹ Remember, it is the ascriber’s context that determines the relevant meaning of ‘know(s)’ (p. 3). The context of my ascription was similar to the context of my sister’s ascription but with one important difference: my sister expects that many of the apparent barns are actually barn façades. This expectation influences what the word ‘know’ means *according to my sister*. Since my sister recognizes that I am unaware of the prevalence of barn façades, she realizes that I have not ruled out the possibility that my niece is pointing at a barn

¹⁹ I presume the principle that an ascriber does not ascribe knowledge unless he or she considers the ascription true.

façade, and so it is unlikely that she would ascribe knowledge to me. However, if she did ascribe knowledge to me her ascription would be false because my epistemic position cannot satisfy what she means by ‘knows.’

The point is not simply that my sister is aware of an alternative that I, as the subject of her ascription, have not ruled out. The point is that my sister is aware of an alternative that has influenced her meaning of ‘knows’ and that I have not ruled out that alternative. In order to appreciate this point, let me amend the factors comprising my sister’s context. Suppose that as I claim to know that my niece is pointing at a red barn, my sister recalls reading years ago about a single barn façade that was erected in another country. This memory does make the possibility that my niece is pointing at a barn façade salient to my sister; however, the saliency of this possibility has not influenced the meaning of ‘know’ that my sister is employing. Although my sister is aware that a barn façade was erected somewhere, years ago, that fact has not caused my sister to expect there are any barn façades anywhere near this countryside. Accordingly, she continues to presuppose that all the structures that look like barns are indeed barns. Therefore, her intended meaning of ‘knows’ will not require that I rule out the possibility that my niece is pointing to a barn façade. In this case, even though I have not ruled out the possibility that I am looking at a barn façade, my sister could *truthfully* ascribe knowledge to me since her ascription does not mean that I have ruled out that possibility.

I would give the same analysis to Moore’s famous, or perhaps infamous, response to skepticism. Even though the possibility that the external world is an illusion is salient to Moore, he insists that he knows that he has two hands and that there is an external world. Moore defends

his conclusion by asserting that *he can know things that he cannot prove*.²⁰ I do not agree with Moore that those beliefs constitutes knowledge, but I do suggest that what Moore meant by ‘know’ was that he need not rule out certain alternatives for a belief to count as knowledge. According to what Moore meant by ‘know’ and Moore’s epistemic position, i.e., his perception of his hands, Moore’s claims to know that he has two hands and that there is an external world are true. Of course, Moore’s beliefs do not constitute knowledge since, as Moore admits, there are alternatives he does not rule out.

There is one final point to make about these related scenarios. Notice that in each scenario I argued that the truth of the ascription, as well as the issue of whether my belief constitutes knowledge, can be determined without determining whether I was looking at a barn or a barn façade. Certainly, my belief that my niece is pointing at a barn is not knowledge if she is in fact pointing at a barn façade; however, in the scenarios presented it can be determined that my belief failed to be sensitive due to the un-eliminated *possibility* that I was looking at a barn façade. This point can be generalized to other scenarios: often the issue of whether S’s belief that P constitutes knowledge can be determined without determining whether P is true because an un-eliminated alternative to P is sufficient to deny that S’s belief that P is knowledge.

Similarly, the truth of P is not directly relevant to the truth of the ascription ‘S knows that P.’ It may be tempting to suppose that the truth of P is essential to the truth of ‘S knows that P’ because if P is false then ‘S knows that P’ is *wrong*, and a true knowledge ascription cannot be wrong! Yet, it is important to realize that the proposition expressed by ‘S knows that P’ does not usually entail that P is true. For instance, my self-ascription that I know my niece is pointing to a red barn expresses the proposition *I have good vision and an unobstructed perception of a red*

²⁰ Moore 1962 p. 148

barn, but that proposition does not entail that there is a red barn. Hence, it is possible that the ascription 'S knows that P' is true even though P is false (though unbeknownst to the ascriber).

III. The Paradox of Skeptical Arguments

A powerful argument for Skepticism takes the following form:

- 1: P only if $\sim H$
- 2: S knows that P only if S knows that $\sim H$
- 3: It is not the case that S knows that $\sim H$
- C: It is not the case that S knows that P

For example:

- 1: S has hands only if it is not the case that S is a BIV
- 2: S knows that S has hands only if S knows that it is not the case that S is a BIV
- 3: It is not the case that S knows that it is not the case that S is a BIV
- (S does not know that S is not a BIV)
- C: It is not the case that S knows that S has hands
- (S does not know that S has hands)

Often, the reaction to such arguments is paradoxical. On the one hand, the argument is to some degree compelling because the argument seems sound, or at least is not obviously unsound. On the other hand, the conclusion of the argument seems not just false, but *obviously* false. If the argument is not obviously unsound, why does the conclusion seem obviously false? Even if upon analysis the argument is determined to be unsound, an explanation is needed for why the argument initially seemed plausible. We can make sense of this paradox by appealing to the distinction between the truth of a knowledge ascription and the question of whether a belief constitutes knowledge.

The conclusion seems obviously false because we have a strong conviction that we know we have hands. However, as discussed above (p.10), this conviction does not signify the judgment that our beliefs satisfy the criteria for knowledge; rather, our conviction represents the judgment that our knowledge ascriptions are true. So, whereas the skeptical argument is claiming that our belief that we have hands does not constitute knowledge, our conviction that we do know that we have hands represents the judgment that the ascription ‘I know that I have hands’ is true. Hence, the paradox can be resolved by pointing out that each aspect of our reaction is responding to a separate issue. On the one hand, the argument is compelling *as an argument about knowledge* because the argument is valid and the premises are plausibly true; on the other hand, the conclusion seems obviously false because in most contexts the claim ‘I do not know that I have hands’ is indeed false. The paradox arises because we take the skeptic’s point about knowledge to be a challenge to the truth of the ascription ‘S knows that S has hands’ – but that is a mistake. It is not obvious to us that each reaction is responding to a different issue because it is natural to assume that the ascription ‘S knows that P’ is true only if S knows that P, in which case the skeptical conclusion *would* challenge the truth of the corresponding knowledge ascription. Though, as I have been arguing, the presumption that ‘S knows that P’ is true only if S knows that P is unwarranted.

Other Contextualists and the Skeptical Argument

Cohen and DeRose each utilize a contextualist account of knowledge ascriptions for two related purposes: to preserve the truth of our ordinary knowledge ascriptions and to explain our reaction to skeptical arguments.²¹ Cohen and DeRose are also each clear that contextualism is a

²¹ Cohen 1999 p. 65; DeRose 1995 p. 4-5

theory about knowledge ascriptions, not knowledge. Cohen offers the following caveat for his treatment of skepticism:

According to contextualism, the truth-conditions of sentences of the form ‘S knows P’ are context sensitive. So strictly speaking, instead of saying that S knows P in one context but fails to know in another, one should really say that the sentence ‘S knows P’ is true in one context and false in the other.²²

DeRose also attempts to clear up potential misunderstandings regarding the relationship between contextualism and epistemology. He states that the debate between contextualism and invariantism is an issue in the philosophy of language, though “it’s a piece of the philosophy of language that certainly has profound importance to epistemology.”²³ If we respect Cohen and DeRose’s caveats then contextualism alone, as a semantic thesis, can only treat the skeptical argument as a series of knowledge ascriptions:

- 1: ‘I don’t know that I am not a BIV’ is true
- 2: If ‘I don’t know that I am not a BIV’ is true, then ‘I don’t know that I have hands’ is true
- C: ‘I don’t know that I have hands’ is true

The contextualist treatment of this argument would be to point out that the conditional that comprises the second premise is always true. In certain extraordinary contexts the antecedent-ascription is true and therefore the consequent-ascription is true. However, in ordinary contexts the antecedent-ascription is false and so the consequent-ascription can be false (and probably is false). Since the first premise and the conclusion represent the components of the second premise conditional, the same point is made regarding the success of the skeptical argument as a whole: if the first premise is true within its context then, by virtue of the second premise, the conclusion is also true; however, if the first premise is false then the conclusion does not follow. Hence, the

²² Cohen 1999 p. 65

²³ DeRose 2009 p. 19

skeptical argument succeeds only insofar as the skeptic is successful in convincing us that the ascription ‘I don’t know that I am not a BIV’ is true, and the skeptic will be successful in convincing us that *that* ascription is true only if we accept what the skeptic means by ‘knows.’

However, construing the skeptical argument as a set of knowledge ascriptions misses the mark. The skeptic qua skeptic is not concerned with the truth or falsity of our knowledge ascriptions; rather, the skeptic is concerned to make the point that our beliefs do not satisfy the criteria of knowledge. The skeptic *may* be interested in claiming that the ascription ‘I know that I have hands’ is false as an extension of the point that my belief that I have hands does not constitute knowledge. If so, the skeptic would be appealing to the idea that the ascription ‘S knows that P’ is true only if S knows that P, and reasoning that since S does not know that P, the ascription ‘S knows that P’ is false. However, that argument is not a necessary extension of the skeptic’s essential claim. A skeptic could embrace a contextualist account of knowledge ascriptions while insisting that our beliefs do not constitute knowledge, i.e., a version of the position I suggest. Hence, contextualism alone, as a semantic thesis about knowledge ascriptions, cannot undermine the skeptic’s claim that our beliefs do not satisfy the definition of knowledge.

Although contextualism alone cannot undermine the skeptic’s point about knowledge, contextualism can contribute to an argument against skepticism in one of two ways. The first option is to combine or merge the semantic thesis of contextualism with a contextual theory of knowledge such that the contexts in which the ascription ‘S knows that P’ is true are the contexts in which S’s belief that P does constitute knowledge. A version of this general strategy is what DeRose has in mind. DeRose intends to expand on ideas already present in “Solving the Skeptical Problem” and propose that knowing P requires that (1) P is true and that (2) S’s belief that P matches the fact of the matter in the ‘sphere of epistemically relevant worlds’ where *the*

*size of the sphere varies with context.*²⁴ These two criteria for knowledge fit nicely with DeRose's version of semantic contextualism and allow DeRose to maintain that in most contexts we do know various Ps, even if we would not know those Ps in other contexts. I am dissatisfied with allowing the criteria for knowledge to be context sensitive and so I recommend the other way that contextualism can assist an argument against skepticism.

The second option is to recognize that the skeptic's conclusion is correct as a point about knowledge, but to shift the focus from the issue of whether a belief constitutes knowledge to whether a knowledge ascription is true, given its context. The skeptic's conclusion is correct because the skeptic is insisting that knowledge is sensitive belief. However, we should recognize that our interest in knowledge is an interest in contextual knowledge ascriptions, not whether a belief is strictly sensitive, independent of any context of ascription. If we recognize this point, then we can focus on the question of whether S knows that P *according to a set of intentions, purposes, expectations and presuppositions*, and we can ignore the question of whether a belief is knowledge independent of such considerations. In doing so, we have a basis for contently considering the skeptic's claims about knowledge to be uninteresting. This response is in line with a tradition that holds that skepticism is true but irrelevant to our daily lives; what has been offered is an explanation for why skepticism seems irrelevant.

One of the hurdles to executing this shift in emphasis is the concern that a knowledge ascription is not valuable unless it is a report that S's belief *really is* knowledge. I believe that concern is misguided. It is a mistake to suppose that the knowledge ascriptions that have been, and will continue to be, valuable to us have been valuable because they report that S's belief satisfies the definition of knowledge. Knowledge ascriptions are valuable as contextual reports

²⁴ DeRose 1995, 2009 p.14

that (for all intents and purposes, etc.) S counts as knowing. Contextual knowledge ascriptions are valuable because our interest in whether S knows that P is an interest in which alternatives can be ruled out and which cannot. The issue of whether a belief constitutes knowledge, independent of a context of ascription, may be interesting but it is not of direct concern. Nonetheless, the definition of knowledge is not expendable because it remains at the core of our knowledge ascriptions. Knowledge is a sensitive belief, and our knowledge ascriptions are contextual reports that S's belief *is* a sensitive belief, though, the factors comprising the context of ascription adulterate the notion of sensitivity. If we abstract from the context of ascription, i.e., we do not excuse the factors that limit the alternatives that must be ruled out, our knowledge ascriptions are revealed to be a claim that S's belief is a sensitive belief.

IV. References

- Cohen, Stewart. 1999. "Contextualism, Skepticism and the Structure of Reasons." *Philosophical Perspectives* Volume 13: 57-89.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *The Philosophical Review* Volume 104, Issue 1: 1-52.
- 2009. *The Case for Contextualism*. New York: Oxford University Press.
- Dretske, Fred. 1971. "Conclusive Reasons." *Australasian Journal of Philosophy* Volume 49, Issue 1: 1-22.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* Volume 23, Issue 6: 121-123.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* Volume 74, Issue 4: 549-567.
- Moore, GE. 1962. *Philosophical Papers*. "Proof of an External World." New York: Collier Books. 144-148.
- Nozick, Robert. 1981. *Philosophical Investigations*. Cambridge: The Belknap Press of Harvard University Press.
- Unger, Peter. 1984. *Philosophical Relativity*. Oxford: Oxford University Press.
- Williams, Michael. 1991. *Unnatural Doubts*. Cambridge: Blackwell.