

University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

Philosophy Dissertations, Theses, & Student
Research

Philosophy, Department of

5-2011

Epistemic Contextualism: A Defense and Analysis

Sruthi R. Rothenfluch

University of Nebraska-Lincoln, sruthirachel@gmail.com

Follow this and additional works at: <http://digitalcommons.unl.edu/philosophydiss>



Part of the [Epistemology Commons](#)

Rothenfluch, Sruthi R., "Epistemic Contextualism: A Defense and Analysis" (2011). *Philosophy Dissertations, Theses, & Student Research*. 4.

<http://digitalcommons.unl.edu/philosophydiss/4>

This Article is brought to you for free and open access by the Philosophy, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Philosophy Dissertations, Theses, & Student Research by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

EPISTEMIC CONTEXTUALISM:
A DEFENSE AND ANALYSIS

by

Sruthi Rothenfluch

A DISSERTATION

Presented to the Faculty of
The Graduate College at the University of Nebraska
In Partial Fulfillment of Requirements
For the Degree of Doctor of Philosophy

Major: Philosophy

Under the Supervision of Professor Albert Casullo

Lincoln, Nebraska

May, 2011

EPISTEMIC CONTEXTUALISM: A DEFENSE AND ANALYSIS

Sruthi Rothenfluch, Ph.D

University of Nebraska, 2011

Advisor: Albert Casullo

Epistemic contextualists maintain that the truth conditions of knowledge ascriptions and denials change according to the context of utterance. In this dissertation, I defend this view against one of its main rivals, classic invariantism, which holds that the contents of such statements remain fixed across contexts. While epistemic contextualists provide a straightforward semantic account of the variability in our knowledge-ascribing behavior, classic invariantists cannot, and therefore must offer some explanation as to why it seems as though the standards for ‘knowing that p ’ shift from one context to the next. To this end, classic invariantists draw a distinction between what speakers convey (the pragmatic content of their assertions) and the proposition literally expressed by their utterance. They maintain then that the standards of ‘knowing that p ’ seem to vary, not because the meaning of ‘knows’ shifts, but rather because what is conveyed or grasped by these utterances vary from one context to the next.

The main thrust of classic invariantism comes from the claim that invariantists can account for fluctuations in uses of ‘knows’ without adopting a context-sensitive analysis of the term. My aim in the dissertation, then, is two-fold: First, I chip away at the general strategy of invoking pragmatic implications to explain away unfavorable linguistic intuitions, thus undermining the classic invariantist’s capacity to explain the phenomenon in question: the variability of our knowledge ascriptions and denials.

Without a proper explanation of our use of 'knows', one that clarifies why we *believe* that 'knows' applies in a given context when it strictly does not, the entire invariantist enterprise fails. After a brief introductory chapter, I discuss, and argue against, a version of moderate invariantism in the second chapter, skeptical invariantism in the third and low-standard invariantism in the fourth. In the fifth chapter, I show the context-sensitivity of 'knows' is linguistically supported by identifying 'knows' within the class of gradable context-sensitive expressions such as 'tall' and 'flat'.

For Michael and Maya

Table of Contents

Abstract	ii
Table of Contents	v
1. Introduction	1
I. Three Theories	2
II. The Role of Contextualism in Epistemology	20
III. A word about Method: Ordinary Language	25
IV. Summary of Chapters	39
2. Contextualist and Invariantist Error Theories	48
I. Contextualism	49
II. Rysiew’s Picture	52
III. What is Said	57
IV. Alternative Error Theory	67
V. Conclusion	72
3. Skeptical Invariantism and Epistemic Distinctions	74
I. Pragmatics and WAMs	76
II. Absolute Terms	79
III. Recalcitrant Intuitions and Emphasis	82
IV. Problems with Emphasis	85
V. Different Forms of Emphasis, Unreliability	90
VI. Conclusion	93
4. Cancelability and Semantic Entailment	95
I. Arguments for and against the Infelicity of K-conjunctions	96
II. Loose Talk and Conversational Implicature	100
III. Loose Talk in Black and Murphy’s Examples	106
IV. Residual Concerns	109
5. Knowing Better	111
I. Gradability and Contextualism	113
II. Knowing that and Knowing why: Cases	119
III. Knowing that and Knowing why: Analysis	122
IV. Linguistic Basis for the Context-Sensitivity of ‘Knows’	137
V. Conclusion	143
6. Conclusion	144
I. Arguments and Implications	144
II. Further Worries	154
Works Cited	165

Chapter One

Introduction

Inclinations to ascribe knowledge to ourselves and others vary from one context to the next. Whether we are willing to attribute propositional knowledge appears to depend not only on the subject's evidence or justification for the relevant proposition, but other non-epistemic factors, such as the importance of the truth of the proposition in question. In addition, our ascribing behavior seems to shift with the variety of counter-possibilities that have been mentioned or seriously considered. Reflecting on radical skeptical scenarios, for example, will often cause speakers to sincerely doubt their knowledge of the most ordinary propositions. To give the reader an idea of the type of linguistic phenomenon I'm describing, let us compare a set of cases: In the first case, I seem to remember getting my car serviced very recently. When my parents ask whether I am taking good care of the car, I confidently tell them that I *know* that I had my Toyota serviced three weeks ago. In the second scenario, everything is as it was before, but this time our family is planning a cross-country trip, and car problems may prove disastrous. In these circumstances, I might not find my memory sufficient. That is, before claiming to *know* that my car was serviced, I might check the records. Note that what has changed in the two cases is the importance of having a safe and reliable car. When the car's reliability becomes critical, I hold myself to more stringent conditions by not attributing knowledge to myself solely on the basis of memorial evidence.

What should we make of such fluctuations in our knowledge ascribing behavior? Are we speaking inaccurately in some cases? That is, in the first scenario described

above, though I ascribe knowledge to myself that my car was serviced only weeks ago, am I using ‘knows’ loosely, perhaps to convey that my epistemic position with respect to this proposition is close enough to knowing for my purposes? Or, does ‘knows’ and the statements that contain it have certain semantic properties that account for such variation in our speech? In the following chapters I argue in favor of the latter, contending that ‘knows’ is a context-sensitive expression, whose meaning shifts according to the context of utterance, a view called *epistemic contextualism*.¹ I will expound and examine alternative explanations of this linguistic phenomenon, ultimately rejecting them on the grounds that they cannot completely and satisfactorily account for the vacillation we observe in judgments about ours’ and others’ knowledge. I will then offer a theory that characterizes the manner in which ‘knows’ is context-sensitive. The details of my argumentative strategy and the sub-conclusions I draw will be delineated in a later section of this introductory chapter. Before we get to this, however, I will set the stage by presenting views of the primary contenders in the debate and identify precisely where my arguments fit in the dialectic.

I. Three Theories

¹ There are, in fact, two ways of construing the view in terms of formal semantics. First, we could think of knowledge ascriptions as involving an indexical reference to standards, in which case the knowledge predicate will actually express different relations across relevantly different contexts. Second, ‘knows’ might correspond to a fixed relation, that is, where ‘knows that *p*’ expresses the same relation in every context. But, the context would determine the standard at which the proposition involving the knowledge relation gets evaluated. Either formulation permits the central contextualist claim which is that the standards of ‘knowledge’ are dependent on the context. This distinction can be found in Stewart Cohen, "Contextualism, Skepticism and The Structure of Reasons," *Philosophical Perspectives 13: Epistemology*, ed. James E. Tomberlin (Atascadero: Ridgeway): 61. Kent Bach also briefly touches on the issue in "The Emperor's New 'Knows'," *Contextualism in Philosophy: Knowledge, Meaning and Truth*, eds Gerhard Preyer and Georg Peter (Oxford: Clarendon Press, 2005): 51-89. He maintains that "The effect is the same either way.... It expresses either a contextually variable two-term relation or a fixed three-term relation, whose third term, the operative standard varies with context" (60).

Broadly speaking, we can divide the relevant semantic theories into three categories: classic invariantism, subject sensitive invariantism and contextualism. Each type offers distinctive explanations of our knowledge ascribing behavior, and are differentiated according to their treatment of the context-sensitivity of our use of ‘knows’. Classic invariantists maintain that knowledge ascriptions and denials correspond to fixed truth conditions across contexts, but that other conditions, such as the warrantability of the assertions (i.e., whether or not they are conversationally appropriate) will depend on the interests and concerns of speakers. Whether or not an attribution or denial of knowledge is warranted is determined by what is communicated or intended by the speaker’s use of ‘knows’, which is taken to depart from its semantic content at least on some occasions. Classic invariantists, then, restrict the variability of ‘knows’ to the non-semantic realm. Subject-sensitive invariantists, on the other hand, maintain that whether or not a subject knows that p at a given time depends on practical factors, so that while S might know that p at time t_1 , she might not know at t_2 , due to shifts in her own situation. Subject-sensitive invariantists, then, allow that the truth conditions of sentences containing ‘knows’ are sensitive to certain non-epistemic contextual features in the subject’s environment. However, sentences of the form ‘ S knows that p at t ’ will have a fixed truth value, regardless of who is making the claim, or the context of the ascriber. Contextualists maintain that the truth conditions of knowledge ascriptions and denials will depend on contextual features in the speaker’s environment. In this section, I will discuss each theory in more detail, focusing on specific accounts and highlighting the differences among them.

Classic Invariantism

The core claim of classic invariantism—that ‘knows’ is insensitive to non-epistemic contextual features, and therefore, that the truth conditions of sentences containing ‘knows’ do not change (given a fixed subject, proposition p and epistemic status e)—allows significant variability in the theories it generates, particularly with respect to the standards of knowledge they posit. Skeptical invariantists set the bar at a very high and even unattainable level, such that ‘S knows that p ’ is true only if the subject has eliminated even the most unlikely counter-possibilities to p , making it the case that most (if not all) of our knowledge ascriptions turn out to be false. A thorough and systematic defense of this position can be found in Peter Unger’s *Ignorance*.² Here, he categorizes ‘knows’, along with ‘flat’, ‘empty’ and ‘vacuum’ as *absolute* terms, which signifies that their correct application requires the complete absence of an opposite corresponding feature. ‘Flat’, for example, applies to an object that is not bumpy or curved to any degree (in other words, absolutely flat). Unlike ‘flat’, ‘knows’ is not a *basic* absolute term. That is, its meaning involves an intermediate absolute term: In order for a subject S to know that p , she must be certain. One is ‘certain’ insofar as no degree of doubt remains in her mind, thereby rendering ‘knows’ a non-basic, absolute expression. One consequence of this analysis, as stated above, is that it turns out that we know hardly anything, perhaps nothing. While elegant, the account cannot succeed without some explanation as to why it *seems* as though we know a great many propositions. Unger explains away these intuitions by invoking a common classic invariantist strategy, which consists of appealing to the pragmatic content of our assertions to account for our mistaken judgments. Unger maintains that while the

²Peter Unger, *Ignorance* (Oxford: Oxford University Press, 2002).

meaning of 'knows' remains fixed, speakers *use* the term to indicate different degrees of epistemic status as they relate to, or rather, suffice for, the practical purposes at hand. He further notes that this is not at all uncommon in the language, but is found to be true of absolute terms in general:

True enough people *call* all sorts of things 'flat'. They use the word to pick out surfaces which compare poorly with those needed for certain precision machinery. They use it even to refer to certain prairies where, even discounting the curvature of the earth, which one really should not discount, there are dips and rises of inches and feet....The normal use is, in each case, to indicate that the object in question is close enough to being flat, for the kind of thing it is, for the purposes which may be presumed in the case, and relative to other factors the context may provide.³

In the same way, though the content of 'knows' does not shift from one conversation to the next, we may nevertheless use it to convey that the subject's epistemic position is close enough to actually knowing for the current situation. Unger, then, identifies 'knows' as a semantically stable expression, while explaining away our variable judgments about whether or not one 'knows' by appealing to non-semantic content.

Moderate and low-standard invariantists posit lower standards of knowledge, making it the case most of our ascriptions turn out to be true on their accounts. Patrick Rysiew, Tim Black and Peter Murphy have independently developed and defended different versions of these views.⁴ We shall briefly look at Rysiew's account to gain an understanding of what such a position offers. Rysiew frames his brand of moderate invariantism in terms of a relevant alternatives theory, subscribing to the view that knowledge involves the elimination of all relevant alternatives, which (in Rysiew's model) are limited to what *normal* humans take to be likely counter-possibilities. If our

³ Unger, 68-69

⁴ Patrick Rysiew, "The Context-Sensitivity of Knowledge Attributions," *Nous* 35, no. 4 (2001): 477-514; Tim Black and Peter Murphy, "Avoiding the Dogmatic Commitments of Contextualism," *Grazer Philosophische Studien* 69 (2005), "Defending a Sensitive Neo-Moorean Invariantism," in *New Waves in Epistemology*, ed. Vincent Hendricks and Duncan Pritchard (Ashgate: Palgrave Macmillan, 2005), 8-27.

subject S has eliminated $\sim p$ possibilities that are likely from a normal human perspective, then it will be true, on Rysiew's view, that S knows that p , regardless of the circumstances of S or the speaker, or the counter-possibilities that might have come up in conversation.

But like their skeptical counterparts, moderate invariantists must also provide an error theory. That is to say, Rysiew must explain why it is that under certain circumstances, we hesitate to ascribe knowledge or assert that the subject does not know that p , despite meeting his proposed standard. Rysiew's approach resembles Unger's in that it involves an appeal to the pragmatic content of the utterance, or what is communicated rather than literally expressed. How so? Rysiew concedes that when remote counter-possibilities are mentioned, we tend to deny that the subject knows the relevant proposition. When we do so, however, we are not using this denial literally, according to Rysiew. That is, we are not conveying that S does not in fact know that p , but rather that S cannot eliminate the just-mentioned or salient counter possibilities, which include alternatives that are *not* likely from a normal human perspective. This is why it often seems correct to say 'S does not know that p ' even though, strictly speaking, S knows that p in virtue of eliminating the relevant alternatives. Hence, Rysiew also relies on what he believes to be the pragmatic content of our assertions to explain away unfavorable intuitions. While at the outset, such a strategy appears adequate to handle fluctuations in our knowledge ascribing behavior, it is in fact ineffective in accounting for such data, or so it will be argued in the three chapters that follow.

Subject Sensitive Invariantism

Subject sensitive invariantism (SSI) incorporates certain practical features of the subject's circumstances in its semantic theory. Factors such as how important the issue is for a subject *S* at a particular time, *S*'s interests and attention affect whether or not *S* knows the relevant proposition at that time. Subject sensitive invariantists, like classical invariantists, can marry their position with various theories of knowledge. In combining SSI with a relevant alternatives approach, for example, subject sensitive invariantists would hold that practical facts in one's environment affect which set of alternatives the subject must rule out in order to know that *p* at any particular time. John Hawthorne (2004) and Jason Stanley (2005) have offered two distinct versions of such a view.⁵

The account nicely accommodates the linguistic behavior cited at the beginning of this discussion, specifically, that our tendency to ascribe knowledge varies with certain non-epistemic factors. When the date of my car's last service was fairly inconsequential, my somewhat vague memory that I had taken it to the dealership three weeks ago was sufficient for me to know that I had my car serviced recently. However, when the car's reliability and performance became critical, due to our impending journey, my memory was not sufficient. Rather, I found it necessary to consult my records before claiming to know that the car had been serviced recently. SSI allows for such variation by incorporating my interests and concerns into the semantic content of 'knows'. But the worry for subject-sensitive invariantists arises in cases where the ascriber is not the putative knower, and more specifically, when the ascriber is in a high-stakes situation, but the subject is not. Stanley describes one such case: Hannah and Sarah are worried

⁵ Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005). John Hawthorne, *Knowledge and Lotteries* (Oxford: Clarendon Press, 2004).

about a future bill, so they want to know whether the bank will be open on Saturday to deposit their checks. To resolve this question, they phone Bill, who claims to have been at the bank two weeks ago on a Saturday when it was open, but has no interest in whether or not the bank will be open. Sarah and Hannah agree that Bill does not know whether the bank will be open this Saturday, which appears to be a reasonable assessment, given their predicament. However, according to SSI, since the bank's hours on Saturday matter little to Bill, his evidence suffices for knowledge. Subject-sensitive invariantists, then, must appeal to an error theory to explain away our mistaken intuitions in such cases.

Such error theories, however, have been effectively refuted by invariantists themselves and proponents of contextualism such as Jonathan Schaffer. As such, I will not be addressing subject-sensitive invariantism in my discussion, except to briefly rehearse some of what is already in the literature in order to set aside this alternative explanation of our knowledge ascribing and denying behavior. Hawthorne has suggested that ascribers in high-stakes situations (like Hannah's) project features of her own context to the subject's, thereby incorrectly holding the subject at high epistemic standards. He invokes psychological data that appear to show that we tend to overestimate the likelihood of a risk by the ease with which we can recall or imagine it. In the case of knowledge-attributions, we tend to overestimate the real danger of counter-possibilities and then deny "knowledge to others where there is in fact no real danger of error".⁶ In the scenario above, Hannah overestimates the likelihood of the bank changing its hours due to her awareness of the costs of being wrong. This error theory has been rejected on the grounds that it overgenerates to cases where the subject is in a high-stakes situation,

⁶ Hawthorne, 164.

but the ascriber is not. Suppose that *we* are considering Hannah's knowledge (keeping fixed Hannah's urgency to have her check deposited). When we intuit that Hannah does not know that the bank is open the next day, we mistakenly over-estimate the likelihood of the possibility that the bank is closed due to the risk conditions Hannah faces, and Hannah really *does know* that the bank will be open (Stanley 2005, Williamson, 2005, Schaffer, 2006),⁷ which, of course, is not what the subject sensitive invariantist will want. This suggests that Hawthorne's strategy of incorporating pragmatic features in the ascriber's physical and psychological environment does not adequately explain our language-ascribing behavior.

Stanley proposes a different account of such cases. He writes that when we discuss whether or not *S* (whose interests and concern differ from our own) knows that *p*, what concerns us is whether *S*'s epistemic state with respect to that proposition would be sufficient for *us* to know in our situation.⁸ So, when Sarah and Hannah phone Bill, what they want to know is if his evidence would suffice for knowledge, given their circumstances. This explains why we think Sarah and Hannah are right. That is, we intuitively recognize that what Hannah and Sarah care about is whether Bill would know, were he facing Hannah and Sarah's concerns. Stanley emphasizes that "the *purpose* [someone in] High Stakes has in asking someone else whether or not *p* is true lies in finding out whether, *if that person had the interests and concerns High Stakes does*, that person would know that *p*. Since *p* is a serious practical question for High Stakes, she is

⁷ Stanley, 101; Jonathan Schaffer, "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism," *Philosophical Studies* 127 (2006): 87-107; Timothy Williamson, "Contextualism, Subject-Sensitive Invariantism and the Knowledge of Knowledge," *The Philosophical Quarterly* 55, no. 219 (April 2005): 213-235.

⁸ Stanley, *Knowledge and Practical Interests* 101-102.

not really worried about that person's own interests and concerns".⁹ Jonathan Schaffer, however, points out that Stanley's explanation would predict that individuals in low-stakes scenarios should then make the same mistake when assessing those in high-stakes situations. That is, a speaker in a low-stakes situation would ascribe knowledge to a subject in a high-stakes situation since, if the high-stakes subject were in her own position she would know. This prediction is unfavorable to SSI on two counts: first, as Stanley himself acknowledges, this is not our intuition in 'low attributor-high subject stakes' situations.¹⁰ Second, this undermines evidence that was originally marshaled as supporting subject-sensitive invariantism. Once again, appealing to the speaker's semantically incorrect but pragmatically warranted understanding of the knowledge claim does not help the subject-sensitive invariantist explain our intuitions in cases where the subject is in a high-stakes scenario, but the putative knower is not.

I will not revisit subject-sensitive invariantism for the remainder of the discussion, in light of two related considerations. First, the main issue in the upcoming chapters concern error or pragmatic theories that are intended to explain our (allegedly mistaken) linguistic intuitions. Second, SSI error theories have already been successfully rejected in the literature in much the same way that I intend to refute classic invariantist theories in the discussion that follows, arguing specifically that proponents of the latter cannot provide an adequate pragmatic explanation of our knowledge-ascribing behavior. (I will further elucidate my project in the last section of this chapter.) Hence, SSI can be set aside for the purposes of this discussion.

Contextualism

⁹ Stanley, 102 (original emphasis).

¹⁰ Stanley, 5.

Contextualists maintain that our use of ‘knows’ varies from one context to the next because the semantic content of ‘knows’ itself is sensitive to the interests of conversational participants and the general direction of the conversation. A sentence of the form ‘S knows (does not know) that *p*’, may be false in one context and true in another despite a fixed subject, proposition, and evidentiary status. While some situations demand extremely high epistemic standards (say, in an operating room), in other situations, we might truly ascribe knowledge to *S* despite *S*’s relatively weak epistemic status with respect to the relevant proposition (perhaps in a casual conversation at a bar). Shifts in meaning may be triggered by various factors: the mentioning of counter-possibilities and the interests and goals of the speakers and the stakes involved.¹¹ Contemporary proponents of the view include Stewart Cohen, Keith DeRose and David Lewis. However, the idea that ‘knows’ may admit of multiple meanings can be found in Norman Malcolm’s 1952 discussion of the distinction between the *strong* and *weak* senses of knowledge as well as Gail Stine’s (1976) “Skepticism, Relevant Alternatives and Deductive Closure,” in which she maintains that ‘it is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts’.¹²

While contextualism provides a viable explanation of our ordinary knowledge ascriptions, it also provides a potential solution to the skeptical puzzle. We will briefly depart from our discussion of ordinary uses of ‘knows’ to focus on the skeptical paradox for two reasons. First, the theory’s ability to resolve the paradox is one its main attractions and I would be remiss if I did not explicate the contextualist’s treatment of the

¹¹ Contextualists discuss different mechanisms to explain shifts in epistemic standards. These differences will be discussed shortly.

¹² Norman Malcolm, “Knowledge and Belief,” *Mind* 61, no. 242 (April 1952): 178-189. Gail Stine, “Skepticism, Relevant Alternatives and Deductive Closure,” *Philosophical Studies* 29, no. 4 (April 1976): 254.

puzzle. Second, while my objective is to offer an account of our ordinary knowledge ascriptions, my discussion in the following chapters makes some reference to the skeptical puzzle. Here, then, is one version of the skeptical paradox:

1. I don't know that I am not a brain in a vat programmed to experience ordinary sensations (BIV).
2. If I don't know that I'm not a BIV, then I don't know that I have hands.
3. I know that I have hands.

Each statement seems plausible, yet mutually inconsistent. At least one of these must be false. Contextualists maintain that skeptics raise the standard of knowledge in the first premise by introducing a highly remote counter-possibility. It turns out, then, that (3) is false relative to these skeptical standards.¹³ However, in normal, everyday contexts, where we are not entertaining the possibility of being a BIV, we both know that we have hands and that we are not BIVs, making (1) false. The puzzle, then, is only apparent.

While most contextualists accept such a resolution of the skeptical argument, there are in fact important distinctions between different contextualist theories, which concern both the mechanisms that trigger shifts in epistemic standards, and what it is that shifts.

Below, I will survey some of the most prominent versions of contextualism.

Keith DeRose endorses an externalist form of contextualism, according to which whether or not S knows that *p* depends on the truth-tracking capacity of her belief. The context of utterance demarcates a sphere of possible worlds surrounding the actual world, within which a subject's belief that *p* must match the fact of the matter about *p* in each of these worlds in order to know that *p* in a given context. As epistemic standards rise, the truth-tracking of one's belief must extend further from actuality for one to count as

¹³ One might also address the skeptical argument by denying closure. See, for example, Fred Dretske, "Is Knowledge Closed under Known Entailment? The Case Against Closure," in *Contemporary Debates in Epistemology* (Malden: Blackwell Publishing, 2005): 13-25.

knowing. In ordinary, non-philosophical contexts, various factors—such as “the importance of being right and what has been said in the conversation”—affect the degree of epistemic strength that the subject must attain in order to know that p , by expanding or contracting the boundary of possible worlds.¹⁴ What actually shifts, then, is the epistemic strength required to know that p , or, in other words, the sphere of possible worlds in which the subject’s belief must track the truth of p . What prompts such shifts (that is, what heightens the epistemic strength required for knowledge) are both practical concerns and what has been uttered in conversation.

Stewart Cohen’s contextualism differs in certain respects from DeRose’s theory and undergoes certain modifications throughout its development. The original formulation was founded on a relevant alternatives framework, which tied the truth conditions of ‘ S knows/does not know that p ’ to salient alternatives to p . However, he diverges from the relevant alternatives structure in subsequent discussions, while still incorporating the idea of the salience of error possibilities. In the first version then, Cohen classifies h as an alternative to a proposition q just in case h is incompatible with q . An alternative is relevant insofar as it satisfies one of two criteria: The external criterion focuses on features in the subject’s environment and states that “an alternative h (to q) is relevant if the probability of h conditional on reason r and certain features of the circumstances is sufficiently high”, where the degree of sufficient probability is

¹⁴ DeRose, “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52, no.4 (December 1992), 916. DeRose offers a slightly different treatment of the skeptical paradox. Here, while the skeptic also ultimately manipulates the conversational context to raise the standards of knowledge, he does so by invoking the rule of sensitivity. That is to say, S knows that she is not a BIV if and only if were this belief false, she would not believe it. A thorough discussion can be found in DeRose, “Solving the Skeptical Problem,” in *Skepticism*, eds. Keith DeRose and Ted A. Warfield (New York: Oxford University Press, 1999), 205.

determined by context.¹⁵ For example, if it is the case that zoos regularly deceive their visitors by placing false labels on their exhibits and disguising animals, then the possibility that we are observing cleverly painted mules is relevant to the belief that we are looking at zebras. The second criterion concerns the subject's evidence in determining the relevance of an alternative: "An alternative (to q) h is relevant if S lacks sufficient evidence (reason) to deny h, i.e., believe not-h", where the level of sufficient evidence is determined by the context in which the attribution of knowledge occurs.¹⁶ If S has absolutely no evidence about the honesty of zoo keepers, then that they might be deceiving visitors becomes a relevant alternative to the proposition that the animals in view are in fact zebras.

What shifts, then, from one context to the next is the standard of relevance which is determined by the external and internal criteria. But what causes these shifts? Cohen maintains that the standard of relevance shifts when the chance of error becomes salient. Let us return to the zoo example to see how this works. Suppose that a zoo visitor mentions the possibility that what we are looking at are not really zebras, but really cleverly painted mules. This utterance might point out that my evidence for believing that these are zebras is insufficient to rule out a situation in which the animals are cleverly painted mules, which expands the standard of relevance to include this possibility. In this context, I neither know that these are zebras nor that these are not cleverly painted mules. According to Cohen, when such doubts or counter-possibilities are not raised, we know both that these are zebras and that these are not cleverly painted mules. This is because in the absence of these challenges, our evidence is sufficient to

¹⁵ Stewart Cohen, "How to be a Fallibilist," *Philosophical Perspectives* 2 (1998): 102.

¹⁶ Cohen, "How to be a Fallibilist", 103.

deny the alternative that these are cleverly painted mules, that is, to prevent it from being relevant.

In “Contextualism, Skepticism and the Structure of Reasons”, Cohen takes a slightly different route, by detaching his account from the relevant alternatives framework (employing concepts such as justification and strength of reasons, instead) but still endorsing the second part of his analysis, which ties shifts in epistemic standards to salient error possibilities. Here, Cohen maintains that context determines how justified a belief must be in order to qualify as justified *simpliciter*, or justified sufficient for knowledge. So, what shifts from one context to the next is the degree of justification required to know that *p* or how strong a subject’s reasons must be to know that *p*, while what causes these shifts remains the salience of the chance of error.

David Lewis’ analysis also involves the elimination of alternatives. According to Lewis, S knows that *p* if and only if “S’s evidence eliminates every possibility in which not-*p*—Psst!—except those possibilities that we are properly ignoring”.¹⁷ Lewis maintains that ‘every’ is normally restricted to a limited domain and offers the following example to illustrate: “If I say that every glass is empty, so it’s time for another round, doubtless I and my audience are ignoring most of all the glasses there are in the whole wide world throughout all of time....They are irrelevant to the truth of what I have said”.¹⁸ The *sotto voce* proviso functions to reiterate the restricted application of ‘every’. Note further that the proviso requires speakers to ignore only those possibilities that it is *proper* to ignore. Which alternatives we may properly ignore are based on certain rules applied to the contexts of both the subject and the ascriber. For example, according to

¹⁷David Lewis, “Elusive Knowledge,” in *Skepticism: A Contemporary Reader*, eds. Keith DeRose and Ted Warfield (Oxford: Oxford University Press, 1999), 225.

¹⁸ Lewis, 225.

Lewis' Rule of Actuality, the possibility that actually obtains in the *subject's* context is never properly ignored, while the Rule of Attention dictates that possibilities that are not being ignored in the conversational (ascriber's) context are not properly ignored. These and other rules demarcate the set of possibilities that the subject's evidence must eliminate in order to know that *p* in a given context. What shifts for Lewis, then, is the set of possibilities or alternatives that may not be properly ignored. The mechanisms responsible for causing such shifts are Lewis' rules, which function either to preclude or include certain alternatives as relevant.

While contextualists take our knowledge attributions and denials at face value, they have to provide an explanation for certain other intuitions that conflict with their thesis. As stated above, contextualists maintain that the meaning of 'knows' changes according to particular factors in the context of ascription. One consequence of this view is that a speaker who utters 'S knows that *p*' in a low-standard context (say in casual conversation) will not contradict a speaker who states 'S does not know that *p*' in a high-standard context (perhaps where the stakes are very high). That is to say, according to contextualism, 'S knows that *p*' and 'S does not know that *p*' can both be true, when uttered in relevantly different contexts. While such compatibility makes sense within the contextualist picture, it certainly seems puzzling that the two statements present no conflict whatsoever. Contextualists, then, must provide some explanation for why some might intuit that 'S knows that *p*' contradicts 'S does not know that *p*' whenever these are uttered, so long as 'S' refers to a single subject and '*p*' to a single proposition.

One response available to the contextualist is to maintain that competent speakers are semantically blind to the context-sensitivity of 'knows'. English speakers are

inclined to think that the two statements above conflict with one another because they are not aware that the meaning of ‘knows’ shifts according to certain features in the conversational context. Stephen Schiffer has argued, however, that such an error theory is implausible, as it seems highly unlikely that competent speakers would be confused about the meanings of contextually variable expressions.¹⁹ Schiffer argues for this by examining three ways contextualists might characterize the context-sensitivity of ‘knows’, which include assigning a hidden indexical to knowledge sentences, identifying ‘knows’ as an indexical and identifying ‘knows’ as a vague term. He maintains that speakers are not typically confused about which proposition is expressed by a sentence that contains a hidden indexical, a vague term or explicit indexicals. Schiffer then concludes that regardless of the model contextualists use to characterize the context-sensitivity of ‘knows’, speakers would not be mistaken about the meaning of ‘knows’. Therefore, if ‘knows’ was context-dependent (and displayed one of these forms of context-sensitivity), then competent speakers would not systematically make mistakes about the content of knowledge ascriptions and denials. He maintains that “speakers would know what they were saying if knowledge sentences were indexical in the way the Contextualist requires”.²⁰

¹⁹Stephen Schiffer, “Contextualist Solutions to Skepticism,” *Proceedings of the Aristotelian Society* 96 (1996):317-333. Schiffer’s argument is in fact directed against the contextualist’s response to the skeptical argument. Contextualists maintain that the skeptical argument *appears* puzzling because speakers are unaware of the context-sensitivity of ‘knows’, which leads ordinary speakers to think that the argument is denying that they have knowledge of ordinary propositions even in non-skeptical contexts. The issue here concerns an apparent contradiction between knowledge ascriptions and denials that are uttered in significantly different (though perhaps not radically skeptical) contexts. However, the contextualist response is similar: speakers erroneously assume that there is a contradiction between knowledge ascriptions in low-standard contexts and knowledge denials in high-standard contexts because speakers are not aware of the context-sensitivity of ‘knows’. Since the underlying ignorance in both cases is the same, Schiffer’s objection applies here also.

²⁰ Schiffer, “Contextualist Solutions to Skepticism,” 328.

Cohen and DeRose have both offered different responses to Schiffer's objection.

Contra Schiffer, Cohen argues that speakers can be unaware of the context-sensitivity of an expression. He writes

You can lead competent speakers to question their everyday ascriptions of 'flatness' by making salient 'bumps' that ordinarily we do not pay attention to....If we implicitly raise the standards high enough then perhaps relative to that context, no physical surface is flat. But of course that does not impugn our ascriptions of flatness in everyday contexts where the standards are more lenient. But then why can we get competent speakers to question their everyday flatness ascriptions by implicitly raising the standards? It must be that although ascriptions of flatness are context-sensitive, competent speakers can fail to realize this. And because they can fail to realize this, they can mistakenly think that their reluctance to ascribe flatness, in a context where the standards are at the extreme, conflicts with their ascriptions of flatness in everyday contexts.²¹

Cohen essentially points out that semantic blindness to context-sensitivity is not as uncommon or atypical as Schiffer seems to suggest. It is then not unthinkable that ordinary, competent speakers might be mistaken about the contextual variability of 'knows', which then explains why they might falsely believe that knowledge ascriptions uttered in high-standard contexts contradict knowledge denials in low-standard contexts. Cohen, however, allows that the context-sensitivity of terms such as 'flat' might be easier to accept than contextualist theories about knowledge and justification.²² His explanation for this disparity invokes the normative nature of knowledge and justification. That is, saying that a belief is justified or that it qualifies as knowledge is to say something good about that belief. However, contextualism offers a deflationary account of these terms in the sense that 'S knows that *p*' uttered in ordinary, low-standard contexts is true even though S has not met our highest epistemic standards. Even though we have a lot of knowledge, it's not exactly what we had envisioned. In the same way, many surfaces are correctly called flat, even though they may contain some bumps. Cohen suggests that

²¹ Cohen, "Contextualism, Skepticism and the Structure of Reasons," 79.

²² He discusses this point in Stewart Cohen, "Contextualism and Unhappy Face Solutions: Reply to Schiffer," *Philosophical Studies* 119 (2004): 185-197.

this result is much easier to accept in the case of flatness than it is of knowledge, since ascriptions of flatness do not carry the normative force of knowledge ascriptions. This is why speakers might be less resistant to the context-sensitivity of terms such as ‘flat’, rather than ‘knows’. Cohen then refutes Schiffer’s argument by showing first, that the context-sensitivity of expressions will not be transparent to competent speakers, but second, that the contextual variability of ‘knows’ may be especially difficult to accept, given its evaluative function. This explains why speakers might be inclined to think that ‘S knows that p ’ is incompatible with ‘S does not know that p ’ even when the two statements are uttered in very different contexts.

DeRose takes a slightly different approach in his response.²³ He begins first by denying the intuition. That is, he maintains that many people would *not* agree that knowledge ascriptions uttered in low-standard scenarios conflict with knowledge denials asserted in more demanding contexts. Rather, if the high and low standard cases are constructed and presented properly, then it is not clear that there is a single intuitive response. As evidence, DeRose cites responses he has received from past students. Secondly, like Cohen, he accepts that the context-sensitivity of ‘knows’ might not be *as* obvious as certain other context-dependent expressions like ‘here’. However, he argues that such semantic blindness does not favor invariantism over contextualism. This is because speakers are equally unaware of the context *insensitivity* of the term, as shown by the fact that speakers will not consistently agree that knowledge ascriptions made in low-standard contexts contradict knowledge denials made in high-standard contexts. Thus, according to DeRose, the relevant intuition is not as widely shared as opponents

²³Keith DeRose, “‘Bamboozled by Our Own Words’: Semantic Blindness and Some Objections Against Contextualism,” *Philosophy and Phenomenological Research* 73, no. 2 (September 2002): 316-338.

seem to believe, and further, does not function as an effective threat against contextualism.

The foregoing discussion presented some of the extant views regarding our knowledge ascribing behavior. Before delineating my own aims in the following chapters and how they will further this debate, I will spend some time addressing worries some may have at the outset, which concern the importance of this project to the study of epistemology and the validity of the method employed throughout the discussion.

II. The Role of Contextualism in Epistemology

We have briefly surveyed three types of analyses that explain our knowledge ascribing behavior—classic and subject sensitive invariantism and contextualism—each of which appear to be reasonable attempts to account for such conversational exchanges. Note, however, that the theories aim at providing an explanation of our *use* of ‘knows’. One might wonder why such theories are important to epistemology, a study that concerns not our communicative intentions, but rather the requirements for genuine knowledge and justification. DeRose accurately sums up this sentiment: “Your contextualism isn’t a theory about knowledge at all; it’s just a theory about knowledge attributions. As such, it’s not a piece of epistemology at all, but of the philosophy of language”.²⁴ While it is true that contextualism is itself neither a theory about the structure of knowledge or justification nor a list of necessary and sufficient conditions for knowledge, it does prescribe a certain manner in which such theories must be interpreted. Mark Heller writes, “Contextualism provides a theory of the word ‘knowledge’ as an overlay for any theory of knowledge, but these are distinct components of a complete

²⁴ Keith DeRose, “Contextualism: An Explanation and Defense,” in *The Blackwell Guide to Epistemology*, eds. John Greco and Ernest Sosa (Malden, Mass.: Blackwell Publishers, 1999), 189.

epistemology".²⁵ One illustration of this overlay comes from DeRose, who explains how a contextualist semantics might frame more traditional positions within epistemology:

If you're a foundationalist, then if you're also a contextualist, you may well come to think of the issue of which beliefs are properly basic (i.e., the issue of which beliefs are justified to a degree sufficient for knowledge independent of any support they receive from other beliefs) and/or the issue of how strongly supported a belief in the superstructure must be in order to count as knowledge or as a justified belief to be matters that vary according to features of conversational context. And if you're a coherentist, then if you're also a contextualist, you'll probably want to hold that how strongly beliefs must cohere with one another in order to count as knowledge (if they're true), or to count as justified, is a contextually variable matter.²⁶

Indeed, the synthesis of contextualism and traditional epistemological theories offers potential resolutions for some devastating problems the latter face. Heller, for example, appeals to contextualism to address the generality problem against reliabilism.²⁷ While it is not clear that his strategy conclusively resolves the difficulty, it does seem to point in the general direction of a solution. Process reliabilists about knowledge maintain that *S* knows that *p* if and only if *p*, *S* believes that *p* and *S*'s belief that *p* is produced by a reliable method, that is, one that produces mostly true beliefs. Reliability, however, is defined as a propensity or tendency, which are features that apply not to a single instantiation of a method, but rather to *types* of processes. Whether or not a belief qualifies as knowledge, then, depends on whether the belief is produced by a process that belongs to a *type* that has the propensity to generate truth. The problem is that the process employed in a particular instance of belief formation can be a token of several different types. To illustrate, consider Conee and Feldman's example of what occurs during the formation of a commonplace perceptual belief: Smith, a tree expert with good

²⁵ Mark Heller, "The Proper Role for Contextualism in an Anti-Luck Epistemology," *Philosophical Perspectives* 13 (1999): 115.

²⁶ DeRose, "Contextualism: An Explanation and Defense," 192.

²⁷ Mark Heller, "The Simple Solution to the Problem of Generality," *Nous* 29, no. 4 (December 1995): 501-515.

vision, looks out of her window, where she sees a plainly visible maple tree nearby. She forms the belief that there is a maple tree near the house, and assuming that everything else in the example is normal, Smith knows that there is a maple tree. Process reliabilists can reach this verdict only if the process that Smith used is reliable. Conee and Feldman, however, point out that that Smith's process was an instance of several types: "visually initiated belief-forming process, process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby, process of relying on a leaf shape to form a tree-classifying judgment, etc".²⁸ Such processes feature different degrees of reliability. Which type must be sufficiently reliable? Reliabilists must provide a way to identify the type whose reliability determines whether a process yields knowledge. Conee and Feldman conclude, however, that there is no principled, epistemically defensible solution in the offing.²⁹

A potential response, however, may lie in a contextualist conception of process reliabilism—proposed by Mark Heller—according to which the extension of 'reliable' is determined by the context of ascription.³⁰ Heller argues that epistemic uses of 'reliable' do not differ from ordinary uses of the term, where its content varies from one context to the next. Despite its flexibility in everyday uses, the term is used without difficulty, given contextual clues as to what 'reliable' applies to. The context contributes in a similar way when the subject matter concerns belief formation. On some occasions, the

²⁸ Earl Conee and Richard Feldman, "The Generality Problem for Reliabilism," in *Epistemology: An Anthology*, eds. Ernest Sosa and Jaegwon Kim (Malden, MA: Blackwell Publishing), 372.

²⁹ Conee and Feldman focus on reliabilism about justification, rather than knowledge, but the authors clarify that nothing important turns on the differences between reliabilism about knowledge and reliabilism about justification.

³⁰ Another way in which contextualism can be incorporated into reliabilism is by allowing the degree of sufficient reliability to depend on the context of ascription.

process used will be described very generally and in other cases more precisely, which will have different implications:

If we demand more of a process in order to count it as reliable, we are demanding that it be reliable in more possible situations. One prerequisite of its being reliable in more situations is that it occur in more situations. Thus, as our standards rise, our description of the process type must become more general. On the other hand, we may sometimes be directly concerned with the question of when this process can occur again, the question of what is essential to this process, the question of which other process tokens are instances of the same process type. In these cases a raising of the standards is a demand that more be counted as essential to the process. The effect would be to require that the process be described more specifically, thereby making it easier for the process in question to count as reliable.³¹

If we accept Heller's claim that the relevant process of belief formation is determined by context, reliabilists are not saddled with the seemingly impossible task of supplying a fixed level of generality across all contexts.

One might worry, however, that speakers' interests may *not* identify a unique process in all cases of belief formation. Non-epistemic uses of 'reliable' tend to unequivocally refer to a particular object or capacity, perhaps because the purpose for which something is called 'reliable' is fairly straightforward. For example, suppose we are talking about icy road conditions and the likelihood of my car skidding and I say, "My car is very reliable" to dismiss worries. No one doubts that I am talking about my car's performance on snow-packed or icy roads. But, as Conee and Feldman correctly observe, there are other instances in which contextual factors do not seem to uniquely identify what 'reliable' applies to. Recall their original example of Smith who forms a belief about an adjacent maple tree. Suppose that Jones, who is sitting in the room with Smith utters, "Smith knows that there is a maple tree nearby". Conee and Feldman write, "nothing beyond the speaker's intentions seems to narrow the candidate pool [of belief

³¹ Heller, 509.

forming processes] in this sort of example”.³² The authors note that while it is possible that Jones had some processes in mind, it is possible or perhaps even more likely, that he was not envisioning any particular process of belief formation: “After all, [Jones] did not say anything about belief-forming processes and there is no reason to think that he was having any thoughts about them”.³³ Conee and Feldman then deliver the fatal blow that “there is no reason to think that in this sort of mundane example, there is such a thing as the contextually determined type”.³⁴

Perhaps the idea is not that speakers must have a particular process in mind, but at least that their interests and concerns help to determine a process. Such a formulation of Heller’s proposal would circumvent Conee and Feldman’s worry that the attributor did not have a particular process in mind. However, even this version would not yield a determinate method of belief formation in all cases. When Jones claims that Smith knows that there is a maple tree nearby, he might simply be telling someone about the types of vegetation that are found near Smith’s house, which would hardly select a single type of belief forming process. Although Heller’s strategy does not clearly eliminate the problem for reliabilists, it suggests one way in which proponents of the view might approach it. They might, for example, refine the position in a way that allows contextual features to determine a set of process-types, each of which feature relevant similarities and subsequently correspond to a similar degree of reliability which can then be assessed for knowledge. Or, reliabilists might be able to provide some rules *within* the contextually determined set that will determine the relevant type. While some additional work is necessary to solve the problem, the task will not be nearly as daunting as

³² Conee and Feldman, “The Generality Problem for Reliablism,” 383.

³³ Ibid.

³⁴ Ibid.

establishing a single level of generality to dictate across *all* interests, goals, etc.

Contextualism, then, contributes to the interpretation of epistemological theories and by so doing provides solutions for some problems commonly attributed to these views.

The debate, as I have presented it here, turns largely on the way in which ordinary, competent speakers use the language. Further, such examples are heavily employed in the course of the following chapters. This methodology is based on the presupposition that ordinary language is a valuable tool in philosophical analysis. But it is not at all obvious that our inclinations in common, everyday conversation should carry so much weight.

III. A word about Method: Ordinary Language

The method I have employed here relies to a great extent on ordinary linguistic behavior and intuitions. This appears to be a prudent decision insofar as theorists want to identify the concept that is of interest to us, that is, provide a semantic analysis of what we are doing when we distinguish some members of our conversational community as ‘knowers’. As Ludlow notes “any investigation into the lexical semantics of ‘knows’ will have to take seriously our intuitions about the proper analysis of knowledge”.³⁵ On the other hand, however, appealing to ordinary usage might not be conducive to a strict semantic analysis of the terms involved. Why not? One concern (alluded to earlier) is that competent speakers often use terms loosely or assert sentences which convey more than the literal content of the combination of the constituents of the sentence. Kent Bach writes,

³⁵ Peter Ludlow, “Contextualism and the New Linguistic Turn in Epistemology,” in *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, eds. Gerhard Preyer and Peter Georg (Oxford: Clarendon Press, 2005): 13.

We hardly ever mean exactly what we say. It is not that we generally speak figuratively or that we are generally insincere (these are different ways of not meaning what one says). Rather, we commonly speak loosely, by omitting words that could have made that we meant more explicit, and we let our audience fill in the gaps.³⁶

Bach goes on to write that spelling things out is usually unnecessary and simply functions to ‘slow things down’. Peter Unger expresses a similar sentiment, asserting that “in everyday affairs we often speak loosely, charitably, and casually; we tend to let what we say pass as being true”.³⁷ He further warns (seeming to speak directly to the method used here) “it is by being wrongly serious about this casual talk that philosophers have come to think it rather easy to know things to be so”. The idea, then, is that because speakers often speak in ways that are intended to communicate more than or diverge in other ways from the semantic content of their assertions, we ought not to rely on such utterances as evidence for our semantic theories.

A particularly threatening form of this argument comes from Patrick Rysiew’s discussion in “Speaking of Knowing,” in which he not only claims that speakers misuse expressions, but applies this general observation to the case of knowledge attributions and denials, and specifies the particular manner in which speakers stray from the literal content of ‘knows’.³⁸ Rysiew maintains that we have sufficient evidence that the relevant assertions are employed for their pragmatic implications (or what they communicate) rather than the propositions they literally express, and for this reason, we ought not to rely on such utterances as evidence for a semantic theory that purports to provide truth conditions for these statements. His argument is intended to raise serious doubts about the type of evidence used not only by contextualists, but also infallibilists and skeptics,

³⁶ Kent Bach, “Speaking Loosely: Sentence Nonliterality”, *Midwest Studies in Philosophy* 25, (2001): 250.

³⁷ Unger, 83.

³⁸ Patrick Rysiew, “Speaking of Knowing,” *Nous* 41, no. 4 (2007): 627-662.

all of whom he categorizes as proponents of *revisionary* theses. If he is correct, Rysiew's argument inflicts a much greater threat to contextualism than those who simply rely on the uncontroversial assumption that many of our conversational exchanges involve implicit content, which is not contained in the meanings of the expressions used, as Rysiew will have applied this unexceptional fact of conversation to the particular case of knowledge attributions and denials. In doing so, he would have provided strong evidence that ordinary language is not a reliable resource for semantic theory. Rysiew's argument, however, turns out to be unsuccessful as the conversational principles on which Rysiew founds his claim do not in fact extend to assertions about knowledge. This worry will be further elucidated after a more detailed presentation of Rysiew's view.

Rysiew subscribes to a common view of assertion, namely, that one represents oneself as knowing that *p* when one asserts '*p*'. Unlike Williamson, however, he does not endorse knowledge as a strict constitutive rule of assertion, according to which one must: assert *p*, only if one knows that *p*.³⁹ Rather, he simply claims that according to the maxims of Grice's Cooperative Principle—most importantly the maxim of Quantity and Quality—speakers convey that they know that *p* by making a genuine assertion. He further maintains that this conception of 'knowing' corresponds to a list of conditions he collectively calls the *ho hum* view. As the title suggests, the *ho hum* view is intended to be entirely uncontroversial and not itself a theory of knowledge, but rather a skeletal framework that may be fleshed out by more substantive theories. As such, these bare, and to some degree neutral, conditions are accepted by most epistemologists. The criteria for knowing that *p* in this view are limited to the subject's having an ungettiered justified

³⁹ Timothy Williamson, *Knowledge and Its Limits*, (Oxford: Oxford University Press, 2000): 238-266.

true belief that p , where the “reference to justified belief... is really just a marker for whatever further property must be added to convert an ungettiered true belief into knowledge”.⁴⁰ Since most epistemologists are also non-skeptics, the theory also attempts to preserve our anti-skeptical intuitions, by requiring relatively moderate conditions on belief and justification.

Recall, however, that the focus of our discussion at this point is not how traditional epistemologists conceive of knowledge, but rather what speakers communicate through assertion. Why think that the two coincide? Rysiew maintains that most competent speakers would come to recognize this as knowledge with some ‘*Theatetan*’ prompting. He further identifies certain parallels between the way in which the majority of epistemologists conceive of knowledge and the ordinary, lay-person’s conception. First, he writes that the ho hum view is essentially ungettiered true belief combined with being in a good epistemic position with respect to the proposition in question. He goes on to write that this is exactly what can be inferred from an assertion “given that the subject is presumed to be conforming to the second sub-maxim of Quality”.⁴¹ That is,

S: ‘ p ’ $\sim\>$ S is in a good epistemic position with respect to p .

where ($\sim\>$) symbolizes the relationship of being communicated or conveyed. Second, in asserting that p , while individuals represent themselves as being in a good position, they do *not* communicate that they are absolutely certain that p , have conclusive evidence, or a perfectly reliable basis for p , which parallels the ho hum view’s moderate conditions on belief and justification.

On the contrary, that people regularly assert in the absence of meeting such very strong requirements, and that they are not thought for that reason to be asserting improperly

⁴⁰ Rysiew, “Speaking of Knowing,” 632.

⁴¹ Ibid.

suggests that the belief and justification they represent themselves as having are not stronger than most theorists' moderate, non-sceptical readings of the belief and justification conditions on (ho hum) knowing.⁴²

The underlying idea, then, is that in asserting that p , speakers commit themselves to being in a strong (but not too strong) epistemic position with respect to p , which mirrors, more or less, the epistemologists' conception of knowledge.⁴³

Next, Rysiew writes that in any conversation there will be certain obvious and mutually available facts, those that are too obvious to actually mention. When one does explicitly point out such an evident fact, it is reasonable to infer that the speaker intends to express something different than the information literally expressed by her words. Bearing this in mind, we can now understand why knowledge attributions and denials (in general) are used to communicate something other than the subject's knowledge of p , at least according to Rysiew. As explained above, in asserting that p , the speaker represents

⁴² Ibid., 633.

⁴³ There is, in fact, a three-way identity being proposed here: That what is (1)communicated and understood by assertion is identical to (2)a common epistemological conception of knowledge (the ho hum view), which corresponds to (3)the way in which ordinary, competent speakers define 'knowledge'. This three-tiered identity is not altogether unproblematic. One worry arises at the level of the ordinary person's understanding of 'knowledge'. It seems that speakers may just as easily be convinced that knowledge involves elevated standards (rather than the moderate conditions on belief and justification that Rysiew envisions) or be persuaded to lower their standards. After all, it seems that this is precisely what is going on in the skeptical puzzle, where speakers are led to believe that knowledge of ordinary propositions requires the elimination of extremely unlikely possibilities, and what, consequently, generates the apparent paradox. That speakers should accept the ho hum view of knowledge seems important to Rysiew's argument. If speakers do *not* define 'knowledge' as an ungettiered justified true belief (which, according to Rysiew, is also what speakers are communicating through assertion), then Rysiew's next premise cannot be sustained. How so? Rysiew's next claim is that asserting 'I know that p ' (that is, using this ascription literally), is to state something that would be mutually obvious by asserting ' p ' alone, and therefore is a departure from the norms of conversation. This supposition cannot be upheld, however, if speakers convey that they have an ungettiered justified true belief with ' p ', but do not equate 'knowing that p ' with having an ungettiered justified true belief that p . They would then not be guilty of explicitly stating information that would be mutually obvious through asserting ' p ' alone since speakers would be communicating very different information with ' p ' and 'I know that p '. While this concern raises some doubt about Rysiew's claim, it does not eliminate the main worry against contextualism that Rysiew has highlighted, namely, that if speakers, generally have an implicit understanding that assertion involves knowledge (whatever that amounts to), then adding 'I know that...' seems to point out information that would be readily accessible to all participants. It is this issue that raises a significant worry for contextualists, because it suggests that ordinary speakers tend to use 'knows' in a non-literal manner.

himself as knowing (ho hum) that p . If this is mutually obvious to both speaker and audience, it would seem tiresome and pointless to utter ‘*I know that p* ’ to express the very same idea. Rysiew concludes that such considerations provide “very good reason to suspect that speakers might use ‘I know...’ to communicate something other than what the sentence they utter expresses”.⁴⁴ He suggests that speakers attribute knowledge in order to convey that the subject is in an especially strong epistemic position with respect to p , i.e., that she is certain that p or has no doubt as to whether p .

Rysiew goes on to write that insofar as speakers are conveying this type of information, (that they have attained some very high epistemic status with respect to p) they “are not ‘talking about’ knowledge at all” when they assert “I know that p ”, and concludes that “we can rightly refrain from letting those data determine the shape of the relevant epistemological theories”.⁴⁵ Rysiew then proceeds to offer a more detailed account of what speakers in fact intend to convey, which incorporates the salient counter-possibilities to p (those that have been mentioned or considered during the conversation).⁴⁶ In addition, Rysiew goes on to show that the same considerations apply to third-person knowledge attributions. These latter issues, while pertinent to Rysiew’s overall view, will not concern us here. What I want to focus on, rather, is the quite simple argument that proceeds from the knowledge account of assertion (loosely speaking) to the conclusion that ‘I know that p ’ is generally not used to communicate that the speaker knows that p . We might formalize his argument in the following manner:

R1: Asserting ‘ p ’ communicates that the speaker knows that p .⁴⁷

⁴⁴ Rysiew, “Speaking of Knowing,” 635.

⁴⁵ Ibid., 628.

⁴⁶ Ibid., 637.

⁴⁷ While the knowledge account of assertion, in one form or another, is a popular view (see Williamson, *Knowledge and its Limits*, Jason Stanley, *Knowledge and Practical Interests*, Hawthorne, *Knowledge and*

R2: Using ‘I know that p ’ literally communicates something that would be mutually obvious by asserting ‘ p ’ alone. (From R1)

R3: When speakers make explicit something that is already obvious, they intend to communicate something over and above the literal content of the statement. (Conversational norm)

C1: Therefore, speakers use ‘I know that p ’ to convey something different than its literal content. (From R2 and R3)

C2: Therefore, theorists cannot rely on ordinary uses of ‘I know that p ’ as evidence for a semantic analysis of ‘knows’. (From C1)

In my response, I will argue that C1, and subsequently, C2 does not follow from the premises listed.⁴⁸ Why not? In asserting ‘I know that p ’ the speaker does not in fact state something that is obvious prior to his assertion. Rather, ‘I know that p ’ states something that would be obvious were the speaker to assert p . (Note, that I am not denying R2, but rather, denying that ‘I know that p ’ is an instantiation of R3.) For this reason, the norm embodied in R3 is inapplicable to such utterances.

Rysiew begins his discussion of R3 by providing two examples in which speakers utter a statement that contains information which is clearly obvious to all the conversational participants, and known to be so by them. In the first case, two English speakers are walking through a heavily wooded forest and one of them states, “That’s a

Lotteries and DeRose, “Assertion, Knowledge and Context,” *The Philosophical Review* 111, no. 2 (April 2002): 167-203.) the view that speakers communicate that they *ho hum* know that p through asserting ‘ p ’ is not. I take it that this way of characterizing what we communicate by ‘ p ’ is defective in certain respects. First, in situations that seem to demand extremely high standards for knowledge (whether we cash this out in terms of truth conditions or conversational warrantability), it is not clear that the speaker is simply representing himself in meeting the conditions for *ho hum* knowledge. Say, for example, in an operating room. Here, it seems that by asserting that p , the speaker communicates that he knows that p according to rather stringent standards. But, in situations where the standards seem to be much lower, the speaker might not even meet the *moderate* conditions for knowledge required by the *ho hum* view. Variability in the conditions for assertion that correspond to the fluctuation of knowledge standards (taken as the truth conditions for ‘knows’) is discussed at length in DeRose’s “Assertion, Knowledge and Context”.

⁴⁸ DeRose has argued that ‘I know that p ’ conveys that the speaker knows that they know that p , but Rysiew rejects this interpretation on the grounds that in asserting ‘ p ’ itself, one represents oneself as justifiedly and truly believing that one justifiedly and truly believes that p . “For as we have seen, that one satisfies the credal-epistemic conditions is a commitment which derives from one’s personal conformity to [the Cooperative Principle which]... applies to what one represents as being the case” (Rysiew, “Speaking of Knowing,” 636). It indeed seems plausible that speakers communicate their knowledge of their knowing that p through asserting ‘ p ’.

tree”. Rysiew maintains that “in such a case, the natural inference for [one] to draw is that [the speaker] is remarking on the obvious in order to get across something over and above, or perhaps quite other than, the information literally expressed by the words [she] utters”.⁴⁹ This seems right. That is, if some proposition *o* is already known by conversational participants, and each participant is aware that the other already knows it, then asserting this very statement is not conversationally appropriate. This impropriety signals to the audience (perhaps tacitly) that that the speaker intended to communicate that the tree in question is particularly large or otherwise impressive, i.e., content that is not entailed by the meaning of the expressions in his statement. A second case involves our use of ‘looks’. Suppose that you and I are both looking at a basket of red apples and I say, “The apples look red.” I have indicated something that is clearly evident to both parties, and known to be so. Perhaps, then, I say this to convey that there is some doubt about the color of the apples.⁵⁰ There are then two steps involved in such exchanges. First, we recognize that the speaker has violated conversational norms by pointing out the obvious. Second, we reason that since the speaker seems to be engaging in rational communication, she must be intending to convey something over and above the semantic content of her utterance. Such communicative intentions are more apparent if the speaker stresses certain terms: ‘*That’s a tree!*’; ‘The apples *look* red.’

The common theme in the two examples, and what essentially motivates our intuition that the speakers intended to convey non-semantic content by their utterances is the fact that they remark on information that is clearly obvious to conversational participants prior to the relevant assertions. This particular feature, however, is absent, or

⁴⁹ Rysiew, “Speaking of Knowing,” 634.

⁵⁰ This example is attributed to H.P. Grice, “The Causal Theory of Perception,” *Proceedings of the Aristotelian Society, Supplementary Volumes* 35 (1961): 121-168.

at least significantly modified in ascriptions of belief and knowledge. Rysiew writes that since asserting ‘*p*’ alone would represent oneself as believing that *p*, speakers (as they appear to do in the preceding examples) would presumably exploit this fact by using ‘I believe that *p*’ to communicate something that is *not* inferable from ‘*p*’ alone—such as that they lack complete confidence of *p* or are not absolutely certain that *p*. And, as further evidence for his thesis, Rysiew notes that this is in fact the case, that is, speakers generally preface a proposition with ‘I believe that’ to communicate uncertainty about the truth of the proposition. Note, however, the subtle difference that emerges in this example. Here, the statement ‘I believe that *p*’ is not explicitly pointing out information that is already believed by all the conversational participants. Rather, it seems to be indicating something that *would be* obvious if the speaker were to utter a different statement, namely, ‘*p*’. Does commenting on information that *would be* obvious have the same effect on generating implicatures (or communicating non-semantic content) as commenting on information that is *already* obvious to speakers and hearers? As the following examples show, it is not at all clear that they have the same effects.

In their independent discussions of what it means to speak non-literally, Bach and Grice offer several examples in which speakers intend to convey, and are understood to communicate, information over and above the semantic content of their phrases. The interesting point to note, however, is that in these examples, if the speaker were to make such information explicit in her statement, she would not be guilty of veering from the norms of conversation. More precisely, we would not have to infer that the speaker intended to convey pragmatic content, because she said something that would have been obvious had she uttered a different statement. Our first example comes from Bach:

(1) 'I haven't taken a bath'

communicates that the speaker has not had a bath that day, or perhaps since she arrived.⁵¹

However, suppose, the speaker were to utter

(1') 'I haven't taken a bath today'.

Or

(1'') 'I haven't taken a bath since I arrived.'

Admittedly, such temporal specifications would be obviously inferable had the speaker simply asserted (1). However, she does not seem to have said anything that is conversationally awkward by specifying a time constraint. Subsequently, it does not seem necessary to suppose that that the speaker intended to convey something different than the literal content of her statement, unlike Rysiew's first two examples. Rather, it seems that (1) is simply a different way of communicating (1') and (1'').

Bach offers another example:

(2) 'Nobody goes there anymore because it is too crowded'

is used to convey that nobody *important* frequents the venue in question, or that nobody *in our circle* goes there. But again, the following explicit formulations seem perfectly acceptable:

(2') 'Nobody important goes there anymore because it is too crowded'.

(2'') 'Nobody we know goes there anymore because it is too crowded'.

The qualifier on 'nobody' is implicit in (2), and would be obvious to the audience were the speaker to utter (2). Despite this, neither (2') or (2'') appear conversationally improper. There is no need to suppose, or no follow-up intuition, that the speaker is

⁵¹ Kent Bach, 'The Semantics-Pragmatics Distinction: What it is and Why it Matters' in *The Semantics-Pragmatics Interface from Different Points of View*, ed. Ken Turner (Oxford: Elsevier, 1999): 6.

attempting to convey some non-semantic content. Rather, it seems that the information conveyed in (2) might be communicated via (2') or (2''), depending on the speaker's intentions. Bach's examples involve a missing element that would make the relevant utterance more specific, but we do not have to rely on just this type of case to demonstrate the problem with Rysiew's claim.

One last example comes from Grice:

A: I am out of petrol

B: There is a garage around the corner.⁵²

According to Grice, B uses his assertion to communicate that the garage is open and has petrol to sell, otherwise she would be infringing the maxim of Relation. But suppose B were to assert instead:

B': There is a garage around the corner that's open now.

While it is true that B's second statement states explicitly what would be understood by asserting his first, it does not seem that B has uttered something conversationally defective. Therefore, we need not assume that the speaker intended to convey information over and above the literal content of the utterance. These examples are intended to show that explicitly stating what *would be* obvious by asserting a different sentence does not generally have the same effect as explicitly stating what *is* obvious. With regard to the latter, there was clearly something strange about the speaker's assertion, which prompted the intuition that he used the relevant phrase ("That's a tree"; "The apples look red") to communicate additional, non-semantic information. However, no such intuition is present in the cases where speakers explicitly state information that would be implicitly

⁵² Grice, Paul Grice, *Studies in the Way of Words* (Cambridge: Harvard University Press, 1989), 51.

conveyed by using a different utterance. Hence, there is a significant difference between the two types of cases.⁵³

One might, however, argue that it is an empirical fact that speakers often assert ‘I believe that *p*’ to communicate that the speaker is not very sure of *p*. Indeed, DeRose identifies ‘I believe’ as a “hedging maneuver” used to indicate that the speaker has attained a less-than-optimal epistemic status with respect to the proposition in question.⁵⁴ Unarguably, this information diverges from the literal content of ‘belief’. Are there then exceptional cases in which the *would be* obviousness of asserting ‘*p*’ alone *is* sufficient to generate an implicature? I want to argue that the implicature is produced in this case not because the sentence literally expresses what would be inferable by asserting only the embedded proposition, but for a different reason, one that Rysiew himself alludes to in a previous example. A speaker communicates additional information in uttering ‘I believe that *p*’ in light of the fact that he would otherwise flout the conversational principle that Grice articulates thus: ‘One should not make a weaker statement rather than a stronger one unless there is a good reason for so doing’.⁵⁵ I want to suggest that in saying ‘I believe that *p*’ the speaker asserts something weaker than what would be communicated through asserting ‘*p*’. In so doing, he generates the implicature that he is not certain or not very confident that *p*. To show that this is the case, we must have some understanding of what it means for one assertion to be weaker/stronger than another.

⁵³ Grice’s second submaxim of Quantity instructs speakers to limit their contribution to be as informative as required, which seems to underlie Rysiew’s claim that asserting a statement that is already mutually obvious is not constitutive of communicative rationality. Again, however, this rule seems to apply to what is *already* known or believed by conversational participants, rather than what would be believed, given a different locution. On one interpretation, the submaxim is read as “Do not say what people already have good reason to believe” (Hawthorne, *Knowledge and Lotteries*, 25). In the case of ‘I believe that *p*’ and ‘I know that *p*’, however, the participants do *not* already believe that *p* or know that *p*. Therefore, the speaker is not flouting Grice’s second submaxim of Quantity by asserting ‘I know that *p*’.

⁵⁴ DeRose, “Assertion, Knowledge and Context,” 187.

⁵⁵ Grice, “The Causal Theory of Perception,” 132.

Grice's discussion of what determines the strength of a sentence is equivocal and muddled. While on one occasion, he identifies the scale as corresponding to the direction of logical entailment, on another he asserts that one statement might be weaker simply because "one has... a strong inclination to regard the first of these statements as weaker than the second".⁵⁶ Other theorists, however (including Rysiew himself) have defined the relative strength of a statement in terms of the degree of informativeness.⁵⁷ That is, an assertion *A* is stronger than an assertion *B* if and only if *A* communicates more information than *B*. This characterization seems to explain Grice's own use and makes sense of the examples he provide.

Now we can understand why the speaker communicates additional, non-semantic information when he states, 'I believe that *p*'. In asserting that *p*, the speaker represents himself as knowing that *p*, as Rysiew and several others have argued. Hence, it communicates to the audience that the speaker is in an epistemic position that qualifies as knowledge. However, if the speaker were to utter instead, 'I believe that *p*' the speaker communicates something weaker,⁵⁸ that he has met at least one (but perhaps not all) of the components of knowledge.⁵⁹ The idea is that the information communicated by

⁵⁶ Ibid., 140.

⁵⁷ See Rysiew, "Speaking of Knowing" FN 12; Also see Tim Black, ("A Warranted Assertability Defense of a Moorean Response to Skepticism," *Acta Analytica* 23 (2008): 187-205) in which he (following Levinson) interprets Grice's second submaxim of Quantity as a restriction on strength. In "Classic Invariantism, Relevance and Warranted Assertability Maneuvers," *Philosophical Quarterly* 55, no. 219 (April 2005): 328-336, Black accounts for the data used by contextualists by stating that in allegedly higher-standard or 'tougher' contexts, where remote counter-possibilities to *p* have surfaced, our knowledge ascription would be "stronger... in that it would be more informative." This is because 'I know that *p*' in such circumstances would implicate that the speaker occupies an epistemic position that rules out the salient counter-possibilities, thereby making his assertion "too strong and unwarranted," 334.

⁵⁸ I take it that this claim does not depend on any particular conception of knowledge or belief insofar as we accept the widely shared view that belief is a component of knowledge. If one were inclined, however, to reject this, then perhaps the same conversational principle might not apply.

⁵⁹ The same charge cannot be leveled against theorists who defend the view that belief or rational belief is the norm of assertion. This is because the information inferred by *p* would not be any stronger than the proposition literally expressed by 'I believe that *p*'. However, Rysiew seems to adhere more closely to the

asserting ‘ p ’ (I know that p) contains more information than the proposition literally expressed by “I believe that p ”. For this reason, the audience reasons that the speaker has uttered a weaker statement to convey some information that is not contained by the semantic content of ‘I believe that p ’, such that the speaker is less than certain about the truth of p . The following example illustrates this point: Toni asks Millie, “What time does the bank close?” Millie responds, “I believe this bank is open until 5”. Toni might (perhaps tacitly) reason in the following way: “Why did Tony say ‘I believe that the bank is open until 5’ rather than simply telling me when the bank closes?’ She must not be sure about the bank’s hours, but is only stating what she *thinks* might be true.” The speaker then asserts something weaker for a reason, namely, to generate the implicature that she isn’t certain as to the truth of p . We can accommodate the intuition that “I believe that p ” generally conveys extra, non-semantic content, but not because it states something that *would be* obvious by asserting p alone, but rather, because it literally expresses a proposition weaker than that conveyed by p .

What then should we say about ‘I know that p ’? According to Rysiew, speakers intend to communicate non-semantic content by using ‘I know that p ’ because the content literally expressed by this statement would be inferable by asserting p alone. However, as noted above, information that *would be* obvious does not have the same affect on communication as information that is already obvious. Moreover, there is no general presumption (as in the case of ‘I believe that p ’) that speakers communicate additional content by saying ‘I know that p ’. Although Rysiew argues that speakers attribute

knowledge account of assertion. Moreover, if Rysiew instead endorsed the view that the norm of assertion is belief, then presumably, he would have to also accept the claim that the speaker communicates that he *believes* that p , rather than knows that p in asserting ‘ p ’. But this view would not help Rysiew in that he could not ultimately conclude that ‘I know that p ’ literally expresses content that is inferable by p alone.

knowledge to themselves to convey a very strong epistemic status (more than that required for knowledge), this does not show that ‘I know that p ’ in fact communicates such information for two reasons. First, we can devise any number of counter-examples to this claim, which function to question its accuracy. That is to say, attributions of knowledge in many cases can seem warranted and true (say in casual conversation) without the speaker or the audience believing that the speaker has achieved some extremely strong epistemic status. We can see this in DeRose’s bank cases. In the first scenario, the husband has some, but not an overwhelmingly high degree of evidence that the bank will be open on Saturday. Nevertheless, his insistence that he knows that the bank is open seems appropriate and true.

Second, Rysiew’s claim that knowledge attributions communicate such information is based on his prior argument that speakers intend to use knowledge ascriptions non-literally. Given that the latter was based on an incorrect premise, he loses support for the former. Thus, I do not believe that Rysiew has successfully shown that speakers use knowledge ascriptions in a non-semantic manner, and, consequently, has not provided any reason to believe that contextualists ought to refrain from relying on ordinary uses of ‘knows’ for their semantic theory. More broadly, that speakers tend to use terms loosely or speak in ways that diverge from the literal content of their assertions does not serve, in and of itself, as a preemptive threat to contextualist methodology.

IV. Summary of Chapters

The main thrust of classic invariantism comes from the claim that invariantists can account for fluctuations in uses of ‘knows’ without adopting a context-sensitive analysis of the term. As stated earlier, their strategy typically involves an appeal to the

pragmatic implications of knowledge assertions and denials to accommodate such variability. There are then two aims for my discussion in the following chapters: First, I will chip away at the general strategy of invoking pragmatic implications to explain away unfavorable intuitions, thus undermining the classic invariantist's capacity to explain the phenomenon in question: the variability of our knowledge ascriptions and denials.⁶⁰ After arguing that this method is ultimately inadequate to explain the context-sensitivity of 'knows' in chapters two through four, I offer a positive analysis of the expression's context dependence in chapter five. In this section, I will present a brief overview of my discussion.

One of the main attractions of epistemic contextualism is its ability to offer a solution to the skeptical puzzle. As stated earlier, contextualists maintain that the standards of knowledge are raised to a high level due to the introduction of the skeptical hypothesis, which makes it the case that we neither know that we are not brains in vats nor that we have hands in this context. Contextualists further claim that outside of this context and relative to ordinary standards of evaluation, we know a great many propositions including the denial of the skeptical hypothesis and less extraordinary claims. The argument, then, presents no paradox at all. But such a solution puts the contextualist in a precarious position: they are forced to provide an error theory to account for the mistaken intuition that the argument *is* puzzling, one which does not attribute massive semantic (or other kinds of) ignorance to competent speakers. In the second chapter, I will examine this issue more thoroughly discussing, in particular,

⁶⁰ To be sure, this is importantly different from the argument I addressed in the previous section, in which Rysiew offers reason to believe that we are using 'knows' nonsemantically. In the following chapters, I argue against invariantists' interpretation of the pragmatic implications of 'I know that *p*'.

Patrick Rysiew's charges against a common contextualist error theory and his invariantist alternative.

Rysiew argues that contextualists impute a large degree of error to ordinary speakers, and for this reason should be abandoned in favor of an analysis that explains the variability of 'knows' without assuming considerable speaker incompetence. He proposes an invariantist alternative according to which 'knows' corresponds to a single and stable definition across all contexts, but speakers *use* knowledge attributions and denials in a way that diverges from their literal content. According to Rysiew, 'S knows that *p*' literally means that *p*, S believes that *p*, and that S has eliminated all the relevant $\sim p$ alternatives, where the standard of relevance is relatively undemanding and fixed. It seems, however, that in certain cases (say where there is a lot at stake), it sounds right to deny that a subject has knowledge, despite meeting these standards. To account for this intuition, Rysiew maintains that speakers say 'S does not know that *p*' in order to express S's epistemic status with regard to the alternatives that have been introduced in the conversation or the *salient* alternatives, which does not bear on whether or not S actually knows that *p*. In short, Rysiew appeals to the pragmatic content of 'S does not know that *p*' in order to explain away our intuition that such denial seems correct in particular cases.

I argue that this strategy ultimately fails as Rysiew's own account is committed to attributing widespread semantic ignorance to ordinary, competent speakers. Specifically, in his account, speakers are unaware of the semantic content of 'knows' as he defines it. Importantly, it is not simply that they happen to use the term incorrectly when they are not careful about their assertions, but they are also unable to recognize the meaning of 'knows' when its content (as conceived by Rysiew) is explicitly articulated. Rysiew's

invariantist theory and pragmatic explanation, then, do not properly suffice to explain our use of 'knows'. I further offer a different contextualist error theory that accommodates our mistaken judgments in the skeptical argument. This chapter therefore recommends contextualism as a preferred analysis of 'knows' by showing that a moderate invariantist's pragmatic efforts to explain the variability of our knowledge ascriptions and denials fail.

In chapter three, I focus on skeptical invariantism. While skeptical or high standard scenarios pose no theoretical problems for this brand of invariantism, as skeptical invariantists openly accept that there is very little if anything we know, they have a difficult time accounting for knowledge attributions in everyday situations. In particular, they cannot accommodate the subtler epistemic distinctions that we care about in non-skeptical situations, due to their excessively high demands. It seems, for example, that I know President Obama has decided to send additional troops to Afghanistan because I have heard this repeated in various news programs, and listened to Obama himself defend his assessment in a recent speech. My toddler, however, does not know this because she has not heard the same reports or speeches and even if she did, she would not understand them. Now, neither of us can eliminate the possibility that a hidden, but powerful force has interfered with our media coverage to lead us to have false beliefs, or any other remote counter-possibilities. Because of this failure on both our parts, skeptical invariantists would deem that neither my toddler nor I know that President Obama has decided to increase troops in Afghanistan. This classification appears preposterous.

Skeptical invariantists explain our strictly mistaken inclinations to ascribe knowledge by focusing (much like their moderate counterparts) on the conversational acceptability of these utterances, or essentially utilizing what DeRose has termed warranted assertability maneuvers (WAMs).⁶¹ That is, while most knowledge ascriptions are literally false, they may nevertheless be warranted insofar as they communicate content that pertains to the current purpose. This explains why we are wont to attribute knowledge so readily in most ordinary circumstances. Peter Unger has argued that *this* type of variability (of the non-semantic variety) is not unique to ‘knows’, but in fact quite common in our language: ‘Vacuum’, ‘flat’, ‘certain’, etc., each have a fixed meaning, but is used to indicate varying degrees of the property for different pragmatic purposes. It seems, however, that this leaves open the possibility that *any* term whose everyday use appears context-sensitive is merely pragmatically variant, but semantically stable. To avoid this charge, Unger offers tests by which to distinguish genuine semantically variable expressions, which he calls the paraphrase and emphasis tests. In this chapter, I examine Unger’s approach to determine whether or not he has provided a viable means to separate terms such as ‘know’, ‘flat’ and ‘certain’, from relative words such as ‘tall’, ‘short’, ‘rich’, etc. I ultimately conclude that Unger fails to provide sufficient support for his tests. Without well-founded, reliable tests, deeming ‘knows’ pragmatically, rather than semantically, variable is arbitrary and unjustified. Here, I show that a skeptical invariantist’s attempt to explain away our knowledge-ascribing behavior by appealing to pragmatic content is unsuccessful.

⁶¹ Keith DeRose, “Contextualism: An Explanation and Defense,” in *The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa (Oxford: Blackwell Publishers, 1998).

The fourth chapter examines a particular argument in defense of low-standard invariantism. Unlike skeptical invariantists, low-standard invariantists do not have to defend unachievable epistemic demands. However, there are certain occasions for which this account appears ill-suited; in particular, cases that require additional support or evidence before a subject can truly be ascribed knowledge. Cohen's airport cases illustrate this problem:

Mary and John are at the airport discussing whether they ought to take a particular flight to New York. They overhear someone asking a passenger, Smith, if he knows whether the flight has a layover in Chicago. Smith consults his itinerary and says, 'Yes, I know—it does stop in Chicago'. It turns out that Mary and John have an important business meeting in Chicago. Mary says to John, 'How reliable is that itinerary? It could contain a misprint. They could have changed the schedule last minute'. Mary and John agree that Smith does not know that the plane will stop in Chicago and subsequently decide to check with the airline agent.⁶²

If, as the invariantist claims, the standards of knowledge are in fact always low, then Mary and John are mistaken in denying that Smith knows. Nevertheless, it seems prudent for them to check with the airline agent, given their desire to attend the business meeting. But Cohen maintains that it sounds inconsistent to say, 'Smith knows that *p*, but we need to check further' or 'We know that *p*, but we need to check further'. Cohen argues, in effect, that it is semantically inconsistent to both ascribe knowledge that *p* and express the need for more evidence. It appears, then, that Mary and John are correct to deny that Smith has knowledge in their context. This raises doubts against an invariantist account that posits fixed and low standards of knowledge across *all* contexts.

Tim Black and Peter Murphy (BM), however, have argued that there may very well be situations in which subjects have met low standards for knowledge, but

⁶²Stewart Cohen, "Contextualism, Skepticism, and the Structure of Reasons," 58.

nevertheless need more information for some other purpose.⁶³ They provide a series of examples in which speakers felicitously claim to know that p while needing to check further to attain a different goal. The oddity of such assertions ('I know that p , but I need to check further'), then, disappears when the speaker cites a different purpose for which they need to check further: a grant proposal, journal publications, etc. Essentially, they show that the implication from 'I know that p ' to 'I need not check further' can be cancelled. According to Grice's cancellability criterion, this means that the two statements are only pragmatically related, rather than (as Cohen maintains) semantically linked. In other words, 'I know that p , but I need to check further' is not a semantically inconsistent statement, though it might *sound* awkward without the appropriate specification. This seems to redeem low standard invariantism, as it allows the subject to have both met the relatively low standards for knowledge, while also recognizing the need for more evidence for some other purpose. I argue, however, that statements of the form 'I know that p , but I need to check further to x ' contain an implicit qualification. Specifically, the speaker intends to convey that that she knows that p relative to one standard, but does not know that p and must check further relative to a higher one. Such statements, then, qualify as cases of loose talk. Grice's cancellability criterion will not effectively distinguish pragmatic implicatures from semantic entailments in instances of loose talk. I argue that because speakers in Black and Murphy's examples are speaking loosely, they have not conclusively shown that 'I know that p ' and 'I need not check further' are only pragmatically related. Therefore, we can sustain Cohen's thesis that it is semantically inconsistent for a speaker to both claim to know a proposition and that she

⁶³Tim Black and Peter Murphy, "Avoiding the Dogmatic Commitments of Contextualism," *Grazer Philosophische Studien* 69 (2005): 165-182.

needs more evidence for that proposition in one and the same context. Hence, BM's appeal to pragmatic implication to support their position fails.

Chapters two, three and four, then, attempt to undercut the driving force behind invariantism by showing that appeals to pragmatic content and implication ultimately do not suffice to defend an invariantist analysis of 'knows', whether the account poses high, moderate or low standards for knowledge. Contextualists seem to have an upper hand with ordinary uses of 'knows', as they take our intuitions about the term at face-value. Intuitions, alone, however, cannot sustain a theory. We need some understanding of the way in which 'knows' is context-sensitive, and an appropriate linguistic basis for its context-sensitivity. Without this, the theory is incomplete, and may even appear implausible. This is what I aim to provide in the fifth chapter.

Contextualists draw analogies between 'knows' and other context-sensitive expressions in order to show that the semantics they attribute to 'knows' is not unique, but rather applies to a particular class of expressions of which 'knows' is a member. Further, by identifying the term within a broader linguistic category, contextualists can show the manner in which 'knows' is context-sensitive. The emerging trend with respect to such identification is to classify 'knows' as a non-gradable context-sensitive expression. In this chapter, I will argue, however, that such an analysis is mistaken, as propositional knowledge ascriptions can come in degrees within a contextualist framework (i.e., where the truth conditions of such statements are sensitive to the context of utterance). In certain high-standard contexts, 'S knows that p ' is true if and only if S knows why p . Knowledge-why, I will argue, can come in different degrees. Propositional knowledge ascriptions, then, can come in different strengths under certain

circumstances. Therefore, I will conclude that the current classification of 'knows' must be rejected in favor of a different one. Specifically, 'knows' must be identified with gradable context-sensitive expressions, which include members such as 'tall' and 'rich'.

The following discussion is intended to demonstrate that epistemic contextualism offers a far superior analysis of our knowledge-ascribing behavior than the various invariantist accounts that have been proposed. In order to sustain fixed standards for knowledge across contexts, invariantists have supplemented their semantic theories with pragmatic explanations that aim to show why speakers might be inclined to speak *as though* the conditions for knowledge vary. This involves a distinction between what speakers convey (the pragmatic content of their assertions) and the proposition literally expressed by their utterance. I will argue that such pragmatic explanations are insufficient to properly account for the systematic variability of our knowledge attributions and denials. Further, I will show that we can identify 'knows' within the broader linguistic category of gradable context-sensitive expressions, thereby reinforcing the view that a contextualist semantics is preferable to classic invariantist theories.

Chapter Two

Contextualist and Invariantist Error Theories

Contextualists maintain that the content of knowledge ascriptions and denials vary according to context. While one of the primary strengths of contextualism is its capacity to treat our variable intuitions and judgments about knowledge assertions charitably, it has been argued that there are some intuitions that contextualism fails to accommodate.⁶⁴ A particularly threatening case involves our reaction to the skeptical argument. Patrick Rysiew has suggested that one contextualist explanation of why the skeptical argument is apparently worrisome commits her to an implausible error theory. This, however, raises a more general problem, namely, whether we can make sense of contextualist data in a way that avoids an implausible error theory.

Rysiew suggests that we might account for our varying knowledge assertions by distinguishing between what is pragmatically conveyed from the semantic content of such utterances. This distinction allows him to accommodate contextualist intuitions while avoiding the contextualist's error theory. He then proposes a theory that explains why our varying judgments about knowledge ascriptions seem right, and offers an invariantist semantic analysis of 'knows'. The viability of Rysiew's semantic theory, as well as the attack on contextualism seems to depend, in large part, on the way in which Rysiew conceives of the pragmatic content, or what is understood by knowledge assertions. Here, I will argue that what is in fact intended and conveyed by our

⁶⁴ See for example, Patrick Rysiew, "The Context-Sensitivity of Knowledge Attributions," *Nous* 35, no. 4 (2001): 477-514; Stephen Schiffer, "Contextualist Solutions to Skepticism," *Proceedings of the Aristotelean Society* 96, no. 3 (1996): 317-333.

knowledge assertions involves more than Rysiew recognizes. Secondly, I will offer an alternative contextualist explanation that accounts for our misguided judgments in the skeptical argument, which circumvents the worries raised by Rysiew.

I. Contextualism

Basic Idea

Contextualists maintain that the content or truth conditions of knowledge assertions of the kind ‘S knows that p ’ or ‘S doesn’t know that p ’ vary according to the context of the speaker. Certain features of the speaker’s context determine how good an epistemic position S must be in with regard to p in order to know that p . Thus, according to the theory it is perfectly consistent for a speaker to affirm that S has knowledge that p in one context, perhaps where the standards for knowledge are quite low, and deny that the same subject possessing the same amount of evidence knows p in another context that demands greater epistemic strength with respect to p . The best support for such a theory comes from our linguistic behavior in everyday, non-philosophical contexts. As the ever-growing number of examples⁶⁵ presented by contextualists and their critics show, speakers tend to raise standards for ascribing knowledge in the presence of certain conversational features such as the mentioning of counter-possibilities, or the importance of being right. Keith DeRose’s bank cases exemplify this behavior.⁶⁶ In Bank case A, a man and his wife arrive at the bank to deposit their checks only to find a long queue. The man suggests that they return to deposit their checks tomorrow. His wife points out that

⁶⁵ Take for example, Matt McGrath’s train cases in "Evidence, Pragmatics and Justification," *Philosophical Review* 111, no. 1 (2002): 67-94. Stewart Cohen’s airport cases in "Contextualism, Skepticism and The Structure of Reasons," *Philosophical Perspectives 13: Epistemology*, ed. James E. Tomberlin (Atascadero: Ridgeway): 57-89.

⁶⁶ Keith DeRose, "Contextualism and Knowledge Attributions" *Philosophy and Phenomenological Research* 52, no. 4 (December 1992): 913-929.

many banks are closed on Saturdays. But, the man responds, “I was here two weeks ago on Saturday, so I know the bank will be open tomorrow”. In Bank case B, the couple find themselves in the same situation, and the man once again suggests coming back the next day. But this time, they have just written a very large check, and face a rather severe penalty if their check is not deposited. His wife raises these points and says further, “Banks do change their hours, do you know that the bank will be open tomorrow?” At this point, the husband says, “Well, no, I’d better check to make sure”. The husband’s evidence for p , and subsequently, his epistemic position with regard to p remained constant. Nevertheless, we are inclined to accept the husband’s knowledge ascription in the first case and his denial in the second. Such behavior supports a contextualist analysis of knowledge ascriptions. While some situations demand extremely high epistemic standards, in other situations, we might truly ascribe knowledge to S despite S ’s relatively low epistemic position with regard to p .

Skepticism

A contextualist analysis of ‘knows’ also serves to provide a solution to the skeptical argument (SA), whose conclusion seems to show that we lack knowledge of ordinary propositions, so ordinary in fact, that we typically take such knowledge for granted. Here is a bare formulation:

1. I don’t know that $\sim H$.
2. If I don’t know that $\sim H$, then I don’t know that O .
3. I don’t know that O .⁶⁷

⁶⁷ This formulation of the skeptical argument can be found in Keith DeRose, “Solving the Skeptical Problem,” in *Skepticism*, eds. Keith DeRose and Ted A. Warfield (New York: Oxford University Press, 1999), 183-219.

H stands for a skeptical hypothesis such as the proposition that I am a brain in a vat electro-chemically stimulated to have normal impressions and experiences (BIV hereafter). O is meant to be an ordinary proposition such as ‘I have hands’.

According to the contextualist, the skeptic automatically raises the standards for knowledge by asserting the possibility of being a BIV. Moreover, it seems equally clear that one is in as strong an epistemic position with regard to O as she is with regard to \sim BIV. The second premise can be understood as a ‘comparative conditional’ used to compare the epistemic position of a subject with respect to different propositions, and it seems true that we are in fact in “at least as good a position to know that the [skeptical] hypothesis is false as we’re in to know the targeted piece of ordinary knowledge.”⁶⁸ By inflating our ordinary standards in the first premise, the skeptic is able to conclude that we don’t know, relative to *such* standards, that we have hands. While the contextualist assents to the truth of the skeptical argument, she denies that this undermines our knowledge in everyday contexts. In everyday contexts, the subject knows that she has hands as the standards for the truth conditions of knowledge ascription are much lower. However, the epistemic strength of her position with regard to this proposition fails to meet the skeptic’s standards. This allows the contextualist to preserve our knowledge ascriptions in ordinary contexts, while recognizing the challenge presented in skeptical contexts.

While the analysis seems to resolve the skeptical puzzle, there is a lingering worry. Because the contextualist interprets the argument as a set of mutually consistent statements, all of which are equally plausible, she must provide an explanation of why it

⁶⁸ DeRose, “Solving the Skeptical Problem”, 203.

is that we nevertheless *think* that the skeptical argument is troubling or paradoxical.⁶⁹ Some have suggested that the only response available to the contextualist is to claim that ordinary speakers *misunderstand* the content of these assertions, in the sense that we do not realize that the skeptic's use of 'knows' invokes extraordinarily high standards. This failure causes speakers to resist the conclusion, since in quotidian contexts, she does have such knowledge. Patrick Rysiew (2001) contends that such an error theory is implausible because it commits the contextualist to the further claim that speakers are not aware of their *own* communicative intentions and interests—as such factors determine the content of the assertions. In the next section, I will further expound Rysiew's criticism and present his preferred analysis, which he takes to avoid the problems associated with contextualism. Then, in the third section, I will argue that his analysis does not accurately capture what is communicated by knowledge assertions. A similar argument will also show that Rysiew is himself committed to an implausible error theory, namely, that speakers are not aware and cannot become aware of the semantic content of 'knows'. In the last section, I will present an alternative contextualist explanation to account for our misguided intuitions in the skeptical argument, and argue that it circumvents Rysiew's objection by attributing minimal error to the competent speaker.

II. Rysiew's Picture

Contextualist Response to Skepticism

The contextualist must explain why the skeptical argument might appear puzzling even though, according to contextual analysis, it is not. Rysiew, following Stephen Schiffer (1996), assumes that the only plausible contextualist explanation is that we

⁶⁹ Originally, I think the worry was raised in Stephen Schiffer, "Contextualist Solutions to Skepticism," *Proceedings of the Aristotelian Society* (1996): 317-333. Rysiew modifies this objection so that it not only generates worries for contextualism, but that it points toward his own invariantist analysis.

misunderstand what is being asserted in the skeptical argument. That is, we think that the skeptic is denying knowledge that we have hands in ordinary standards, assuming that what the skeptic means when he talks about knowledge is what *we* mean by the term. The skeptical argument appears puzzling, then, because we certainly don't want to conclude that we don't know that we have hands relative to ordinary standards.

Such a mistake seems reasonable enough when the role of the skeptic is played by an opponent or any interlocutor, but Rysiew contends that this error theory is entirely implausible when we ourselves are asserting the skeptical argument. Recall that according to the contextualist, anyone who asserts the first premise raises the standards for knowledge and automatically shifts the content of 'knows'. Contextualists further maintain that the propositions expressed by these assertions are determined by features of the speaker's conversational environment, including his interests, communicative intentions, goals, etc. It appears then, that our failure to understand the literal content of our assertion would entail ignorance about our own interests and communicative intentions. We wouldn't know *what we mean* by our assertions. That we might be mistaken about our own communicative intentions is the implausible error that contextualists must attribute to ordinary speakers.

Prima facie, such a charge might appear misdirected. It seems that we might acknowledge our intentions and interests, but still be mistaken about the content of our assertions due to some degree of 'semantic blindness'. That is, the contextualist could respond that we fail to acknowledge the rules that govern our assertions, which causes us to make knowledge claims that would be true in low-standard situations while false in high-standard contexts. Ordinary speakers are simply confused about the way in which

such features determine truth conditions. The contextualist, however, cannot appeal to such semantic blindness, because it is part of their theory that the prevalent standards are transparent to conversational participants. There is implicit agreement about the standards for the use of 'knows' both relative to high and low standards:

In neither LOW nor HIGH is there any disagreement among any speakers about whether any subject knows. When S-HIGH [speakers are in high standards] says that the subject does not 'know' she meets with no resistance at all from their conversational partners, who have also sensed that unusually stringent standards are appropriate in the context, and have adjusted their use of 'knows'[There is] no element of arm-twisting, arguing or convincing, but only of informing ...that one subject does not 'know' something.⁷⁰

Thus, while it might be the case that speakers are semantically blind to some degree, this should not affect their capacity to identify the correct truth conditions for knowledge assertions. Rysiew's worry then is sustained. That is, if we *are* mistaken about the literal content of a knowledge utterance, then we must also be mistaken about the features that fix the content, such as our own communicative intentions. This is because insofar as we know our practical interests, etc., contextualists maintain that we automatically acknowledge the corresponding standards. Thus, ignorance about the latter would seem to show ignorance about the former. Indeed, this seems peculiar. It appears then, that while contextualism offers one way to account for our variant judgments about knowledge ascriptions, it suffers from an implausible error theory, a defect that demands a different explanation of our intuitions.

Rysiew's Alternative

Rysiew offers an alternative explanation to account for our varying knowledge assertions. Instead of allowing that the semantic meanings of these assertions shift, Rysiew maintains that it is what *we understand or convey* by these assertions that

⁷⁰Keith DeRose, "'Bamboozled by our own words': Semantic Blindness and Some Arguments against Contextualism" available on website: <http://pantheon.yale.edu/%7Ekd47/Bamboozled.doc>.

changes according to our context. He explains that we typically identify ‘what is said’ in an utterance with a “conceptually and pragmatically enriched proposition closely related to the proposition literally expressed—the proposition the speaker is most plausibly regarded as having intended to convey”.⁷¹ While the literal content of an assertion depends on the ‘syntactic character’ of the sentence, we tend to grasp the most *salient* proposition. For example, in uttering ‘I’ve had breakfast’, the speaker does not literally indicate a particular time or day that this occurred. However, we usually take this assertion to ‘say’ that the speaker has had breakfast *today*. This is what is ‘said_{Loose}’ according to Rysiew. In the same way, speakers will tend to identify ‘what is said’ in a knowledge assertion *not* according to its literal content, but with the most salient proposition, making it the case that what we understand in a knowledge assertion does not always reflect its literal content. This suggests that our knowledge ascriptions may vary because of what is understood or conveyed by speakers, rather than the semantic content of these assertions. We can now understand our intuitions in the bank cases in a different way: When we believe that the speaker has said something true in the first *and* the second bank case, in claiming ‘I know that *p*’, and ‘I don’t know that *p*’ respectively, we mean that what he said_{Loose} was true.

Rysiew employs a relevant alternatives (RA) framework to distinguish the pragmatic and semantic content of knowledge assertions. According to his brand of RA, a knowledge ascription is true if the subject has eliminated all the *relevant* counter-possibilities to what is known. The favored standard of relevance remains constant: they are limited to what *normal* humans take to be likely alternatives. Thus, if our subject has

⁷¹ Rysiew, “Context-Sensitivity” 486.

eliminated $\sim p$ possibilities that are likely from a normal human perspective, then it will be true that S knows that p in a low-standards context, as well as high-standards contexts, say, where an unlikely $\sim p$ possibility has been raised:

The mere mentioning of a certain not- p alternative, or the mere fact that the speaker has certain not- p possibilities in mind when he says, “S knows that p ” doesn’t affect what that sentence means (what it literally expresses).⁷²

What the sentence literally expresses, according to Rysiew, is that the subject believes that p , p , and that the subject is in a good enough epistemic position to know that p in virtue of being able to eliminate the normal counter-possibilities. Why then are we inclined to change our judgment in the face of additional counter-possibilities? That is, why do we judge that the speaker knows that the bank will be open in the first case, and then deny it in the second? Rysiew maintains that while certain unlikely counter-possibilities that emerge during the course of conversation do not affect the semantic meaning of knowledge assertions, they function to change their *pragmatic* content, or what *we understand* by these assertions. The speaker asserts that he knows (or does not know) that p in response to the possibilities mentioned to convey that his epistemic position with respect to p is such that he can (cannot) eliminate these possibilities. The audience judges that knowledge ascriptions are true when the subject has eliminated the salient possibilities. Although such $\sim p$ alternatives are not relevant to the semantic content of the utterance, they are nevertheless salient to the conversational participants. Thus, in the second bank case, it seems right for the speaker to claim ‘I don’t know that the bank is open tomorrow’, because his epistemic position is evidently not good enough to eliminate the salient possibility that his wife just mentioned, namely, the bank could have changed its hours. Asserting the literal truth, however, (that he knows that the bank

⁷² Ibid., 488.

is open) would have imparted that he can rule out this possibility. We tend to take a knowledge ascription to ‘say’ that the subject has ruled out all the salient possibilities, and deny its truth when we see that the subject has not done so.

Rysiew’s account appears to accommodate our varying knowledge ascriptions while avoiding the pitfalls of contextualism. The variability of our assertions is due to the pragmatic content of the utterances, or what they convey with respect to the salient counter-possibilities. Their semantic content concerns the subject’s relation to the relevant alternatives, thereby making knowledge ascriptions true if and only if the subject has eliminated these alternatives and, subsequently, is in a good enough position to know that *p*. In his account, speakers are not mistaken about what *they mean*. They might, however, make knowledge claims that are literally false because they are not attentive to, or “don’t much care about” the purely semantic content of their assertions.⁷³ Thus Rysiew seems to have accounted for contextualist data without attributing unreasonable error to the speaker.

III. What is Said

In this section, I will argue that Rysiew has not offered a preferable alternative to contextualism, as his own account is committed to an implausible error theory, namely that competent language speakers are not aware and cannot easily become aware of what he takes to be the semantic meaning of their knowledge assertions.⁷⁴ Rysiew’s argument against contextualism and for his own invariantism depends largely on his conception of

⁷³ Ibid., 496.

⁷⁴ In the following section I will further argue that contextualism can account for our intuitions in the skeptical argument by resorting to an alternate that does not attribute an unreasonable error to ordinary speakers.

the pragmatic content of knowledge assertions. His characterization of what we tend to grasp, or what we identify with ‘what is said’ in an utterance is intended to make sense of contextualist data without attributing an implausible error to ordinary speakers.

According to his view, one utters

(1) I don’t know that *p*

only to indicate that she cannot rule out the possibility that was raised in the course of conversation. In the second bank case, when the speaker’s wife mentions the possibility that the bank might have changed its hours, the speaker answers ‘(I guess) I don’t know that the bank will be open tomorrow’ to indicate that his epistemic position is not good enough to rule out the counter-possibility that his wife has raised. His response addressed the issue that was salient in his wife’s question. However, if this captures the man’s communicative intentions, then it seems reasonable to suppose that the speaker would admit that he does in fact know the proposition, given an opportunity to do so. Why? If his utterance was only intended to speak to the most salient issue in the conversation, then it seems that he would be willing to recognize his *actual* position with respect to the proposition in case he were called to do so. In other words, while I concede Rysiew’s point that what we understand in certain utterances often reaches beyond the strictly semantic meanings of the literal elements of the utterance, it seems that these are cases in which speakers are able to recognize, upon minimal prompting, that they have read more into the sentence than in fact was there. Moreover, if need be, they have the capacity to reconstruct a sentence that *literally* expresses their communicative intentions, given ordinary language competence and fairly normal subject matter. Consider, for example, a teacher who asks before beginning his lecture, ‘Is everyone here?’ Such

utterances, according to Kent Bach, strictly and literally express an unrestricted proposition, but are typically used to convey something more specific.⁷⁵ The specification of the content “is not uttered and does not correspond to anything in the syntactic structure of the uttered sentence. So it is not explicit....Such utterances are understood by way of a pragmatic process of expansion”.⁷⁶ In this case, the students will understand the teacher to be asking whether all of the students registered for his class are present. His students would accordingly respond affirmatively (given that all the students are present), even though they have asserted a literal falsehood. They are perfectly aware that *everyone*, in the sense of every single man and woman, is not in the classroom. Nevertheless, they respond in this manner in order to address what they take to be the salient question, namely whether those who are registered for the class are present. This seems to be a conversationally appropriate reply. Now suppose that the teacher, instead of proceeding with his lecture, draws attention to the purely literal content of his assertion by pointing out that certain individuals who are not enrolled in the class, such as the president of the United States, his wife, etc, are not in fact ‘here’. It seems clear that his students would immediately acknowledge that they had responded to what they assumed was the speaker’s intention, rather than the literal content of his utterance. They would retract their response by claiming that it is *not* the case that everyone is here, and clarify that what they meant was that those who are enrolled in the class are present. Moreover, their ability to withdraw their response and revise it in a syntactically accurate manner is to be expected from competent speakers of the language.

⁷⁵ Kent Bach, “The Semantics-Pragmatics Distinction: What it is and Why it Matters” in *The Semantics-Pragmatics Interface from Different Points of View*, ed. Ken Turner (Oxford: Elsevier, 1999).

⁷⁶ *Ibid.*, 70.

When we turn to high-stakes situations that involve knowledge assertions, however, speakers do not seem quite as apt to revise their response upon similar questioning. To show this, let us modify the second bank case so that it reflects what occurs in the scenario above. In the second bank case, Rysiew stipulates that the speaker has in fact eliminated the relevant counter-possibilities, and his epistemic position is thereby good enough to know that *p*. (Recall that relevant counter-possibilities are just those that we, as normal humans, would think likely to occur.) Rysiew takes this to mean that the husband is in a position to rule out the likely alternative that the bank is regularly open only on weekdays by having been to the bank on recent Saturdays.⁷⁷ Now suppose that once the husband has asserted that he does not know that the bank will be open tomorrow, a third party asks, ‘So you haven’t been to the bank on a Saturday recently?’ The third party’s question has the same effect as the teacher’s in the previous example, which is that it indicates the semantic content of the speaker’s utterance (as envisioned by Rysiew) by pointing out one of its implications. Having visited the bank on a recent Saturday would have sufficed to eliminate the likely alternative that the bank is regularly closed on the weekends. On Rysiew’s account, if the speaker does *not* know that the bank will be open on Saturday, this entails that he cannot eliminate the likely alternative that the bank is usually closed on the weekends because he has not gone to the bank on a recent Saturday. But when the third party points this out, the husband will not have the same reaction as the students in the previous example. In the latter, when the teacher pointed out the literal extension of his question, his students acknowledged the fact that the president and his wife were absent, and further recognized this fact as having

⁷⁷ Rysiew, 490.

semantic implications for their response. In the second bank case, however, it seems fairly unlikely that the speaker would shift his knowledge claim in light of the third party's utterance. That is, when the third party indicates the literal extension of his utterance, the speaker acknowledges that he *has* been to the bank on a recent Saturday (and can therefore eliminate *likely* counter-possibilities), but is not thereby inclined to retract his knowledge denial. Revising his previous assertion indeed seems rather counter-intuitive and not at all like the students' revision in the previous case. While he would acknowledge that he has eliminated the relevant possibility alluded to by the third party's assertion, he would *not* withdraw his earlier claim.

In case the reader's intuition concerning the speaker's behavior clashes with my own, I can offer two reasons to think it fairly unlikely that the speaker would retract his knowledge denial. First, I take it that the couple's subsequent decision in the two bank cases is supposed to hinge on his claim. In the first case, it seems that the couple decides to return to the bank the next day, while in the second, the speaker decides to investigate further. If, in fact, the speaker sincerely thought that his epistemic position were good enough to know the proposition, then re-checking on this fact appears to be unnecessary and somewhat irrational by their own lights. Second, the possibility of incurring harsh penalties might reasonably cause him to re-evaluate evidence that he once thought sufficient to know the proposition. Although the speaker considered his epistemic position good enough to know the proposition initially, acknowledging the grave consequences might, at the very least, cause him to reconsider his support. These are reasons to doubt that drawing attention to the literal content of his assertion (as conceived by Rysiew) will compel the speaker to retract his denial, as in the case mentioned above.

What does this show? The fact that speakers in such high-stakes situations are unwilling to retract their knowledge denial when a hearer indicates the allegedly literal extension of his utterance has at least two implications concerning the speaker's communicative intentions. First, in uttering (1), the speaker is in fact not *merely* conveying that he cannot eliminate the possibility that has just been raised, but is also communicating that he does not believe that his epistemic status is good enough to *know* that *p*. That is, he is making a sincere claim about his overall epistemic status with regard to the proposition. This is why he is not inclined to withdraw his knowledge assertion when he is given such an opportunity. Rysiew might concede this point without thereby having to abandon his theory. He might claim that this shows that individuals are regularly mistaken about their actual epistemic position in certain situations. However, the second implication of the speaker's reluctance to withdraw his knowledge claim is that the speaker cannot become aware of the semantic meaning (as Rysiew conceives of it) of his assertion. How is this a further implication? If the speaker was able to understand the literal or semantic content of 'know' (according to Rysiew's account), then he would be inclined to change his mind about his epistemic position, recognizing, in fact, that he *does* know the proposition in question. However, since he does not do so, it appears that the speaker is not aware, and cannot become aware of the semantic content 'knows'. This indeed is a strange consequence. It appears rather peculiar that competent, rational speakers of the language cannot become aware of the meaning of 'knows' in certain contexts. Rysiew, then, must also adopt an implausible error theory in order to explain our linguistic behavior in high-stakes situations. Here is a more precise

formulation of my argument: (Assume that the speaker has eliminated the relevant counter-possibilities and is in a high-stakes situation)

1. In uttering ‘I don’t know that p ’ the speaker intends to convey (i) he cannot rule out the salient counter-possibilities to p or (ii) he cannot rule out the salient counter-possibilities to p and his epistemic position is thereby not sufficient to know that p .
2. If the speaker intends to convey (i), then he would admit to knowing p upon minimal prompting.
3. The speaker will not admit to knowing p upon minimal prompting.
4. Therefore, the speaker intends to convey (ii). (From 1, 2 & 3)
5. According to Rysiew, ‘I don’t know that p ’ literally expresses that one has not eliminated the *relevant* counter-possibilities to p , and as such, his epistemic position is not good enough to *know* that p .
6. If the speaker were aware or could easily become aware of the semantic content (as Rysiew defines it) of ‘I don’t know that p ’, the speaker would admit to knowing that p upon minimal prompting.
7. The speaker is not aware and cannot easily become aware of the semantic content of ‘I do not know that p ’ (as Rysiew defines it). (From 3, 5 & 6).
8. But it is implausible that competent English speakers are not aware or cannot become aware of the semantic content of ‘I don’t know that p ’.
9. Therefore, Rysiew’s analysis is committed to an implausible error theory. (From 7 & 8)

It appears, then, that Rysiew does not succeed in offering an account of contextualist data without having to adopt an implausible error theory.

Plausible Responses

Rysiew will want to deny my conclusion that speakers cannot become aware of the semantic content of an ordinary expression like ‘knows’. In response to a slightly different criticism, Rysiew maintains that claiming to know that p in the presence of uneliminated $\sim p$ possibilities requires an uncomfortable cancellation. ‘I know that p ’ conveys that I have eliminated the salient possibilities. If there are salient alternatives that I have not eliminated, then I would have to claim the following, “I know that p , but I

have not ruled out all the salient counter-possibilities to p "⁷⁸, thereby imparting that I have both ruled out the alternatives and that I have not ruled them out. The second clause "raises the then-mentioned possibility to salience. So while the 'I know' communicates that the speaker is confident that p and thus that he (believes he) can rule out the salient not- p possibilities, the concession clause communicates that there's a salient not- p possibility which the speaker *can't* rule out".⁷⁹ The speaker sounds as if he is contradicting himself, at least with regard to what is understood by the two clauses.

This response (that I have attributed to Rysiew, but which involves crucial elements of his reply to other criticisms) concedes that speakers in high-standards contexts might refrain from changing their knowledge assertion, but rejects the conditional in the second premise of my argument above. This premise states that if the speaker intended to convey only that he cannot eliminate the salient possibilities raised in the conversation by uttering 'I don't know that p ', (rather than an assessment of his overall epistemic status with respect to the proposition), he would admit to knowing p upon minimal prompting. That is, were his audience to question him further or point out the literal extension of his utterance, the speaker would revise his statement. He would clarify that even though he can't eliminate the salient possibilities, he nevertheless knows the proposition. I then claimed that the best explanation for his reluctance to do so is that he takes his utterance to reflect his evaluation of his epistemic status with respect to the proposition, such that he believes that it is insufficient to know that p . Rysiew's response blocks my conclusion by denying the consequent. The idea is that even if it is the case the speaker intends to convey that he can't rule out the salient possibilities (*without*

⁷⁸ Another way of saying this would be 'I know that p , but it's possible that q ', where q is a salient $\sim p$ possibility. Rysiew maintains that this would be equally uncomfortable since it sounds inconsistent.

⁷⁹ Rysiew, "Context-Sensitivity", 493.

further intending to convey anything about his epistemic position), he would *not* be inclined to admit to knowing that p upon minimal prompting. Asserting that he knows that p would require the speaker to convey seemingly inconsistent information, namely that he both has and has not ruled out the salient possibilities in question. Thus, the speaker's unwillingness to retract his knowledge denial neither reflects an assessment of his epistemic position nor does it show that the ordinary speaker is unaware of the semantic content of the expression.⁸⁰

According to this response, the discomfort associated with canceling an implicature causes speakers to refrain from revising their knowledge claim. But if this phenomenon explains our reluctance in the case of knowledge assertions, it would seem that any expression whose pragmatic content is widely associated with it will have the same result. That is, speakers would be unwilling to cancel the implicature in other cases since doing so would convey something that sounds inconsistent or contradictory. Such discomfort, however, does *not* appear to produce the same result in other cases. Consider the following example. Suppose I say, 'John and Lisa are married'. Although I have not literally stated that the two are married to each other, such an expression typically imparts this information. But John and Lisa are not married to one another. Suppose after I have uttered this statement, someone asks, 'What was their wedding like?' Her question, is of course, conversationally appropriate since I seemed to convey that they were married to each other. But, since it is clear that she has misunderstood what I intended to say, I would clarify the matter by stating, 'John and Lisa are not married to each other'. I might even say, 'John and Lisa are married, but not to each other'. According to Rysiew, since

⁸⁰ While this response is similar to the one Rysiew offers, it was originally intended as a response to a slightly different objection. I have modified it so that it pertains to my objection.

such an utterance sounds contradictory, (because I appear to convey both that John and Lisa are and are not married to each other), I will refrain from canceling the implicature. Admittedly, my audience may be surprised, and perhaps even accuse me of misleading them. But as I think is clear, I would in fact cancel the implicature in the interest of telling the truth and clarifying a mistaken understanding. This example shows that first, when I perceive that my audience has misunderstood my intent, I will find it necessary to cancel the implication of what I have said. Second, I will do so even if this requires an 'uncomfortable' cancellation for the sake of clarifying my communicative intentions. Communicating successfully takes precedence over uncomfortable cancellations.

Let us now turn to the knowledge assertion in the second bank case. According to the example above, the speaker would clarify his communicative intentions at the expense of uttering something seemingly inconsistent. In fact, in spite of such discomfort, it seems that speakers are prepared to do just this in the interest of clarifying his communicative intentions. The discomfort associated with canceling the pragmatic implicature of a knowledge assertion, then, would not block the speaker from revising his knowledge claim. This suggests that the speaker's reluctance to revise his knowledge claim is not due to linguistic anxiety. Rather, the speaker is unwilling to change his knowledge assertion because he sincerely meant that his epistemic position is insufficient to know that *p*. Moreover, given his reluctance, it appears that the speaker is not aware, and cannot become aware of the semantic content of his utterance in Rysiew's account. Hence, Rysiew does not succeed in offering a theory that accommodates contextualist data without invoking an implausible error theory.

IV. Alternative Error Theory

Are we at an impasse? It appears that neither Rysiew nor the contextualist can account for our knowledge ascriptions without adopting an implausible error theory. Rysiew assumes that the only way for the contextualist to explain why the skeptical argument appears puzzling is to claim that the speaker has misunderstood the conclusion. That is, we find it troubling because we believe that the skeptic is using 'knows' according to our ordinary standards. But, the contextualist maintains that the content of a knowledge sentence is determined by the speaker's context (including her interests, intentions, etc), which the speaker then uses to accurately determine whether or not to ascribe knowledge to the subject. Confusion, then about the literal content of her ascriptions amounts to ignorance about her own interests, communicative intentions, and so on. I concede Rysiew's point that the contextualist cannot hope to explain the difficulty associated with the skeptical argument by appealing to a misunderstanding of which propositions are expressed therein. Indeed, the theorist must avoid attributing such an error to the ordinary person. I will end by offering an alternative route for the contextualist.

The Skeptical Context

I think then, that this is the wrong explanation of what occurs in the skeptical context. Instead, in uttering the premises and conclusion of the skeptical argument, the speaker is in fact aware of the context and understands the propositions that are expressed. That is, she accepts that each of the knowledge assertions are made relative to extremely high standards. The skeptical argument nevertheless appears puzzling because she overestimates her epistemic position with respect to the embedded proposition in the

conclusion. This causes her to be mistaken about the truth *value* of the conclusion, while still grasping the content of the knowledge ascription. How can this be? I want to suggest that the skeptical context presents a unique situation with extremely, if not untenably high standards for knowledge, such that even an extraordinarily strong epistemic position with regard to a given proposition *p* is rendered insufficient to ascribe knowledge that *p*. At the same time, however, we not only have some evidence, but evidence that has gathered throughout our lives—and upon which we unhesitatingly act—for the proposition that we have hands. This might cause us to be resistant to the claim that our epistemic position with regard to having hands is not strong enough for the high standards invoked by the first premise. We recognize, in other words, the content of our knowledge assertion, which is what allows us to deny knowledge in the first premise. It is not that we don't understand the *meaning* or content of 'knows', but rather, we are misguided about the truth value of the statement in the conclusion.

We might better understand what occurs in the skeptical argument by appealing to DeRose's discussion of attributor and subject factors.⁸¹ According to DeRose, the semantics of knowledge sentences depend on what he calls "attributor" factors, which fix their truth conditions, and "subject" factors, which determine the truth value of such assertions. The former determine how good an epistemic position the subject must be in with regard to *p* to count as knowing that *p*, while subject factors reflect whether or not the truth conditions are satisfied.⁸² Thus, 'S knows that *p*' asserted in context *C* is true only if *S* meets the epistemic standards determined by *C*. The ordinary speaker, then, might have difficulty accepting the truth of the conclusion in the skeptical argument, not

⁸¹Keith DeRose, "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52, No.4 (December 1992): 921-923.

⁸² *Ibid.*, 922.

because she has misunderstood the attributor factors, or the proposition that is being expressed, but because she takes her evidence, the features in her context qua subject, to support the fact that she has hands. It is an overestimation of the strength of her subject features, then, that leads her to deny the skeptical conclusion, that we don't know that we have hands. This is why the skeptical argument *appears* to be troubling and paradoxical. We are inclined to agree with the first premise that we don't know that we are BIVs, given the high standards, however, even in such high standards, we are not inclined to accept the conclusion.

What explains the fact that we do not make the same mistake with respect to the first premise, as we might in the conclusion? We are resistant to the conclusion not merely because of the preponderance of evidence for the proposition that we have hands, but due to the fact that we tend to *overestimate* our evidence. We mistake the epistemic strength of our position with regard to this proposition. While we have just as much evidence for not being BIVs, we are not inclined to overestimate our evidence with regard to the proposition that we are not BIVs. Why not? This is because of the very fact that having hands is an *ordinary* proposition. We reason and behave in such a way that assumes this. For example, suppose I am trying to determine how I can carry all the grocery bags to my apartment in a single trip, I might reason that I can place some bags in certain positions or at certain locations on my arms. Or, take the case of a rock climber who is planning his next move. He reasons that he can extend his hand to a certain cleft in the rock. A mother who is attempting to carry her infant while still holding the hands of her other children will have to strategize how she might use her arms. Given that we

use this fact in ordinary reasoning, we seem to consistently assume that we have enough evidence for this claim.

In contrast, it does not appear that not being a BIV factors into ordinary tasks. It is certainly the case that I would not be rock-climbing or carrying grocery bags if I were a handless BIV, however, the point is that this fact is not typically entertained or considered in our practical or theoretical behavior. Even though we might have the same degree of evidence with regard to both propositions (that we have hands, and that we are not BIVs), we are much less likely to overestimate our epistemic strength with respect to the skeptical hypothesis, since it is clear that we are less likely to reason or behave using the belief that we are not BIVs. Moreover, even if the skeptical hypothesis were disguised in more ordinary scenarios, such as the possibility that we are dreaming, we would be less inclined to overestimate our position with respect to such a possibility. This is because once again that all our perceptions and experiences are not illusory seems to be fairly remote from our ordinary behavior and reasoning. We do not typically reason by employing the fact that we are not dreaming. While the proposition that we have hands is paradigmatically ordinary, in the sense that we reason according to this fact, the skeptical hypothesis is not. This is what causes us to mistake our epistemic position in the first and not the latter, and makes the argument appear paradoxical. That is, while the conclusion states that we don't know that we have hands, it appears to us that our epistemic position with respect to this proposition *is* sufficient to grant knowledge.

This alternative avoids Rysiew's objection as it is not committed to the claim that speakers are confused about the content of their assertions, and hence, their communicative intentions. We recognize the high standards of the skeptical context and

thus are able to know what is *meant* by each of the premises and the conclusion. We automatically grasp this when the possibility of being a BIV is introduced. However, accepting the conclusion that we don't know that we have hands sheds doubt on our overwhelming evidence that we have hands. Moreover, because the speaker is able to grasp the literal content of her assertion, the response skirts the worry I raised for Rysiew's account. It is clear that competent speakers have access to the semantic meanings of their utterances. While such an explanation appears to circumvent these problems, there are probably further worries. In the next section, I will address one such problem, and hope to show that it is only apparent, and neither counts against the explanation that I have offered with regard to the skeptical argument, nor contextualism generally.

A Plausible Worry

Even though the contextualist is not saddled with the claim that ordinary speakers are completely unaware of the content of their ascriptions, it appears that both the contextualist and the invariantist must say that the speaker is mistaken with regard to certain judgments about knowledge ascriptions. I have suggested that ordinary reactions to the skeptical argument are mistaken because we overestimate the strength of our epistemic position with regard to having hands, which ultimately leads us to deny the conclusion that we don't know that we have hands in high standards. This is problematic for the contextualist because it appears that she takes some intuitions at face value while dismissing others as misguided. But if contextualists cannot explain why some judgments, but not others, are taken at face value, their analysis seems rather arbitrary.

As previously discussed, however, it is not the case that I have *arbitrarily* determined which intuitions to take at face value. It seems clear that we have more reason to accept our intuitive judgments about some matters of fact over others. In the conclusion of the skeptical argument, it is clear that first, we have a lot of evidence that seems to show that we have hands. Moreover, having hands is an *ordinary* proposition, in the sense that it makes reference to something that is familiar. We are inclined to use this fact in decisions that we make. These features together mistakenly lead us to believe that we *know* we have hands relative to high standards. Presumably not *all* knowledge ascriptions will meet such criteria. Our inclinations or judgments in the skeptical argument are unique. Thus, even though our intuition in this case might not be taken at face value, this does not show that I have arbitrarily selected the intuitions to accept and those to deny. Rather, it is the particular and unique nature of the skeptical context that renders our judgment inaccurate. In contrast, practical ordinary judgments in our daily lives are taken to correspond to the actual truth conditions or semantic content of knowledge ascriptions.

V. Conclusion

I have argued that Rysiew does not successfully show that his invariantist semantics is preferable to contextualism. In making knowledge assertions about themselves, speakers seem not only to impart information about their relation to salient alternatives in the conversation, but also their belief about their epistemic status. This is because speakers are unwilling to change their knowledge assertion once the literal extension of their utterance (as Rysiew understands it) is pointed out. Their reluctance to

shift their assertion also reflects the fact that speakers, on Rysiew's account, appear to be unaware and unable to become aware of the semantic content of their assertions in certain contexts. Rysiew's view, then, is committed to an implausible error theory according to which competent speakers of a language cannot become aware of the semantic content of their utterances. I further suggest an alternative account that explains our reactions in the skeptical argument, which avoids attributing an implausible degree of error to ordinary speakers.

Chapter Three

Skeptical Invariantism and Epistemic Distinctions

As discussed in the previous chapter, epistemic contextualists argue that the truth conditions of knowledge-ascribing and knowledge-denying statements vary according to the speaker's context. Features that affect their truth conditions include doubts or challenges raised in the course of conversation, as well as the interests and concerns of conversational participants. This, in turn, allows speakers to use 'knows' in a way that reflects their immediate concerns, whether their situations call for demanding or fairly low epistemic standards. Suppose that I believe that there is milk in the fridge. I remember putting it there last night and have evidence to suggest that no one else drank milk after me. Suppose further that it is true that there is milk in the fridge. In such a situation, it seems clear that I know that there is milk in the fridge. Admittedly, I have not eliminated the possibility that my husband or I might have sleep-walked at night and unconsciously consumed all the milk or that there was a robbery during which only the milk was stolen. However, these counter-possibilities appear excessive to my situation. For my purposes, my memory seems sufficient for knowledge, thereby distinguishing my epistemic status from someone who lacks such evidence. Contextualism straightforwardly makes such fine epistemic distinctions by appropriately lowering the standards of 'knowledge' to correspond to our immediate concerns.

This inherent quality in contextualism highlights a deficiency or weakness in skeptical invariantism (SI). SI cannot acknowledge subtle differences in epistemic states due to two essential and intrinsic features. First, the account posits exceptionally high standards for knowledge that exceed the evidentiary levels that we customarily seek and

remain fixed regardless of our purposes or interests. Proponents of the view maintain that to know that p , one must be able to eliminate *all* possibilities in which $\sim p$. Second, because subjects will fail to meet these standards, it turns out that most if not all knowledge ascriptions are false. Both aspects make it impossible, at least *prima facie*, to use ‘knows’ in a way that recognizes finer distinctions in epistemic conditions.

According to SI, my epistemic status with regard to the whereabouts of the milk would be similar to that of someone who lacked my memorial evidence. As neither of us could eliminate all the $\sim p$ possibilities—such as a robbery during which the milk was stolen or that an evil demon has suddenly destroyed the milk—skeptical invariantists would declare that neither of us knows that the milk is in the fridge. Mark Heller (1999) notes that

it is not that the skepticism itself is an unacceptable conclusion. Rather it is that such standards fail to draw the distinctions that are important to us.... If I hope that my wife knows where she left the car, I do not care about her ability to rule out worlds in which an evil genius is feeding her false memories.... What I care about is her ability to rule out worlds in which her memory works as it actually does...but she leaves her car in a different location. If she can rule out such worlds, then she is likely to find the car.⁸³

Such an asymmetry in the accounts would presumably recommend contextualism over its invariantist rival.

Skeptical invariantists, however, have supplemented their semantic theses with an account of ordinary usage that accommodates the finer epistemic distinctions highlighted by contextualists. They maintain that while most knowledge ascriptions are strictly false, their pragmatic content, or what they convey, is true. This pragmatic strategy, however, proves inadequate since it seems that we can indiscriminately classify any apparently context-sensitive term as pragmatically, rather than semantically variable. We might, for

⁸³ Mark Heller, “The Proper Role for Contextualism in an Anti-Luck Epistemology,” *Philosophical Perspectives* 13 (1999): 5, 8.

example, apply this approach to words such as ‘tall’ and ‘nearby’, which are almost universally recognized as contextually dependent. Skeptical invariantists, then, must provide criteria or tests by which to distinguish ‘knows’ from terms that are contextually-sensitive. Peter Unger has offered what he calls the paraphrase test, which—when combined with emphasis—is designed to isolate ‘knows’, a pragmatically variable term, from semantically variable terms. In what follows, I will argue that Unger’s distinction is not sufficiently supported. Without an accurate and reliable method (or criterion) to distinguish semantic context sensitivity from apparent (i.e., pragmatic) context sensitivity, there is no reason to believe that ‘knows’ and other terms such as ‘flat’ fall into the second, rather than the first, category. This conclusion further exemplifies my overarching thesis that invariantist pragmatic explanations of our ordinary use of ‘knows’ cannot be sustained. This chapter, then, will have the following structure: After a more thorough exposition and analysis of the pragmatics method, I will expound Unger’s paraphrastic test. I will argue, however, that his proposal is unintuitive, and therefore, lacks sufficient support. Furthermore, attempts to generate the right type of intuitions by strategically emphasizing certain words fails as this proves to be an unreliable linguistic device. Hence, I conclude that skeptical invariantism cannot acknowledge the finer distinctions in epistemic status, which are easily accommodated by contextualists.

I. Pragmatics and WAMs

As stated above, skeptical invariantists sustain their elevated truth conditions for knowledge by relegating ordinary uses of ‘knows’ to matter that is imparted or conveyed to the audience, i.e. its pragmatic content. This content then delineates criteria for the warranted assertability, rather than the semantic truth, of knowledge ascriptions. Keith

DeRose has termed this invariantist approach the ‘warranted assertability maneuver’ (WAM), and maintains that it is often employed to handle intuitions that appear to support contextualism: “an intuition that an assertion is true can be explained away by means of the claim that the assertion, while false, is warranted”. Alternatively, WAMs can also be used to explain “why an assertion can seem false in certain circumstances in which it’s in fact true”.⁸⁴ That is, skeptical invariantists allow that the claim ‘I know that p ’ is felicitous or warranted because it implies that I have sufficient evidence for the task at hand. However, the statement is nevertheless false as I have not eliminated *all* the $\sim p$ possibilities. Peter Unger has further asserted that a divergence between assertability and truth conditions, or pragmatic and semantic content, is not peculiar to ‘knows’, but is fairly prevalent within the language. He illustrates this phenomenon by comparing the use of ‘vacuum’ in a scientific laboratory versus a peanut factory:

Someone in a peanut factory may declare that vacuums have been established in certain cans.... His intent is to convey the idea that, for those purposes, there is no important difference between the insides of those cans and a space which really is a perfect vacuum. The meaning of ‘vacuum’ remains the same whether on the canner’s lips or in a laboratory report....The regions in the laboratory are much closer to being real vacuums than those in the peanut cans. But that difference, while important to the research is not important to keeping the peanuts fit for eating in several months.⁸⁵

Although the canner’s interests warrant his claim, the insides of the cans do not strictly qualify as vacuums, rendering his assertion literally false. Unger maintains that the language is rife with such terms—words that are often misused because they are typically (if not always) applied to objects or persons that fail to satisfy their stringent truth conditions. He classifies such terms, which also include ‘flat’, ‘certain’, etc., as absolute

⁸⁴ Keith DeRose, “Contextualism: An Explanation and Defense,” in *The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa (Oxford: Blackwell Publishers, 1998), 13.

⁸⁵ Peter Unger, *Ignorance* (Oxford: Oxford University Press, 2002), 51-52.

to indicate an absolute limit.⁸⁶ ‘Knowing’ also falls in this category as it requires (absolute) certainty. This makes it the case that at most, hardly anyone knows anything.

By appealing to the pragmatic content of ‘knows’, invariantists appear to accommodate ordinary uses of the term to make fine epistemic distinctions while preserving high and invariant standards for knowledge. But Unger’s approach does not ultimately suffice. Why not? It leaves open the possibility that *any* term whose everyday use appears context-sensitive is merely pragmatically variant, but semantically stable. For example, observations that support Unger’s pragmatic explanation of ‘knows’ can be cited regarding our use of ‘tall’. In ordinary contexts, we employ ‘tall’ to communicate information that is, in large part, contextually determined. Asserting the statement ‘Lisa is very tall’ in the midst of other kindergarteners seems to mean something different than when the same sentence is uttered in the midst of basketball players. As with ‘knows’, we can assign such use to pragmatics and brand them false, but warranted. That is, we might classify ‘tall’ as an absolute term such that it is correctly applied to a person just in case no one else is taller than she is. It would then follow that at most, hardly anyone is tall. This raises problems for Unger and other invariantists, since tall is widely accepted as a context-sensitive term whose content varies according to features such as a salient comparison class. Hence, as it stands, this pragmatic strategy is inadequate without some further clarification about why terms such as ‘know’ require clauses about assertability while others do not.⁸⁷ Here is a more formal presentation of the worry:

⁸⁶ I will present a more precise definition of absolute terms in the next section. Here, the important point is that Unger identifies certain terms as denoting absolute limits, which may nevertheless be used in less strict ways given their pragmatic content.

⁸⁷ A similar charge could be raised against the contextualist: We often use terms to convey more than their literal content. In some of these cases, there is broad consensus that the semantic content of such terms remains constant, despite such variation in usage. Contextualists contend that our variable usage of

1. The apparent context-sensitivity of word usage is relegated to shifts in pragmatic, rather than semantic, content. (Unger's theory)
2. The use of 'tall' is apparently context-sensitive. (Observation)
3. 'Tall' shifts in its pragmatic content, but has a stable semantic content. (From 1 & 2)
4. 'Tall' has an unstable semantic content. (Widely accepted, *reductio ad absurdum*)

The argument resembles a *reductio ad absurdum*, given the contradiction in (4) and the second conjunct of (3).⁸⁸ Invariantists can avoid this result only by qualifying (1). In other words, they must provide tests or criteria that differentiate 'knows' and other absolutes from semantically context-sensitive terms such as 'tall'. The remainder of the chapter will present and evaluate Unger's efforts to distinguish apparent from actual context-sensitive terms.

II. Absolute Terms

Unger begins by comparing relative terms to basic absolutes, maintaining that while the former admit of degrees of applicability, the latter "denote an absolute *limit*". He defines absolute terms in relation to the absence of corresponding relative features, explaining that absolute limits are "approached to the extent that the relevant relative property or properties are absent" in the object.⁸⁹ For example, one is (absolutely) certain insofar as no degree of doubt remains in his mind; an object is flat just in case it is

'knows' corresponds to the term's semantic content. However, they have not provided criteria that distinguish 'knows' from other terms, whose variable use does not signify a context-sensitive content. While I do not address this question directly in this chapter, note that the direction of the first three chapters in this dissertation is aimed towards rejecting context-independent analyses of 'knows'. That is, attributing semantic context-sensitivity, rather than pragmatic variance to 'knows' seems to be the best approach to account for its systematic variability in ordinary contexts.

⁸⁸ I limit myself to a resemblance because it might be argued that the argument does not result in a strict contradiction. While it is *almost* universally accepted that comparative adjectives like 'tall' are contextually sensitive, Herman Cappelen and Ernie Lepore, *Insensitive Semantics* (Oxford: Blackwell Publishing, 2005). have argued that 'tall' does not fall into the very exclusive list of contextuials that includes only indexicals and common adjectives such as 'local' and 'domestic'.

⁸⁹ Unger, 55.

not all bumpy or curved. While the strict semantic content of absolute expressions remains fixed across contexts, they may be *used* with different communicative intentions to suit the purpose at hand. Unger classifies ‘knows’ as a non-basic, absolute term to indicate that its meaning entails an intermediate, basic absolute term: certainty. While most absolutes correspond to the absence of some opposite feature, such as ‘flat’ and ‘bumpy’, Unger maintains that it is not an exhaustive test as some absolutes do not. For example, the terms ‘full’ and ‘empty’ do not readily admit of a relative feature to help define them.⁹⁰

What we need then is a more effective means by which to differentiate absolutes from other terms in the language. To this end, Unger introduces the paraphrase test, according to which terms that admit of Unger’s paraphrase are classified as absolutes. He presents two types of paraphrases: one which interprets the effects of certain modifiers on absolute terms and a second that applies to comparative locutions that contain absolute terms. Unger maintains that in the case of absolutes, modifiers such as ‘very’ and ‘extremely’ serve to indicate how *closely* the object *approximates* the limit that is denoted by the absolute term, rather than the degree to which the object exhibits this property. For example, ‘That table is very flat’ expresses the proposition that the table is *nearly* flat or very close to absolute flatness. Such a paraphrase, however, does not apply to relatives. Consider, for example, the statement ‘The surface is very bumpy’. Here, the statement does not merely indicate how close the object is to being bumpy. Rather, it expresses the proposition that the table is bumpy to a large degree. Note two important negative corollaries of the paraphrase: first, a statement that contains a

⁹⁰ Unger, pp 55-56.

modified absolute such as ‘The table is very flat’ does not indicate the degree to which the object or person reflects the property denoted by the term. Second, the statement actually depicts the object as not having it at all. That is, a table described as ‘very flat’ does not satisfy the conditions of (absolute) flatness, but is instead only asymptotically flat.

Unger extends his paraphrase to comparative locutions. Consider the following statement:

‘*a* is not as flat as *b*’

We might take this sentence to mean that *a* is not as nearly flat as *b*. This suggests, however, that both objects are only *nearly* flat, and hence, not flat. This is an incorrect translation as the original statement leaves open the possibility that *b* is flat. The apt translation is this: “The first surface is either not flat though the second is or else it is not as nearly flat as the second”.⁹¹ While the paraphrase applies to absolute terms such as ‘certain’, it does not capture what is said in similar comparisons using relative terms.⁹²

What kind of support does Unger provide for his paraphrases? In his discussion, Unger appeals to semantic intuitions as his primary source of support. When differentiating between the effects of modifiers on absolutes and relatives, Unger writes, “intuitively, we may notice a difference here. For only with our talk with ‘flat’ do we have the idea that these locutions are only convenient means for saying how closely a surface approximates... a surface which is (absolutely) flat”.⁹³ He calls his translation of statements containing absolute terms “intuitively plausible”, while he finds the

⁹¹ Unger, 59.

⁹² Unger admits that the method may not consistently identify relative terms and further analysis may be required. The test does more to establish terms as absolute rather than relative. He discusses this on pp 60-61.

⁹³ Unger, 57.

paraphrase “*quite obviously*, a terribly poor interpretation of what we are saying” in the case of relatives.⁹⁴ Further, he maintains that “on *reflection* there is a difference between confidence and certainty”, the former being relative and the latter absolute.⁹⁵ Unger then, relies wholly on our intuitive judgments as he offers no further evidence to validate his paraphrase.

III. Recalcitrant Intuitions and Emphasis

The paraphrastic procedure appears to qualify (1) in the argument above. That is, terms that admit of the modifier and comparative paraphrases are absolute and subsequently require a pragmatic account of ordinary use. We might then note that ‘tall’ fails the paraphrase test. To describe someone as very tall is not to say that she is *nearly* tall (such that no one is taller than her) but rather to indicate that she is tall to a high degree. The paraphrastic procedure preserves the widely accepted thesis that ‘tall’ is a context-dependent term and in turn appears to salvage Unger’s WAM from a contradiction.

It is not clear, however, that Unger is completely out of the woods. As indicated earlier, Unger grounds his method primarily on semantic intuitions. But it seems quite plausible for statements containing absolute terms to deviate from Unger’s interpretation and relatively simple to imagine cases in which this is true. Consider the following conversation:

Shelley: We need a surface that is very flat for our art project.
Mike: That table is flat enough to hold our project steady and level.

⁹⁴ Ibid. (emphasis mine).

⁹⁵ Ibid.,65.

In response to Shelley's request, Mike points to a table that he considers suitably flat, one that can provide a steady and level surface. Why is this exchange problematic for Unger's account? Unger's paraphrase would read Shelley's statement as 'We need a surface that is nearly (absolutely) flat'. Further, according to the two corollaries, Shelley's request would neither specify a degree of flatness, nor ask for a table that is flat at all. Note, however, that Mike's response does not reflect this understanding. Instead, Mike offers a table that is *sufficiently* flat without referring to absolute flatness at all. In other words, he interprets Shelley's request as asking for a table that exhibits an appropriate degree of flatness. Such a response appears reasonable and suggests that our intuitions do not support Unger's paraphrase. As Unger relies wholly on intuition to validate his distinction of absolute terms, he lacks an adequate defense of his approach.

Unger, to some extent, acknowledges that a term's absolute character might elude us initially:

If one approaches the key sentences I've exhibited in *one* way, everything I say seems to make good sense. But it is very easy and natural to approach these sentences in *quite another* way. In this other way, at least very much of what I have said seems and sounds just plain wrong.⁹⁶

The 'other' way involves "taking in sentences quickly, with no special attention given to them or any of their constituents".⁹⁷ According to Unger, the problem can be remedied by emphasizing the terms of interest, which in turn alerts the audience to attend to certain parts of the sentence and grasp their meaning. For example, if someone brought a jug when you asked for a mug, you might repeat the request by saying 'Bring me the *mug!*'⁹⁸ Or if they brought the mug to someone else, you might say, 'Bring *me* the mug!' Both

⁹⁶ Unger, *Ignorance* 74.

⁹⁷ Unger, 75.

⁹⁸ *Ibid.*, 75.

utterances express the same proposition and neither alters the meaning of the original, unemphasized sentence: ‘Bring me the mug!’ Rather they focus the listener’s attention on a certain part of the statement.

How does this help Unger’s theory? First, he maintains that applying emphasis to absolute terms enables speakers to concentrate on their semantic content without changing their meaning. Second, emphasizing these terms in the context of modifiers or comparative locutions generates intuitions that are consistent with his paraphrase. Hence, Unger has the support he needs to defend his paraphrastic test, thereby allowing him to isolate absolute terms and avoid the *reductio* discussed earlier. He offers the following example to illustrate:

‘This is flat; that other is flatter’

Such a comparison is inconsistent according to Unger’s account. His paraphrase interprets comparative utterances in the following manner: either one object must be flat while the other is not, or neither object is flat. Hence, if an object *x* is truly flat, then there cannot be some object *y* that is flatter. Either *x* is not flat or *y* is not. However, the statement’s inconsistency seems less than obvious. Emphasis proves useful in this case by prompting us to focus more carefully on the terms of interest:

‘That is *flat*; that other is *flatter*’

The comparison now appears awkward. Emphasis, then, seems to generate judgments that support Unger’s paraphrase.⁹⁹ While we do not ordinarily demand our interlocutor to avoid such statements or make a similar demand on ourselves, we immediately recognize an inconsistency when we concentrate on the main components of the

⁹⁹ Moreover, Unger notes that there is no comparable infelicity associated with ‘That is *bumpy*; that other is *bumpier*’.

sentence. The paraphrastic procedure combined with emphasis then seems to effectively identify absolute terms.

IV. Problems with Emphasis

Unger, then, appears to have resolved the original worry against an account that posited fixed and unattainable standards for knowledge. He (i)acknowledges fine distinctions in epistemic status by appealing to the variable pragmatic content of ‘knows’ and (ii)provides a test (paraphrase plus emphasis) that identifies absolute terms, thereby differentiating terms which have variable pragmatic contents from terms that do not.

There are, however, worries that emerge with Unger’s reliance on emphasis to secure the necessary intuitions. Emphasis is conducive to his account only if it consistently allows speakers to attend to the meaning of the relevant term(s) without altering their content.

If, instead, placing stress on a word changed its meaning, then our judgments about these terms and the statements which contain them would correspond only to the content of the emphasized terms and the sentences in which they appear. Unger then could not appeal to our intuitions about emphasized terms to support his paraphrase of unemphasized words and statements, as the two would have dissimilar semantic contents. Hence, emphasizing a term must preserve the semantic content of the unemphasized term.

Accordingly, Unger insists that emphasis works in just this manner:

Shouldn’t we have *some* device(s) to attract attention to a term which work in a way that does not effect the term’s meaning, or the meaning of sentences of which that term is a part? Shouldn’t we have a device for every term, and every sentence, that is, a device whose application is wholly general.... It seems to me that emphasis, especially with the device of stress, does just this desired and expected job.¹⁰⁰

¹⁰⁰ Unger, 76.

There are, however, a number of compelling examples that challenge this assertion. Contra Unger, emphasis appears to affect both the proposition expressed by complex statements and the meaning of individual terms in sentences.

Fred Dretske has discussed several ways in which focal shifts can affect the semantic analysis of certain expressions.¹⁰¹ He designates sentences that contain emphasis ‘contrastive’, explaining that they reveal a dominant contrast or a “featured exclusion of certain possibilities”.¹⁰² For example, “Clyde gave *me* the tickets” indicates that Clyde gave the tickets to me rather than someone else, whereas ‘Clyde gave me *the tickets*’ implies that he gave me the tickets rather than the money or check.¹⁰³ While such shifts in emphasis do not alter the content of this statement, Dretske maintains that they do affect the analysis of larger expressions in which this sentence is embedded. Consider the following two statements:

‘Clyde gave me *the tickets* by mistake’

‘Clyde gave *me* the tickets by mistake’.

Suppose that Clyde was told to give me the money, not the tickets. Both sentences contain the same terms, yet it seems that we would assent to the first and deny the second. Why? In the first statement, we recognize that ‘by mistake’ modifies ‘the tickets’ and indicates that Clyde was supposed to give me something else, whereas ‘by mistake’ modifies ‘me’ in the second sentence, indicating that Clyde’s error consisted in giving them to me, rather than someone else. Emphasis affects the proposition expressed by the sentence by specifying Clyde’s mistake. Such a device, then, does not appear as semantically innocent as Unger promises.

¹⁰¹ Fred Dretske, “Contrastive Statements,” *Philosophical Review* 81, no. 4 (Oct. 1972):411-437.

¹⁰² Dretske, 412.

¹⁰³ *Ibid.*, 413.

Dretske further asserts that emphasis restricts the scope of the operator in a complex sentence. Suppose that Clyde, who has grown accustomed to and rather fond of his bachelorhood, discovers that he stands to inherit a substantial sum of money once he gets married.¹⁰⁴ After some searching, he settles on Bertha, who is out of the country for 11 months of the year for business. While it seems true to say ‘The reason that Clyde *married* Bertha was to qualify for the inheritance’, it seems false that ‘The reason that Clyde married *Bertha* was to qualify for the inheritance’. The person that he chose, Bertha, is significant only to the extent that she would least affect his lifestyle, not because she alone would allow him to qualify for the inheritance. Dretske maintains that while the embedded statement (‘Clyde married Bertha’) retains its content in the two inscriptions, the scope of the operator (‘the reason’) shifts according to the site of emphasis. That is, whether Clyde’s inheritance explained the marriage or his choice of bride is determined by which portion of the embedded sentence/operand was stressed. Again, it appears as if emphasis affects semantic content.

In addition to the types of semantic influence that Dretske cites, emphasis can also alter the meanings of individual terms in less complex sentences. Consider the following example: Suppose that at a party, I say to a group of people, ‘Sean and Lisa are friends, but not *friends*’. The intended meaning of the second iteration clearly differs from the first. My audience will, in all likelihood, understand a proposition whose truth conditions differ from the sentence ‘Sean and Lisa are friends’. That is, they will take my statement to mean that Sean and Lisa friends, and nothing more than that. Without understanding the utterance in this way, the statement will be incoherent. Here, emphasizing ‘friend’

¹⁰⁴ Ibid., 417-418.

changed its content. Such cases are not difficult to find. The sentence ‘Lisa is *short*’ expresses a much greater deficiency in height than ‘Lisa is short’. While the latter conveys that Lisa is somewhere below average, the first statement suggests that she is well under its range. Emphasizing terms then appears to shift the content of individual terms as well as complete expressions.

Before considering Unger’s response, let us review how such examples undermine Unger’s method. Recall that in order for emphasis to effectively identify absolute terms, the device must allow speakers to attend to them without changing their content. If, however, emphasizing a part of the statement alters its content, then our judgments about this sentence will not help Unger’s theory, as such intuitions will not pertain to the meaning of the unemphasized terms and their statements. Hence, our acceptance of the paraphrase under conditions of emphasis would not validate Unger’s theory of absolute terms. The examples above suggest that emphasizing a part of the statement does alter its content. Therefore, Unger cannot appeal to emphasis to support his paraphrase.

Unger, however, presents a strong defense of his method. First, he proposes an alternative explanation that accounts for the difference between statements that vary in foci, one that does not appeal to shifts in semantic content. Second, he provides reason to reject the view that emphasis affects semantic content or truth conditions. I will discuss these in turn. In the case of Clyde and Bertha, we encounter these:

- (1) The reason that Clyde *married* Bertha was to qualify for the inheritance.
- (2) The reason Clyde married *Bertha* was to qualify for the inheritance.

Unger asserts that that the reason, the *whole* reason, that Clyde married Bertha was to qualify for the inheritance and sustain his life style. As such, Unger maintains that both (1) and (2) are false in virtue of their omission of part of the reason Clyde wed Bertha. The statements are, nevertheless, dissimilar in a significant respect. In the first version, the speaker highlights their marriage as an important event and then indicates that the reason for this event is that Clyde would qualify for the inheritance. This is not the case with (2). Here, the speaker does not explicitly point to the marriage and hence the reason that follows appears unsuitable. Unger maintains that while “the different emphases...indicate what the speaker or writer represents as the thing to focus on”,¹⁰⁵ (1) and (2) express the same proposition with identical truth conditions. In other words, we mistake the pragmatic content, or what is communicated, for the proposition expressed, and presume (1) true and (2) false. Unger might then say something similar about the *friend* example. Here, the audience mistakenly believes that ‘friend’ corresponds to different content when emphasized because they confuse the speaker’s implied meaning with the literal extension of these statements. While I might have used ‘Sean and Lisa are friends, but not *friends*’ to convey that they are just friends and nothing more, the proposition I literally expressed is incoherent. Hence, our intuitions in each of these cases are easily explained away.

In his second argument, Unger shows that reiterations of a statement, irrespective of emphasis, express the same proposition as the original sentence. Suppose that Dretske says ‘The reason that Clyde *married* Bertha was to qualify for the inheritance’, but I mishear him and think that Clyde and Eartha are married. Dretske corrects me by stating

¹⁰⁵ Unger, 78.

“I said that the reason Clyde married *Bertha* was to qualify for the inheritance”.¹⁰⁶ Unger maintains that Dretske says the very same thing the second time despite the shift in foci. Moreover, Dretske intends to use the new sentence in just such a way. From this, Unger concludes that varying intonation and other forms of stress do not alter the content or truth conditions of statements.

V. Different Forms of Emphasis, Unreliability

Unger, then, appears to successfully stave off worries about emphasis. He first explains away our intuitions about the semantic effects of emphasis by differentiating between communicative intentions and semantic content. Second, he provides examples in which sentences that vary in foci express the same proposition. Unger’s defense, however, does not ultimately save his view. Why not? Let us begin with his second response. Admittedly, Unger’s example is fairly intuitive. We do not suppose that Dretske has said anything different when he repeats himself, despite the fact that he stresses ‘married’ the first time and ‘Bertha’ the second time. Note, however, that this is a singular instance of emphasis, specifically, one in which the audience has misheard part of the utterance. In the example, I misunderstand that Clyde has married Eartha rather than Bertha, causing Dretske to restate the same sentence, all the while sustaining a fixed semantic content. But this particular feature is absent in the problematic cases presented above. More importantly, Unger’s reliance of emphasis in his own theory does not resemble such a case. In Unger’s account, emphasis does not serve to clarify terms that have been misheard or misread. He does not, in other words, emphasize ‘flat’ because

¹⁰⁶ Ibid., 78-79

the audience might think he said ‘cat’ or ‘rat’. Rather, the emphasis is intended to direct disproportionate attention to the specified terms in order to convey that they bear some special significance. For example, the sentence ‘He is *certain* it is a Cadillac, but he’s *more certain* it is an automobile’ causes the audience to attend carefully to the italicized terms not because there is any risk that the reader will otherwise misread it, but because Unger intends that the terms be understood in a particular way. In Dretske’s examples, emphasis functions in the same way. ‘Clyde gave *me* the tickets by mistake’ indicates the precise manner in which he erred: he gave them to me rather than someone else. Hence, Unger’s example is an isolated instance and subsequently does not show that emphasis is semantically innocent in the relevant cases.

Unger’s alternative explanation of our problematic intuitions, according to which we mistake the speaker’s meaning for the semantic content of the statements, also ultimately fails to redeem emphasis. Suppose that Unger is right: We are misguided about the semantic implications of emphasis and that his alternative explanation captures what is going on. The problem is that emphasis nevertheless produces apparently erroneous judgments about the meaning(s) of the emphasized term(s). This seems indisputable, as shown by both the problematic cases discussed above as well as Unger’s efforts to explain away these intuitions. If emphasis diverts speakers from a term’s literal extension, then the method seems to be unreliable. Since emphasis has the tendency to produce faulty semantic intuitions it ought not to be used to verify the paraphrase or other features of allegedly absolute terms. Unger’s method does not provide dependable access to the “actual meaning of our words”, and subsequently, does not properly distinguish absolute terms.

One might argue that emphasis is an unreliable strategy only if we fail to distinguish between conversational implicature and semantic content.¹⁰⁷ Although emphasizing a term might convey information that deviates from the literal meanings of terms, we can nevertheless safely rely on the device by simply differentiating the semantic content from the speaker's communicative intentions. While we might at first have assumed that 'Sean and Lisa are friends, but not *friends*' means that Sean and Lisa are friends and nothing more, upon some reflection, we recognize that the proposition, as stated, is semantically incoherent. Acknowledging this distinction allows us to see that emphasis does not change the semantic content of terms, thereby preventing the device from generating incorrect intuitions. If this is right, emphasis proves to be a reliable linguistic tool that can be used to verify Unger's paraphrastic procedure.

The problem, however, is that this fix ultimately gets things backwards. Why? Recall that Unger introduces emphasis as way to direct the audience's attention to the literal meaning of the emphasized term, thereby generating intuitions that are consistent with his paraphrase. For example, emphasis is what allows us to appreciate the oddness of the following statement:

'This is *flat*, but that is flatter'.

According to Unger, stressing 'flat' permits the reader to focus on the semantic content of the term (i.e., absolute flatness, with no bumps), which then invokes the intuition that the statement is awkward. But the considerations above seem to suggest that emphasizing a term may lead the speaker to understand something else, the pragmatic implicature of the utterance in which it appears—or more generally, the speaker's intended meaning—

¹⁰⁷ This question was raised by Al Casullo.

instead of the semantic content of the statement. The current suggestion, then, is to separate literal meaning from the speaker's intended content to assist us in properly applying emphasis and understanding its effects. Suppose that we are able to attain a means by which to do so, call it the Criterion of Semantic Content (CSC). For example, we apply CSC to 'This is *flat*' and are able to distinguish the semantic content of this utterance from its pragmatic implicature. It turns out that 'This is *flat*' has semantic content S and pragmatic content P in context C. Perhaps differentiating the content in this way allows us to recognize that emphasis affects only P, but not S. In this way, emphasis can be applied in a way that will avoid confusion between semantic content and the speaker's communicative intentions. This suggestion will not ultimately help Unger. In addition to the fact that there is currently no consistent and universal method resembling CSC for all utterances, the process is circular: Applying emphasis was intended to direct speakers to distinguish the literal meaning of the emphasized term. ('This is *flat*' allows speakers to recognize that 'flat' means absolute flatness rather than some artificial degree of flatness.) Now, it turns out that we are relying on this distinction in order to properly apply emphasis. That is, we use CSC to distinguish the literal meaning of an expression from what is conveyed, and then rely on this distinction to apply the principle of emphasis.

VI. Conclusion

The main worry for skeptical invariatism was its inability to make fine distinctions in epistemic status. Unger offered a pragmatic explanation that accounted for such uses and further presented a test to distinguish pragmatically variable terms in the

language. However, his test was not intuitively supported. To attain the necessary intuitions, Unger appealed to the principle of emphasis. But applying emphasis to the relevant terms appeared to shift their content. Unger then attempted to defend his method by attributing intuitions about the semantic implications of emphasis to semantic confusion. However, emphasis then seemed to offer only unreliable access to the meaning of terms. Moreover, differentiating the semantic content of an utterance from its conveyed content cannot redeem the test as this ends up being circular. Hence, we are left without a viable means of distinguishing absolute terms, those that are semantically stable but pragmatically variable, from other words. Skeptical invariantism, then, remains unable to accommodate subtle epistemic distinctions and for this reason is inferior to a contextual analysis of 'knows'.

Chapter Four Cancelability and Semantic Entailment

Assertions of the form ‘I know that *p*, but I need to check further’ or ‘X knows that *p*, but we need to check further’ appear, *prima facie*, strange. Their oddity might be explained by any number of presuppositions and platitudes that we hold about knowledge: that it is infallible or that it represents a maximal epistemic status. Lewis writes, for example, “to speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory”.¹⁰⁸ Dretske notes that “those who view knowledge as a form of justified (true) belief typically...[speak] not simply of justification, but of full, complete, or adequate justification”.¹⁰⁹

The doubtful acceptability of such statements is particularly significant in the debate between low standard invariantists, who maintain that the standards of knowledge are low and fixed across contexts, and epistemic contextualists, who posit flexible standards. On the contextualist side, Stewart Cohen maintains that the semantic inconsistency of ‘I know that *p*, but I need to check further’ reveals the implausibility of low-standard accounts of knowledge.¹¹⁰ If standards of knowledge were perpetually undemanding, then an assertion that both ascribes knowledge and expresses a need for further evidence would be perfectly acceptable. However, Cohen argues that such

¹⁰⁸ David Lewis, “Elusive Knowledge,” in *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted A. Warfield (Oxford : Oxford University Press, 1999), 221.

¹⁰⁹ Fred Dretske, “The Pragmatic Dimension of Knowledge,” *Philosophical Studies* 40, no. 3 (Nov. 1981): 364.

¹¹⁰ Stewart Cohen, “Contextualism, Skepticism, and the Structure of Reasons,” *Philosophical Perspectives* 13 (1999):58-59.

statements are incoherent, thereby raising doubts against low standard invariantism. Black and Murphy defend invariantism by arguing that such statements while awkward-sounding are semantically unproblematic.¹¹¹

Here I will argue that such statements (what I will call k-conjunctions) appear infelicitous because they are, as Cohen argues, semantically inconsistent. That is, one cannot both assert that she knows that p and that she requires additional evidence for p relative to a single epistemic standard without contradiction. I will begin by presenting Cohen's argument for the semantic deficiency of such statements in which he appeals to H.P. Grice's cancelability test. Next, I will expound Black and Murphy's argument against Cohen, which allegedly shows that these statements are semantically consistent by employing the very same test. In the second section, I will discuss a constraint that significantly restricts the application of the cancelability test, namely that the test proves unreliable in cases of loose talk. Lastly, I will argue that Black and Murphy's examples are cases of loose talk and thus fall outside the scope of the cancelability test. As a result, these examples are not sufficient to refute Cohen's contention that the statements are semantically, rather than pragmatically deficient. Hence, I will have shown that Black and Murphy's attempt to characterize the oddity of 'I know that p , but I need to check further' as a symptom of its pragmatic implications fails.

I. Arguments for and against the infelicity of k-conjunctions

Cohen's Argument for Semantic Deficiency

Cohen's argument relies in part on his now widely dissected airport example¹¹²:
Mary and John are at the airport discussing whether they ought to take a particular flight

¹¹¹ Tim Black and Peter Murphy, "Avoiding the Dogmatic Commitments of Contextualism," *Grazer Philosophische Studien* 69 (2005).

to New York. They overhear someone asking a passenger, Smith, if he knows whether the flight has a layover in Chicago. Smith consults his itinerary and says, ‘Yes, I know—it does stop in Chicago’. It turns out that Mary and John have an important business meeting in Chicago. Mary says to John, ‘How reliable is that itinerary? It could contain a misprint. They could have changed the schedule last minute.’ Mary and John agree that Smith does not know that the plane will stop in Chicago and subsequently decide to check with the airline agent. Does Smith speak truly in ascribing knowledge to himself or are Mary and John correct in denying this claim?

Cohen suggests that we might reason that Mary and John are wrong—perhaps because their standards for knowledge are too high—and Smith is right. This would favor a low standard invariantist account of knowledge attributions, according to which ‘S knows that *p*’ is true just in case relatively low standards of knowledge are met by S. Cohen, however, offers good reason to resist this option: Given Mary and John’s concerns, it seems prudent for them check with the airline agent rather than relying solely on the itinerary. But this makes it unlikely that Mary and John can attribute knowledge to themselves or Smith, since it sounds inconsistent for Mary and John to claim, ‘Smith knows that the flight has a layover in Chicago, but we need to check further’.¹¹³ Moreover, if Smith’s itinerary is good enough for him to know, then it should be good enough for Mary and John to know. Hence, they would also have to assert, ‘We know that the flight stops in Chicago, but we need to check further’, which, according to Cohen, sounds equally strange. Rather, he maintains that the right thing to say is that Mary and John cannot attribute knowledge to themselves or Smith since they cannot

¹¹² Cohen, “Contextualism, Skepticism, and the Structure of Reasons”, 58.

¹¹³ *Ibid.*, 59.

feliculously (and therefore consistently) claim to know that p while also expressing the need for more evidence that p . Cohen's first attempt to argue for the semantic deficiency of k-conjunctions, then, amounts to nothing more than their oddity. Such an argument, however, is certainly insufficient to establish Cohen's conclusion. As suggested earlier, the impropriety of k-conjunctions can be explained by appealing to pragmatic implications, rather than semantic error. What we need is a reason to think that the statement's apparent infelicity amounts to semantic inconsistency.

Cohen goes on to provide a relatively stronger argument for his case, which is intended to thwart the pragmatic strategy. He adopts Grice's cancelability criterion, according to which semantic entailments, unlike pragmatic implications, are not cancelable. For example, 'Jones is an above-average soccer player' implicates that he is not a great soccer player. We can easily cancel this implication by stating 'Jones is an above-average player, in fact, he's a great player'. This shows that the relation between 'Jones is an above-average player' and 'Jones is not a great player' is merely pragmatic. There are other implications, however, that cannot be cancelled. For example, one cannot cancel the implication from 'The tabletop is square' to 'The tabletop has four sides', as the former semantically entails the latter. It is this kind of relation that Cohen assigns to the clauses of k-conjunctions. He maintains that the implication from 'I know that p ' to 'I need not check further' cannot be cancelled, which in turn shows that the relation between these two statements is indeed semantic.

Black and Murphy's Argument for Pragmatic Deficiency

Black and Murphy deny that 'I know that p ' semantically entails 'I need not check further' by offering a series of examples in which the implication is felicitously

cancelled. Since the statements meet Cohen's condition for pragmatic implicature, Black and Murphy contend that the apparent strangeness of k-conjunctions is due to pragmatic failure rather than semantic inconsistency. Their examples follow:

- (1) A student learns that Einstein knows that $E=mc^2$ from a reliable text book. Skepticism about testimony aside, the student knows that Einstein knows that $E=mc^2$. But when asked to provide evidence for the equation on a physics exam, such testimonial support is insufficient. Here, it seems perfectly felicitous for the student to assert, "Einstein knows that $E=mc^2$, but I ought to acquire more evidence that $E=mc^2$."¹¹⁴
- (2) A scientist knows that p on the basis of her evidence that p . But she increases the likelihood of securing a research grant if she acquires more evidence that p . Moreover, her inductive knowledge does not entail the truth of p . Acquiring more evidence will then raise the degree of the probability of p closer to 1. For these reasons, it seems appropriate for her to claim, 'I know that p , but I ought to acquire more evidence that p '.¹¹⁵

The authors further provide an example that shows that the propriety of k-conjunctions is not restricted to inductive knowledge, but extends to deductive knowledge as well.

- (3) A mathematician knows that p on the basis of a deductive proof. However, mathematicians generally prefer multiple proofs, and furthermore, this mathematician aims to publish her findings in a journal. Both considerations provide reason for her to seek further evidence that p . Once again, asserting 'I have a proof that p , but I ought to acquire more evidence that p ' seems perfectly felicitous.¹¹⁶

Black and Murphy maintain that we find k-conjunctions in these contexts felicitous because we can identify a specific reason or purpose for which the subject is checking further. Without such a specification, 'I know that p , but I need to check further' implicates that the subject needs to check further in order to know that p , which suggests that the subject does not already know that p . This implication explains the apparent inconsistency of the statement in some cases. By recognizing that the subject is seeking additional evidence, not for acquiring knowledge, but for some other purpose, such as publishing in a scholarly journal or securing a grant, we can felicitously cancel the implication from 'I know that p ' to 'I do not need to check further'.

¹¹⁴ Black and Murphy, "Avoiding Contextualism," 8.

¹¹⁵ *Ibid.*, 10.

¹¹⁶ *Ibid.*

The same approach can redeem k-conjunctions in the airport case. Although Smith's itinerary is sufficient for Mary and John to know that the flight has a layover in Chicago, they nevertheless require additional evidence to make it more likely that they will meet their business contact. We can then interpret their assertions in the following way: 'We know that the plane stops in Chicago, but we need to check further for the sake of meeting our business contact', thereby canceling the implication that they need not check further. By Cohen's own lights, then, 'I know that *p*' only pragmatically implicates 'I need not check further' and k-conjunctions that cancel this implication can be classified as pragmatically, rather than semantically deficient.

II. Loose Talk and Conversational Implicature

The debate between Cohen and Black and Murphy appears quite straightforward. Cohen maintains that 'I know that *p*' and 'I need not check further' are semantically related in virtue of failing the cancelability test. Black and Murphy counter that k-conjunctions are pragmatically deficient by showing that they satisfy the cancelability criterion. Both sides, however, ignore a certain weakness that Grice attributes to the criterion, namely, that the reliability of the test is compromised in instances of loose or relaxed speech.¹¹⁷ The upshot of this qualification is that satisfaction of the cancelability test proves to be a necessary but not a sufficient condition for pragmatic implication. In order to show that one statement pragmatically implicates another, the cancellation must occur in cases where speakers are not engaged in loose talk.

In the following section I will argue that the speakers in Black and Murphy's examples are speaking loosely when they appear to cancel the implication from 'I know

¹¹⁷ Paul Grice, *Studies in the Way of Words* (Cambridge: Harvard University Press, 1989), 44.

that *p*' to 'I need not check further'. Hence, their use of k-conjunctions is not sufficient to prove the presence of pragmatic implication. Before I can do so, however, it will be necessary to determine precisely what Grice meant by 'loose talk'. Grice himself does not provide an elucidation of his usage and instead seems to rely on a common understanding of the term. Unfortunately, the expression fails to identify a general univocal extension. Indeed, theorists have associated 'loose talk' with various types of utterances including those involving approximations, vagueness, etc.¹¹⁸ The best we can do is to glean an understanding of Grice's conception from the examples he provides.

Macbeth

While watching a scene in *Macbeth* in which it is common knowledge that Macbeth is hallucinating, an audience member felicitously states, 'Macbeth sees Banquo'. Given that Banquo is not before Macbeth, the speaker seems to have cancelled the implication from 'S sees *o*' to the presence of *o*. But, according to Grice, "we should not conclude from this that the implication of the existence of the object said to be seen is not part of the meaning of the word see" or that there is a sense of the term which lacks this implication.¹¹⁹ Rather, Grice maintains that 'see' is here used loosely and for this reason does not indicate that 'S sees *o*' pragmatically implicates the presence or existence of *o*.

There might be some concerns peculiar to the semantics of 'see' that hinder a perspicuous understanding of what Grice had in mind. It might be argued that 'see' admits of more than a single meaning or sense, some of which do not semantically entail the presence of the object seen. For example, the term is often used to mean 'envision',

¹¹⁸ See, for example, Dan Sperber and Deirdre Wilson "Truthfulness and Relevance," *A Quarterly Review of Philosophy* 111, no. 443 (July 2002): 583-632.

¹¹⁹ Grice, 44.

which clearly does not entail the existence of the object seen. Alternatively, some maintain that the meaning of 'see' is to be appeared to in some specified manner. If either definition supplies a part or the whole of the semantic content of 'see', then the existence of the object cannot be entailed by 'see'. Moreover, since the apparent felicity of 'Macbeth sees Banquo' in the example supports this conclusion, the case cannot serve to illustrate the limitations of the cancelability criterion.

Grice does not offer an argument for what he takes to be the conventional meaning of 'see'. I will follow suit. My purpose, recall, is to gain some understanding of what Grice meant by loose talk. Let us then cast worries about the semantic content of see to one side and defeasibly grant that 'S sees *o*' semantically entails the presence of *o*. Now we can ask why the example above would fail to refute this tentative thesis despite canceling the implication. Before considering a plausible explanation, let us consider Grice's second example.

The Tie

Two men are discussing whether to purchase a tie that they both know to be medium green.¹²⁰ Despite their shared knowledge, while viewing the tie in different lights and angles, they say such things as 'It is light green', or 'It has a touch of blue'. Even though each speaker acknowledges that the tie is medium green, they utter and accept statements that attribute other colors to the tie. According to Grice's rule, it would seem that the tie's being medium green only pragmatically implicates (rather than strictly entails) that the tie is not another color as the speakers have felicitously cancelled this implication. Clearly, this is incorrect. An object cannot be both green and blue at the

¹²⁰ Ibid.

same time. The speakers in this example seem to have stated and accepted the cancellation of a semantic entailment, which exposes the shortcomings of the cancelability test. Why is the cancelability criterion inapplicable in such cases? What, in other words, does it mean to use words or phrases in a loose way and subsequently immunize them against the cancelability criterion?

Sentence Non-literality

The speaker in each case above conveys a statement that is distinct from the proposition expressed by the syntactical structure of her utterance. Her assertion appears felicitous because while the literal structure of these utterances contains a semantic cancellation, the content that is conveyed does not. ‘Loose talk’ then seems to pick out specific cases in which the literal content of a statement diverges from what is communicated. We might refine our understanding of his conception by appealing to a phenomenon Kent Bach (2001)¹²¹ calls sentence non-literality, where one is literally saying one thing, but meaning something else. Under this category, Bach makes a further distinction between figurative and sentence non-literality. The latter designates the non-literal use of an entire sentence (rather than the non-literal use of components in the sentence) and it is this that Bach identifies with loose talk. A speaker engages in loose talk when she intends to convey a proposition that contains an implicit qualification that is absent from her literal statement. This qualification might take the form of an implied temporal indexing or other contextual references. Bach offers the following examples: When one says ‘I haven’t taken a bath’, she most likely does not mean to say that she has never bathed. Rather, she intends to convey a qualified proposition, perhaps one of the

¹²¹ Kent Bach, “Speaking Loosely: Sentence Non-literality,” *Midwest Studies in Philosophy* 25 (2001): 249-263. “You don’t say,” *Synthese* 127 (2001):11–31.

following: ‘I haven’t taken a bath *today*’, or ‘I haven’t taken a bath since I worked out’. A host who tells his guests ‘There’s more beer in the fridge’ would be taken to be misleading his guests (perhaps inadvertently) if the only beer in the fridge was contained in a small puddle. Note, however, that the speaker has not said anything false, as there was, after all, some beer in the fridge. But he implicitly conveys a qualified version of the statement, specifically that there is more beer *to drink* in the fridge.

There are different ways that conversational participants recognize that the speaker means more than what she has stated. Bach maintains that listeners are typically able to infer the speaker’s implicature by recognizing that the utterance taken at face value would violate certain Gricean maxims: “for it is on the presumption that the speaker is being truthful and relevant (and intends to be taken to be) that the hearer figures out what the speaker means in saying what he says”.¹²² For example, the sentence ‘I have not eaten’ will not generally be used to mean that the speaker has never eaten, but rather something more precise: he has not eaten supper, or that he has not eaten today, etc. The audience can figure out the implicature, or what the speaker intends to convey, by recognizing that the speaker could not have meant what he literally stated.

There are two points to note about recognizing instances of loose talk (or what Bach also refers to as conversational implicature). First, Bach’s account does not require the audience to undergo an explicit reasoning process, which involves first grasping the literal statement, and then inferring the communicative intentions of the speaker. Rather, such steps might be tacitly performed and the relevant proposition realized without conscious recognition of the literal statement. Second, Bach emphasizes that violations

¹²² Bach, “Speaking Loosely: Sentence Non-literality,” 255.

of these Gricean maxims comprise some of the triggers that might be used to communicate and grasp meaning, but they do not exhaust all the ways in which speakers convey and the audience understand the speaker's intention. Essentially, the audience has to recognize a connection between what the speaker states and what he means to communicate, which might involve several different features of the context, including the goals of the conversation.

Some distinctions between conversational implicature and other kinds of speech are called for. First, we must be careful not to confuse loose talk with Gricean pragmatic implicature. Loose talk is a narrower conception than implicature as the latter also involves figurative and other forms of non-literal speech. Moreover, loose talk is confined to utterances that omit certain qualifications or specifications that would make the statement literally precise. Conversational implicatures, however, are not restricted in the same way. We might then think of conversational implicatures as one species of the genera conversational implicature as there are ways in which the former resembles the latter. Bach's conception is also distinct from what others have termed loose talk, including expressions of vagueness or approximations in which one speaks roughly but nevertheless means what he says. Exaggerations and understatements that distort the truth to create a certain effect also constitute a different category as Bach's brand of loose speech does not involve using expressions in these ways.

Gricean Loose Talk

Bach's sentence non-literality is consistent with Grice's conception of loose talk. Consider the first example in which the speaker states 'Macbeth sees Banquo'. Clearly,

she did not mean that Macbeth sees something that is identical with who or what they know as ‘Banquo’ in the play. Rather, the speaker meant to convey that Macbeth perceives an image of Banquo or that Macbeth is undergoing a hallucination of Banquo, neither of which implicates the presence of Banquo. Despite the absence of such qualifications in her literal utterance, the audience understands that the speaker is not using ‘see’ in a conventional sense, but rather to indicate that Macbeth ‘seems to see’ Banquo, producing a statement that is different than the one articulated by the speaker. This specification of the way in which Macbeth perceives Banquo is implicit in the speaker’s utterance.

In Grice’s second example, two men make claims such as ‘The tie is blue’ and ‘It is light green’ despite knowing that the tie is medium green. The speakers convey, however, that the tie is blue or light green only under particular conditions, perhaps, ‘*under this light*’ or ‘*at this angle*’. Therefore the statement that is conveyed does not contain a contradiction as the speakers mean to cite specific circumstances under which the tie is a different color. Once again, Grice’s example clearly portrays Bach’s notion of loose talk as speakers assert statements that contain implicit qualifications.¹²³

III. Loose Talk in Black and Murphy’s Examples

In uttering k-conjunctions, speakers in Black and Murphy’s examples implicitly convey a qualification that is absent from their literal utterance, namely, two different standards of epistemic evaluation—one in which they know that p and another in which they do not know that p and must check further. This feature indicts such cases as

¹²³ In fact, Bach himself suggests that Grice might have meant this in his discussion. He writes, ‘Grice occasionally alluded to what I’m calling implicature as when he remarked that it is often ‘unnecessary to put in qualificatory words’.

instances of loose talk, and thus renders them inadequate to establish that 'I know that p ' pragmatically implicates, rather than semantically entails, 'I need not check further'.

Why suspect that k-conjunctions contain such a specification? It seems that in each of the cases, we can reasonably question why the speaker needs more information for the specified purpose despite knowing that p . That we are entitled to challenge the speakers may not be obvious because we tend to immediately grasp the specifications that the speaker has omitted. We tacitly recognize, for example, that the scientist appeals to higher standards when she expresses the need to check further, even though she fails to mention such a qualification, as the process of applying for research grants is widely recognized as having stringent standards of evaluation. But imagine someone unfamiliar with the rigmarole of grant applications, who overhears the scientist's claim. It seems quite appropriate for him to ask why, if the scientist has performed enough experiments to know that p , grant referees might demand even more evidence. The researcher can respond in one of two ways: she can either (1) retract her knowledge ascription or (2) admit that there are two different standards in play. If she chooses (1), then it is clear that her use of the k-conjunction was not sincere and no cancellation has occurred. On the other hand, she might opt for (2), replying that while she knows that p according to a certain standard—perhaps set by her own lab or colleagues—the referees' standards are more stringent than her own or the wider scientific and social community. She means to say 'I know that p , but according to the grant referees' tough standards, I don't. So, I will need to gather more evidence.' The speaker essentially claims to know that p in one context, but admits she does not know in another, more demanding context, for which she will have to acquire more evidence.

That subjects are speaking loosely in Black and Murphy's cases might be difficult to accept because, as suggested earlier, we tend to automatically recognize that grant referees and highly reputed journals employ rather high standards of evaluation to grant that p or that one knows that p . Suppose, however, that the scientist is planning to visit a young niece rather than apply for a research grant. She asserts the same sentence: 'I know that p , but I ought to acquire more evidence'. We are clearly at a loss as to why visiting her niece would demand further evidence. It would seem reasonable for us to question why, if she has enough evidence to know that p , she must gather more evidence for visiting her niece. The scientist must respond by either withdrawing her knowledge claim or recognizing two different standards. She might uphold her claim to know by explaining that her niece is exceptionally skeptical. By doing so, the researcher explicitly recognizes that her niece represents a higher standard of knowledge. The scientist might say, '*Ignoring my niece's unattainable standards or With the exception of my niece's skepticism, I know that p* '. She claims to know that p in a context of lower standards, but not in one that imposes higher standards of knowledge.

When a speaker utters k-conjunctions, she implicitly conveys two different standards of knowledge, even though this qualification is absent from her literal statement. This is made clear by the fact that in each of these cases, the audience is well within her linguistic right to question why the speaker needs more evidence for the specific purpose, given that she has already claimed to know that p . This point is obscured in Black and Murphy's cases because we take it for granted that the speaker is talking about two different standards. However, shifting a few non-essential features of the case makes it evident that a reasonable specification is absent from the speaker's

literal utterance. The speaker can then clarify her utterance only by executing one of two responses: (1) the speaker may deny the original knowledge ascription or (2) distinguish the low standard context in which she knows from the high-standard context in which she does not know and must check further. Since the speakers in Black and Murphy's examples intend to convey qualifications absent from their literal statement, it is clear that the subjects are engaged in loose speech. Hence, their uses of k-conjunctions do not suffice to show that 'I know that p ' merely pragmatically implicates 'I need not check further'.

IV. Residual Concerns

I have argued that the subjects in Black and Murphy's examples are speaking loosely when they utter k-conjunctions because they fail to articulate a particular qualification that is nevertheless communicated to their audiences. Hence, their examples lie outside the scope of the cancelability test, and therefore do not suffice to show that 'I know that p ' merely pragmatically implicates 'I need not check further'. My opponents, however, might respond that that subjects are speaking loosely when they utter the k-conjunction is irrelevant because the proposition ultimately conveyed also cancels the implication. When we clarify the speaker's statement so that her qualification is explicitly stated, the resulting proposition cancels the implication from 'I know that p ' to 'I need not check further'. The speaker maintains that she knows that p , but needs to check further, given the demands of a specified purpose. Therefore, the limitations of the cancelability test do not undermine Black and Murphy's argument that the two statements are pragmatically related.

Black and Murphy's conclusion, however, cannot be sustained. Upon closer inspection, it is clear that speakers are not asserting genuine cancellations. Rather, they claim to know that p , but then recognize a need for further evidence from the perspective of a *different* judge, one who (or a group of who) will employ tougher standards of knowledge ascriptions. The same can be observed with other terms that incorporate contextually determined standards of application. While the statement 'Billy is tall, but he is not tall' sounds contradictory, (just like 'I know that p , but I need more evidence for p ') qualifying the second clause by introducing a different evaluative perspective removes the inconsistency: 'Billy is tall in comparison to the rest of his class, but not according to professional NBA standards'. The speaker is not *cancelling* her claim that Billy is tall, but rather asserting that he is tall according to one comparison class, but not another. In the same way, the scientist in Black and Murphy's example is not cancelling the implication from 'I know that p ' to 'I do not need more evidence'. Rather, she specifies that she needs more evidence according to a more stringent epistemic standard in which she does not know that p . The scientist means to say, 'I know that p , *but given the referees' standards*, I need to check further'. Thus, we can uphold Cohen's thesis that one cannot both ascribe knowledge to oneself and claim to need further evidence from a single context or perspective.

Chapter Five Knowing Better

Contextualists draw analogies between ‘knows’ and other context-sensitive expressions in order to strengthen and support their claim that knowledge ascriptions vary in their truth conditions from one context to another. Cohen, for example, maintains that ‘knows’ is similar to ‘tall’, ‘rich’ and ‘poor’, each of which can come in varying strengths and depends on certain contextual factors for its application.¹²⁴ Cohen’s analogy strengthens the contextualist thesis by showing that the semantics they attribute to ‘knows’ is not unique, but rather applies to a particular class of expressions of which ‘knows’ is a member. Further, it supports contextualism by showing the manner in which ‘knows’ is context-sensitive: Like ‘tall’ and ‘rich’, ‘knows’ comes in varying degrees of strength and relies on the context for the appropriate standard of application. It can therefore be classified as a gradable context-sensitive expression. If ‘knows’ cannot be identified within a broader linguistic category like this one, its context-sensitivity becomes suspect as contextualists would have to attribute to it a sui-generis semantics. Thus, such group-membership and analogies are crucial to the contextualist program.

Jason Stanley¹²⁵ has rejected Cohen’s analogy by showing that ‘knows’ fails certain linguistic tests of gradability: knows cannot be modified (using modifiers such as ‘very’ or ‘really’) and does not allow for comparative constructions (using ‘more’ and ‘better’). This is a potentially devastating objection for contextualism because it

¹²⁴ Cohen, “Contextualism, Skepticism and the Structure of Reasons,” 60.

¹²⁵ Jason Stanley, “On the Linguistic Basis for Contextualism,” *Philosophical Studies* 119 (2004): 119-146.

challenges the justification of the context-sensitivity of ‘knows’ by appeal to gradable expressions such as ‘tall’ and ‘rich’, and subsequently raises a more general doubt about the context-dependence of ‘knows’. In response to Stanley’s charge, proponents of contextualism have identified ‘knows’ within alternative *non-gradable* linguistic categories such as modified comparatives (‘sufficiently tall’) and adverbially modifiable verbs (‘snore’). More specifically, they accept Stanley’s argument against the gradability of ‘knows’, but argue that contextualists are not committed to propositional knowledge ascriptions coming in different degrees, and then go on to offer alternative models that display the alleged non-gradable context-sensitivity of ‘knows that *p*’. This allows contextualists to identify ‘knows’ within a linguistic category (albeit a non-gradable one), and therefore, strengthen and support their semantic claim.

I want to argue that such recent defenses of contextualism are misguided, precisely because they do not allow for the degree variability of propositional knowledge ascriptions. I will argue here that attributions of knowledge can in fact come in degrees within certain contexts. My argument will proceed as follows: (1) In some contexts, statements of the form ‘S knows that *p*’ are true if and only if S knows why *p*. (2) Knowledge why can come in degrees. (3) Therefore, propositional knowledge ascriptions can also come in degrees within these contexts. If this is the case (that is, if it is possible for propositional knowledge attributions to come in degrees within a contextualist framework) then we have reason to reject non-gradable accounts of the context-sensitivity of ‘knows’, in favor of Cohen’s original analogy, which correctly identifies ‘knows’ as a gradable context-sensitive expression.

This chapter will then have the following structure: first, I will briefly rehearse some of the current positions regarding the (non)gradability of ‘knows’, including an exposition of Stanley’s original tests for gradability and Daniel Halliday’s and Michael Blome-Tillman’s responses to his claims. In the second section, I will offer cases that show that knowing that p under certain circumstances requires the subject to know a larger set of information that explains the truth of p . I will provide an analysis of such cases in the third section, which will be divided into four parts: first, I will argue that knowledge-why must be understood in terms of Schaffer’s anti-reductionist account of knowledge-wh. That is, S knows-wh iff: S knows that p as the true answer to the indirect question Q . Second, I will argue that knowing the true answer to the question why p will involve two particular features, (i) knowing a subset e of statements within a set x that explains the truth of p and (ii) having true beliefs about the interrelationships between propositions within e and p . These two features allow knowledge-why to come in degrees. Third, I will argue that in the cases described in section two, ‘S knows that p ’ is true if and only if S knows why p , given my account of knowledge-why. These points will allow me to conclude that propositional knowledge ascriptions (insofar as they are equivalent to knowing why) can, in fact, come in degrees. Fourth, I will address plausible objections to my account of knowledge-why and my claim that the high-standard scenarios described in section two require S to know why in order to know that p . In the final section, I will argue that ‘knows’ is correctly identified with terms such as ‘tall’ and ‘rich’ which reflect the gradable context-sensitivity of ‘knows’.

I. Gradability and Contextualism

Stanley

Gradable expressions come in different degrees of strength. A gradable expression F denotes a property that is measurable along a scale, such that something is more or less F depending on its position on the scale.¹²⁶ Cohen maintains that ‘knows’ is a gradable term because it entails ‘justified’ which comes in different degrees. This supports Cohen’s claim that ‘knows’ can be identified within the category of gradable, context-sensitive expressions, which both rely on the context for their application and can come in varying degrees of strength. Stanley, however, provides two tests for gradability that seem to show that ‘knows’ cannot be categorized as a gradable context-sensitive expression.¹²⁷ This suggests that contextualists cannot appeal to terms such as ‘tall’ and ‘flat’ to justify the context-dependence of knowledge ascriptions, and therefore raises doubts about the central contextualist claim that the standards of knowledge shift. Stanley’s two linguistic tests for gradability are as follows. The first test requires that gradable expressions admit of modifiers:

- a. That is very flat.
- b. That is really flat.
- c. John is very tall.
- d. John is really tall.

The second test requires that the term have a natural comparative construction such as ‘flatter than’, ‘taller than’, etc. Stanley demonstrates that ‘knows’ passes neither test by pointing out the oddness of statements such as “John very knows that penguins waddle” and “Hannah knows that Bush is president more than she knows that Clinton was

¹²⁶ Such a definition of ‘gradability’ can be found in Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005), 77 and Christopher Kennedy, “Vagueness and grammar: the semantics of relative and absolute gradable adjectives,” *Linguist Philos* 30 (2007): 3.

¹²⁷ Stanley in fact presents a further objection against transferring the semantic properties of ‘is justified’ to ‘knows’. I will discuss this in section 3.

president”.¹²⁸ He considers some locutions which would suggest that ‘knows’ is in fact gradable, but argues that such cases are only apparent:

1. John knows Bill better than Mary does.
2. I know very well that Bush is president.
3. Hannah knows logic better than John does.

Stanley maintains that such propositions do not in fact display the gradability of propositional knowledge ascriptions, as (1) and (3) indicate an acquaintance relation which is not held between a person and a proposition. Stanley classifies (2) as a pragmatic indicator, since the corresponding negative construction is unacceptable:

4. I don’t very well know that Bush is president.

Stanley concludes that knowledge ascriptions cannot be classified as gradable, or as coming in varying strengths, for these linguistic reasons.

Responses to Stanley

Halliday: Non-gradable modified comparatives

Daniel Halliday defends contextualism by accepting Stanley’s charge that ‘knows’ is a non-gradable expression, but argues that this does not necessarily deprive contextualism of linguistic support.¹²⁹ He maintains that contextualists are not committed to the gradability of ‘knows’, but only to the claim that the standards of ‘knows that *p*’ are flexible across contexts.¹³⁰ These two features (gradability and the flexibility of standards) can and do come apart: while comparative adjectives such as ‘expensive’ and ‘tall’ display both, ‘knows’ does not.¹³¹ This explains why the former

¹²⁸ Stanley, “On the Linguistic Basis,” 125.

¹²⁹ Daniel Halliday, “Comparatives and Gradability,” *Philosophical Studies* 132 (2007): 381-393.

¹³⁰ Halliday distinguishes these two features as gradable₁ and gradable₂. I chose not to use his classification as it could potentially confuse matters. For our purposes and according to the way I’ve defined gradability (as the capacity to come in different strengths), Halliday identifies ‘knows’ as a non-gradable expression.

¹³¹ Halliday further argues that appealing to modified comparatives also accommodates a certain platitude about knowledge, which is that “once an individual’s epistemic position relative to a proposition becomes

provides a poor analogy for 'knows'. Halliday argues instead that 'knows' belongs to the category of modified comparatives such as 'sufficiently tall', which does not come in varying degrees of strength, but whose standards of application are sensitive to the context. For example, while a child may be 'sufficiently tall' for entry into a ride at an amusement park, she might not be 'sufficiently tall' in a context where conversational participants are attempting to reach a Frisbee that has landed on the top branches of a tall tree. Halliday's account, then, defends contextualism by (i) conceding that 'knows' is not a gradable expression but (ii) arguing that contextualists are not committed to the gradability of 'knows that *p*' and (iii) offering a non-gradable context-sensitive linguistic model of 'knows'.

Halliday presents a plausible analogy, and a convincing defense of contextualism. I will argue, however, that such an analysis must be rejected as it is in fact possible for propositional knowledge ascriptions to come in degrees within a contextualist framework, one in which the truth conditions of knowledge-ascribing statements are sensitive to features in the context of utterance. Specifically, in some contexts, 'S knows that *p*' will be true if and only if S knows why *p*, which allows attributions of knowledge *in such contexts*, to come in degrees. Even though gradability might be conceptually distinguishable from context-sensitivity, Halliday's analogy will not work if it does not fully capture the semantics of 'knows that *p*'. Therefore, we must identify 'knows' with

strong enough for that individual to be truly ascribed knowledge of that proposition, then further improvement in the epistemic position of that individual is redundant... [and] does not amount to their knowing that *p* in any way being improved" (Halliday, 386). It is not clear, however, that this is true of all modified comparatives: One might say, for example, "She is very tall, but he is taller" or "This is really flat, but that is flatter". This, in turn, raises doubts about his analogy. Halliday might respond that 'knows' belongs to a certain subset of modified comparatives that entails a threshold such as "most expensive". While this is an improvement, his theory cannot accommodate the main worry discussed above, which is that in some contexts propositional knowledge ascriptions can come in degrees. Accounts that link 'knows that *p*' with non-gradable expressions cannot accommodate this feature.

a different linguistic category, which reflects the gradability and context-dependence of ‘knows’. Before discussing this, however, I want to present one more response to Stanley’s charge which also classifies ‘knows’ as a non-gradable expression.

Blome-Tillman: Non-gradable Verbs

Michael Blome-Tillman¹³² defines semantic gradability as having a link to a scale that measures the gradable property associated with the expression in question. It is important to note that in his view, ‘gradable’ terms do not themselves come in different strengths. Rather, they merely require a link to a scale possibly via one of its components. This is distinct from my definition (and others’) of ‘gradability’ according to which the term itself can come in different degrees of strength. I will then refer to Blome-Tillmann’s notion of gradability as ‘semantic gradability’. He argues that while ‘knows’ is semantically gradable (that is, linked to a scale that measures a gradable property), it need not pass the tests of *syntactic* gradability that Stanley proposes.¹³³ He maintains that “the relation between a verb’s semantic links to scales and its syntactic gradability is not as strict as Stanley expects it to be”.¹³⁴ Consider, for example, ‘snore’, whose truth conditions vary along a scale of loudness: what may count as ‘snoring’ in one context, say, in a concert hall, may merely count as ‘heavy breathing’ in another context, say, in a dormitory. Although ‘snore’ unarguably fails Stanley’s tests, we may nevertheless indicate its underlying scale with the following constructions:

4. x snores very/quite/extremely loudly.
5. x snores more loudly than y.¹³⁵

¹³² Michael Blome-Tillmann, “The Indexicality of ‘knowledge’,” *Philosophical Studies* 138 (2008): 29-53.

¹³³ As will be shown shortly, Blome-Tillmann’s account is similar to Cohen’s original claim insofar as he ties the semantic gradability of ‘knows’ to justification, which can come in different strengths. However, he stops short of claiming that ‘knows’ itself comes in degrees because of its relationship with justification.

¹³⁴ Blome-Tillmann, “The Indexicality of ‘knowledge’”, 45.

¹³⁵ *Ibid.*

While ‘snore’ is linked to an underlying semantic scale, its semantic gradability is not syntactically indicated in the same way as gradable adjectives. That is, individual events of snoring can be ordered according to their loudness. From this, Blome-Tillman concludes that expressions that are semantically gradable may be either syntactically gradable in the manner portrayed by Stanley’s tests or adverbially modifiable along a scale *s*. We can then indicate the semantic gradability of ‘knows’ via adverbial phrases that establish orderings along a scale of epistemic strength:

6. x knows with very/quite/extremely good evidence/justification that *p*.
7. x knows with better evidence/justification than *y* that *q*.¹³⁶

Thus, Blome-Tillmann concludes that ‘knows’ belongs to the linguistic category of semantically gradable verbs. Just as ‘snore’ links to a scale of loudness, ‘knows’ links to a scale of epistemic strength, and can accept adverbial modifiers along this scale. He agrees with Halliday that contextualism does not commit us to the view that ‘knowledge’ itself comes in degrees, but rather that knowledge-states “can be ordered with respect to their epistemic strength”.¹³⁷

While Blome-Tillmann’s analysis allows ‘knows’ to correspond to a scale, note that he explicitly denies that ‘knows that *p*’ can itself come in different strengths. On his account, then, it is possible to attribute varying degrees of justification or evidence, but not knowledge. In what follows, however, I will argue that propositional knowledge ascriptions can come in degrees. This suggests that Blome-Tillmann’s analogy must be rejected in favor of one that accommodates the gradability of ‘knows’.

¹³⁶ Blome-Tillman, 46.

¹³⁷ *Ibid.*, 47.

The common response, then, to Stanley's charge is to (a) accept that 'knows' is not gradable and cannot come in different strengths, but (b) claim that contextualists are not committed to the gradability of 'knows' and (c) offer alternate non-gradable linguistic models that reflect the manner in which 'knows' is sensitive to the context. These three moves allow proponents of contextualism to preserve the central contextualist claim that the standards for 'knowing that p ' vary according to the context. This defense of contextualism, however, is not sustainable. This is because contextualists are in fact committed to the degree variability of propositional knowledge, as I will demonstrate in the following sections. For this reason, the non-gradable linguistic models proposed by Halliday and Blome-Tillman must be rejected as they do not capture this feature of the semantics of 'knows'. Rather, we must identify 'knows' with a different linguistic category. In the next section, I will set the stage by offering examples in which it appears that the truth of knowledge ascriptions seems to depend on the subject's knowledge of the *explanation* for the relevant proposition, or their knowledge of why it is true. The third section will provide an analysis of what is going on in such cases, showing specifically that statements of the form 'S knows that p ' are true in such circumstances if and only if S knows why p , subsequently allowing propositional knowledge ascriptions to come in different degrees of strength.

II. Knowing that and Knowing Why: Cases

Here, I will present cases in which it seems that the subject knows that p if and only if she knows a broader set of statements that explain the truth of p , or in other words, why p is true. I will describe two versions of each case: the first will be a low-standard

scenario in which the truth conditions for knowing that p do not involve knowing why p , and the second, a type of high-standard scenario in which they do.

PTSD A

After coming home from his deployment in Afghanistan, a soldier begins to experience wild mood swings, depression and nightmares. He tells his wife that he has Post Traumatic Stress Disorder from his time in battle. But his wife, responds, “Are you sure? You haven’t been sleeping well; maybe you just need more rest?” The husband responds, “No I know that I have PTSD, because my doctor told me at my appointment earlier today”.

PTSD B

The circumstances are the same as the case above, but here, we focus on the doctor’s perspective. The doctor presents his patient to a number of his colleagues telling them that the patient he just saw is suffering from PTSD. At this point, one of his colleagues responds, “Are you sure? A lot of patients coming back from war have prior illnesses and their existing mental illnesses are exacerbated during their time in battle”. The doctor responds, “I know that he has PTSD because he displayed the classic symptoms and has no prior history of mental illness. Moreover, I’ve seen a lot of war veterans who are in the same shape being effectively treated for PTSD”.

Both the doctor and the patient’s self knowledge attributions seem appropriate, despite a substantial difference in the epistemic status of their belief that the patient has PTSD. While the patient has little knowledge about diagnosing criteria or the medical treatment of other war veterans, his belief that he has PTSD is justified via testimony from a reliable source. In contrast, the doctor’s colleagues raise counter-possibilities that are quite plausible: that the patient might have a different mental illness, that war worsens the symptoms of previous mental illnesses. The doctor, however, correctly claims to know that the patient has PTSD because he can eliminate such plausible, but nevertheless challenging counter-possibilities. It is his knowledge of a network of explanatory statements that explain the patient’s diagnosis and their interrelationships that verifies the

truth of his knowledge claim. I will say more about this in the next section. Before that, let us consider one more case:

Montessori Schools A

A group of friends have met for brunch and their conversation turns to education. One of the speakers, Rose, says “The Montessori school system is better than others for children who are independent learners”. Her friend, however, says, “I sent my children to good public schools, which I think provide the same opportunities”. Rose, however, responds, “I know that the Montessori school system is particularly good because my daughter, an education specialist, was just telling me”.

Montessori Schools B

Here, we consider the case from the perspective of Rose’s daughter, Lily. Lily is giving a presentation on the best types of schools for children with varying temperaments. Lily announces, “We know that the Montessori school system is better than others for children who are independent learners”. Those attending the conference will be other researchers and parents who are seeking information on where to send their children to school. Such informed and interested audience members raise a number of more difficult counter-possibilities that were not relevant in the previous case including, “How can children succeed in the future if they don’t learn the basic subjects taught in a more traditional way?” But Lily insists “We know that the Montessori System is more suitable for such kids because we have studied the way in which they absorb and process information. The learning environment offered by the Montessori system effectively promotes their intellectual and social development, which will help them succeed in the future, even though their learning environments may change then”.

Like the *PTSD* cases above, both Rose and Lily seem to be truthfully ascribing knowledge to themselves despite different epistemic statuses with respect to their belief that the Montessori School System provides an excellent learning environment for independent learners. What does this difference consist in? I will argue in the next section that Lily is able to sustain her knowledge ascription that p because of her knowing why p .

III. Knowing that and Knowing Why: Analysis

Both sets of cases were intended to highlight an important feature of the truth conditions of knowledge ascriptions. Under certain circumstances, an individual's claim to know that p will depend on her grasp of other propositions that explain the truth of p . Here, I will provide an analysis of these cases: First, I will offer an account of knowing-why. This will include two parts: (i) I will adopt Schaffer's account of knowledge-wh, but (ii) show that knowledge why involves features that diverge from knowing that and which allow knowledge-why to come in different degrees. Second, I will argue that the subjects in both B cases above must know why p in order to know that p . That is, 'S knows that p ' in these contexts if and only if S knows why p . Given the degree-variability of knowledge-why, and the equivalence between 'S knows that p ' and knowledge-why in such contexts, I will conclude that propositional knowledge ascriptions can come in degrees. Lastly, I will consider objections to both my analysis of knowledge-why and my claim that subjects in the B cases in fact must know why p in order to know that p . My responses will hopefully not only successfully defend my position, but also clarify it to a greater extent.

Reductionism vs. Anti-Reductionism

In both case B's described above, S's knowledge that p is based on knowing why p , or so I will argue. What do I mean by knowledge-why? One option is to treat the latter as a species of knowing that. Reductionists such as Lewis, Higgonbotham, and Stanley and Williamson have offered accounts which reduce 'knowledge-wh'

(knowledge why, where, what, who, when, how) to knowledge-that statements.¹³⁸ They argue that the proposition expressed by statements that include ‘knowledge-wh’ can be explained in terms of knowing-that. The relation denoted by knowledge-wh is not significantly different from the knowledge-that relation. Here is one formulation of their view, cited in Brogaard (2009):

Reductionism: S knows-*wh* iff there is a proposition *p* such that S knows that *p* and *p* answers the indirect question of the *wh*-clause.¹³⁹

Since we are here only interested in knowing-why, I will restrict my comments for the remainder of this section on this subset of knowledge-wh statements. The idea, then, is that S knows why *p* just in case there is some proposition *e*, such that (i) S knows that *e* and (ii) *e* answers the question ‘Why *p*’? This is a reductive account in the sense that statements of the form ‘S knows why *p*’ denotes the same relation as knowing-that, with the further condition that the proposition known answer the relevant wh-question. Such a reductive analysis seems inadequate, however, to knowing why: It might be, for example, that I know the criteria for PTSD as listed in the DSM, that the patient exhibits these symptoms and that individuals suffer from PTSD upon undergoing a traumatic event for a prolonged time period. Further, these propositions together answer the question of why the patient is diagnosed with PTSD. But, I might not register the link between these factors and the fact that the patient’s correct diagnosis is PTSD. This can happen because I am not attending to the explanatory, probabilistic and other types of

¹³⁸ David Lewis, “Whether Report,” in *Papers in Philosophical Logic* (Cambridge: Cambridge University Press, 1998), 45-56; James Higginbotham, “The Semantics of Questions,” in *The Handbook of Contemporary Semantic Theory*, ed. Shalom Lappin (Oxford: Oxford University Press, 1996), 361-383; Jason Stanley and Timothy Williamson, “Knowing How,” *Journal of Philosophy* 98 (August 2001): 411-444.

¹³⁹ Berit Brogaard, “What Mary did Yesterday,” *Philosophy and Phenomenological Research* 78, No. 2 (March 2009): 440.

connections between these statements. Perhaps my cognitive apparatus is faulty or my reasoning skills impaired or I simply do not have true beliefs about the reliability of the DSM in mental illness diagnosis or the fact that war veterans experience similar psychological problems. If I fail to grasp the way in which these statements cohere or the way in which the set of them explain the patient's diagnosis, then it certainly seems as though I don't know why the patient has PTSD. That is, even though there exists an e such that (i) I know that e and (ii) e answers the question of why p ?, I still don't know why p .

Brogaard provides the following counter-example along the same lines:

Suppose that S knows that Lisa left the party early but has never heard of Lisa's boyfriend Jim. Suppose furthermore that Lisa's boyfriend Jim is upset because Lisa left the party early. If to know why Jim is upset is to know that Lisa left the party early, then S knows why Lisa's boyfriend Jim is upset. But if someone were to ask S why Jim is upset, then S would be unable to provide an answer.¹⁴⁰

Even though Lisa knows the very proposition that answers the question of why Jim left the party, it still seems as though she does not know why Jim left the party, as she has not made the connection between Jim's being upset and Lisa's leaving. Both examples suggest that there is something more to knowing why.

Schaffer's anti-reductionist account avoids this worry. For Schaffer, having why knowledge requires not only knowing a proposition but knowing that the proposition is the true answer to the why-question. He writes that "it is not enough to know the

¹⁴⁰ Berit Brogaard, "I know, therefore I understand," *typescript*: 8. The reductionist might respond that "Lisa left the party" is not a complete answer to the wh-question. Rather, the correct answer is: 'Jim was upset because Lisa left the party'. Hence, it will turn out that even on their account, S does not know why Jim is upset. Note that such a requirement would impose certain conditions on what qualifies as an answer to a question: not only must the statement provide the answer to the question, but it must contain some further indication of the question itself. As such, this is an ad hoc response because in most cases, people will intuit that something qualifies as an answer whether or not it explicitly indicates what it is trying to answer.

proposition that just so happens to be the answer—one must know the answer as such’.¹⁴¹

S knows-wh iff KspQ, where Q is the indirect question of the wh-clause, and *p* is its true answer.¹⁴²

KspQ is to be read as: S knows that *p* as the true answer to Q. Brogaard writes that this necessitates the following requirement on wh-knowledge:

Knowledge-wh: to know-wh is to know that *p* is the true answer to the indirect question of the *wh* clause.¹⁴³

Schaffer’s analysis is an improvement over the reductionist account as it handles the two counter-examples described above. That is, in order to know why *p*, the subject must know that *e* as the true answer to the question why *p*. To illustrate: someone who knows the criteria for PTSD as listed in the DSM, and that the patient exhibits these symptoms and even that other war veterans (and others subjected to a traumatic event for a prolonged period of time) have been diagnosed with and treated for PTSD, but who does not link these as answers to the question ‘Why is the patient diagnosed with PTSD’ will not count as knowing why he is diagnosed as such. Schaffer’s account fares better in Brogaard’s example as well: it will turn out that S does not know why Jim is upset because S does not know that Mary left as the true answer to the question, ‘Why is Jim upset?’

¹⁴¹ Jonathan Schaffer, “Knowing the Answer,” *Philosophy and Phenomenological Research* 75, No. 2 (September 2007): 392.

¹⁴² Schaffer, “Knowing the Answer,” 392.

¹⁴³ Brogaard, “I know, therefore I understand,” 8.

What it takes to know the true answer

On the anti-reductionist account, knowing why requires not only knowing a proposition, but knowing that a particular proposition answers the why question. Schaffer's account then identifies that the subject must know the connection between the answer and the wh-question, and goes some way toward providing a satisfactory analysis of knowing why in the cases above. Schaffer's analysis, however, brings us only so far. It is a general analysis of knowledge-wh questions and as such, does not explicate what it takes to know that something is the true answer to the question why. This is important because knowing the true answer to a *why* question, I take it, can involve knowing not a single proposition, but rather a set of statements that collectively explain why p and the way in which they explain why. Here is a formal proposal:

S knows e is the true answer to why p iff: p , S believes that p and S knows a subset e within a set x that explains the truth of p , and has true beliefs about the explanatory, probabilistic, and other types of relationships between the members of the subset and p .

There are, then, two things to note about the analysis. First, it shares certain conditions of 'S knows that p ', including the truth of p and S believes that p . But there are two important differences between the way I've construed knowing-why and knowing-that: Knowledge why involves (i)knowing some members within a set of statements that contribute to a full explanation of p (that is, for each propositional member $m_1 \dots m_n$ in subset e , S knows that m) and (ii)having true beliefs about the interrelationships therein.¹⁴⁴ The last two features allow knowing-why to come in degrees. That is, the subject might know only a minimal subset that explains the truth of

¹⁴⁴ The use of 'knowing' as it appears in the definition will be further clarified in the objection section below. Here, I simply want to draw attention to the way in which the truth conditions of 'S knows that p ' when uttered in a low-standard context, differs from its truth conditions in the types of high-standard scenarios depicted above, where 'S knows that p ' if and only if it is true that S knows why p .

p , or may know a larger subset. For example, S knows why the door is open just in case S knows that ‘Joan always leaves the door open’ is the true answer to the question ‘why is the door open?’ But one might further know that the door is difficult to latch, and one cannot close it without pulling hard, the door does not close on its own and that Lisa is too careless to pull the door hard enough as the answer to ‘why the door is left open’. Knowing-why p , then, can consist of knowing a larger or smaller group of statements within the explanation for p . Further, the subject might grasp the interrelationships between these statements to different degrees. What this shows is that such knowledge can vary along two different axes: the size of the subset of known propositions and the number of true beliefs that subjects have about the connections between these statements. To take another example, knowing why Billy is crying might involve knowing that Joan hit Billy, this event hurt Billy and Billy usually cries when he is hurt. But the subject might know why Billy is crying by knowing that Billy has suffered from tissue damage, there are signals s sent to the brain that indicate tissue damage, which in turn activate neurons that are responsible for the release of tears and corresponding emotions. Thus, one can know why Billy crying by knowing different subsets within the explanation for why he is crying and grasping the interrelationships between the members of this set to different degrees.

Application of Analysis to Cases

Let see how this analysis helps explain what is going on in the cases depicted above. In both version B’s, the speakers truly ascribe knowledge to themselves of the relevant propositions in virtue of knowing why these propositions are true. In other

words, statements of the form ‘S knows that p ’ is true in these cases because S knows some members $m_1 \dots m_n$ of a set x that explains the truth of p and S has true beliefs about the explanatory, probabilistic and other relations within this subset and p . Why think that knowing that p amounts to knowing why here? In the types of high-standard cases above, the interlocutors have background knowledge, interests and/or stakes invested in the truth of p , which give rise to challenges and counter-possibilities that cannot be answered unless the putative knower has a firm grasp on a sufficient explanation for p .

In Montessori School B, the other researchers and parents have a vested interest in the claims made by the speaker, and background beliefs that allow them to raise certain types of alternatives that call for some explanation of the speaker’s knowledge claim. Knowing that the Montessori school system provides a better learning environment for independent learners, then, requires knowing that the Montessori School System provides a learning environment L , that L is conducive to the intellectual and social development of children who are I , and children who are independent learners are I , etc. It is by knowing a broader set of integrated statements that explain the truth of p and having true beliefs about the connections among these statements that the subject can truly be ascribed knowledge that p in such circumstances. This is in contrast to case A, where the subject was truly ascribed knowledge because of her testimonially justified belief that Montessori Schools are better for independent learners. Here, the conversational participants did not raise the types of doubts or challenges noted in B.

We can observe this disparity again in the PTSD cases. In PTSD B, the doctor’s colleagues are well-aware of the alternative diagnoses that might apply in the case of the war veteran. The counter-possibilities they raise, then, require knowledge of why it is

that the patient has PTSD. In order for the doctor to sustain his knowledge claim in the second case, he was required to know that the DSM listed mood swings, and other symptoms for diagnosis of PTSD, that the patient exhibited these classic symptoms, and that war veterans generally suffer from PTSD. It is by *knowing* these integrated statements and the explanatory and probabilistic connections therein that the doctor could truly ascribe knowledge to himself that the patient was suffering from PTSD.¹⁴⁵ This is unlike PTSD A, where the patient's knowledge attribution seemed appropriate despite not knowing the explanation underlying his diagnosis. His wife does not have the appropriate background beliefs to raise doubts that would require the knower (the soldier) to know why *p*.

Further, we can see how subjects featured in both case B's might know why to different degrees. In PTSD B, the doctor knows that the patient is suffering from PTSD because of his knowledge of the nature of the disease, including the criteria listed in the Diagnostic and Statistic Manual, the symptoms typically exhibited by war veterans who

¹⁴⁵ My analysis (though independently formed) bears a strong resemblance to Stephen Hetherington's recent proposal that propositional knowledge comes in degrees due to its identity with knowledge how. (Stephen Hetherington, "Knowing (how it is) that *p*: Degrees and Qualities of Knowledge," *Veritas* 50, No. 4 (December 2005): 129-162). That is, his claim is that knowing that *p* is identical to knowing how it is that *p*. Hetherington maintains that the latter comes in degrees and thus allows the former to do so. There are, however, important differences between his account and mine. First, Hetherington does not provide an argument for why we must identify the notions of knowing how and knowing that across all contexts. In my discussion, I have suggested that knowing that amounts to knowing why in some cases as dictated by the interests and concerns of the conversational participants in the context of utterance. In contrast, it is not clear at all why there is a two-way entailment in Hetherington's paper. Second, in some instances of propositional knowledge, it does not make sense to ask 'how it is that *p*'. For example, it seems at least a bit awkward to say, "How is it that the Montessori School System provides a better learning environment for independent learners?" However, it is perfectly felicitous to ask the following question: "Why do Montessori schools provide a better learning environment for independent learners?" The second observation might be more or less idiosyncratic. The first point remains, however: Hetherington has not provided any reason to identify 'knowing how it is that *p* with knowing that *p*. That is, in most cases, these two notions can easily come apart: "I know that the machine works, but I don't know how it is that the machine works" seems to be a perfectly legitimate statement, which suggests that knowing how it is that *p* is not clearly equivalent to knowing that *p*. Thus, he has not shown why contextualists or anyone else must subscribe to this identity.

have PTSD, that the symptoms displayed by the patient resemble those in both the categories above, etc. More succinctly, he knows this set of statements as answering the question ‘Why is the patient’s diagnosis PTSD?’ Suppose, however, that the doctor is training a resident, who has also closely studied PTSD. She too will know the criteria listed in the DSM, but may not have previous experience with war veterans. She too knows why the patient’s illness can be diagnosed as PTSD, as she knows some of the members of the explanation for the patient’s diagnosis. However, the doctor knows a greater set of statements that explains the truth of p by knowing the experiences of other war veterans, the manner in which such patients exhibit these symptoms, etc. That is, the doctor’s knowledge of why the patient has PTSD is superior to his student’s. Insofar as ‘knowing that p ’ is equivalent to knowing why p , in these cases, the doctor knows that p to a higher degree. Here is a concise presentation of my view that propositional knowledge ascriptions can come in degrees:

- P1. In some contexts, ‘S knows that p ’ iff S knows why p .
- P2. Knowing why p can come in degrees.
- C1. In some contexts, knowing that p can come in degrees.

Objections

What I have presented above is reason to believe that there are instances in which propositional knowledge ascriptions can come in degrees. These are cases in which knowledge-that is equivalent to knowing why. Following Schaffer, I believe that S knows why p just in case she knows e as the true answer to the question why p . But I further maintained that ‘knowing e as the true answer to why p ’ involves knowing statements that constitute a subset e of the set x which explains the truth of p , and having true beliefs about the interrelations between the members of e and p . That knowing-why

involves knowledge of a number of propositions and their relationships to one another allows knowledge-why to come in degrees. In the next section, I will argue that this has certain implications for the correct linguistic model for ‘knows that p ’. Before I do so, however, let me address a few worries regarding what I have said so far, which will also help me further clarify my view. The first objection will address my analysis of knowledge-why, and the second, my claim that subjects in the B cases are required to know why p in order for their propositional knowledge ascriptions to be true. The third objection will question my move from the variability of knowing explanatory statements to the variability of knowledge-why.

OBJECTION ONE: I have ultimately defined ‘knowing that’ in terms of knowing that:

The first worry with my account is that ‘knows that’ appears both in the analysans and the analysandum. That is, I’ve suggested that ‘knowing that p ’ is equivalent to knowing why p in the cases above. But, I’ve defined knowing why p in terms of knowing that $m_1, m_2 \dots m_n$, where m denotes individual members of the subset e and having true beliefs about the relationships between these members and p . How should we understand ‘knowing that’ as it appears in the definition of knowledge-why? This iteration of ‘knows that’, I want to suggest, will also be context-sensitive. That is, the standards for knowing each m will depend on contextual factors such as accepted background information, the interests and stakes of the conversational participants and the salient counter-possibilities that pertain to each proposition. Here is a concrete example: Let us call the proposition ‘The Montessori School System provides a better learning environment for independent learners’ p and one of its explanatory statements m

(i.e., a member of the set that explains the truth of p such as the proposition, ‘Other things being equal, schools that allow for social and intellectual maturation are better than those that do not’). The standards for ‘S knows that p ’ will be different than the standards for ‘S knows that m ’. While the former requires that the subject know a set of integrated statements as the true answer to why p , the latter will correspond to different truth conditions that will vary according to certain contextual factors such as what is accepted as background information and what is not, salient alternatives, and the interests and stakes of the conversational participants. “Knowing that m ” in case B will require a particular degree of epistemic strength, but given certain background assumptions shared by the audience and speaker, will probably not necessitate the standards of ‘knowing that p ’.

The standards of knowledge are then relativized to individual propositions. This will have the implication that ‘knows’ corresponds to different standards depending on the proposition in question. This is not an unsupported view. Daniel Halliday, for example, maintains that where knowledge of more than one proposition is being ascribed, the standards of knowledge ascriptions can vary.¹⁴⁶ Halliday also writes, however, that these propositions must be suitably unrelated. While the propositions in the cases discussed here are relevant to one another, I take it that it is possible that the standards for knowing p and knowing members of the set that explain p come apart due to differences in the contextual factors associated with them.

OBJECTION TWO: In both B cases, the speaker (*qua* subject) has a greater degree or different type of evidence/justification, not a different type of knowledge.

¹⁴⁶ Halliday, “Comparatives and Gradability,” 388.

One might argue that my analysis of the subject's knowledge in the cases above is mistaken. It is not that subjects in the B cases have a different type of knowledge, but rather that they possess a stronger degree of evidence or a unique type of justification. In addition to showing that my analysis of the cases is simply mistaken, this objection has a further implication: if the difference between the subjects in the two versions of the cases is the type or degree of justification/evidence possessed by the subject, then I cannot draw the conclusion that propositional knowledge ascriptions come in degrees. Why not? To understand why I cannot move from the degree variability of justification to the degree variability of knowledge, we have to revisit the debate between Stanley and Cohen. In Cohen's original analysis, the gradability and context-sensitivity of 'knows' piggy-backed on the gradability and context-sensitivity of justification. More precisely, Cohen argued that knowledge comes in degrees because one of its components, justification, comes in degrees.¹⁴⁷ However, Stanley later demonstrated that these aspects of justification cannot establish the gradability and context sensitivity of propositional knowledge ascriptions. It might, for example, be the case that justification is context-dependent, but "that the most demanding property denoted by 'is justified' relative to any context is one such that knowing a proposition requires standing in that most demanding justification relation to that proposition".¹⁴⁸ The context sensitive and gradable properties of 'is justified' need not transfer over to 'knows', since S's justification for *p* must reach a fixed standard or threshold to qualify as 'knowledge that *p*' in a given context. That is, while 'is justified' might be gradable, justification sufficient for knowledge (which is the more accurate component of knowledge) need not be. That 'is

¹⁴⁷ Cohen, "Contextualism, Skepticism and the Structure of Reasons," 60.

¹⁴⁸ Stanley, *Knowledge and Practical Interests*, 80.

justified' can come in greater and lower strengths, then, does not entail that knowledge comes in different degrees. This is why we cannot establish the gradability of 'knows' by appealing to the gradability of justification. Objection two, then, not only raises worries about the accuracy of my analysis in the cases presented above, but if it is right, also challenges the view that propositional knowledge ascriptions can, in fact, come in different strengths.

I want to argue, however, that the epistemic distinction between the subjects in A and B cases is not one of justification or evidence. One can have justification e for p , on certain conceptions of justification, even though $\sim e$. That is, if I seem to remember turning the stove off an hour ago and firmly believe that no one has turned it on since, then I am justified in believing that the stove is off. However, it turns out that I did not in fact turn the stove off. It doesn't seem right to say that I am completely unjustified in my belief that the stove is off. But in the cases above, the subject's knowledge that p requires the truth of the members within a set of statements that explain the truth of p . Therefore, the subjects in such cases do not simply possess a unique type or high degree of evidence/justification, since such an analysis does not require the truth of the explanatory statements.

One might object, however, that in my response, I am assuming a definition of justification that it not equivalent to justification sufficient for knowledge. We are only interested in the latter as PTSD and Montessori School Systems present cases in which the subject's justification for p classifies her belief that p as knowledge. My opponent then might interpret the B cases in the following way: S knows that p in such cases iff: p , S believes that p , and S's justification consists in the (partial) explanation for the truth of

p . It seems, then, that there is no reason, to identify knowing that p in such cases with a different form of knowledge, since the explanatory, probabilistic and other relations highlighted in my account can be absorbed into the requirements of justification that p .

Before I respond to this objection, let me point out the distinction between my proposal and the one cited in the objection. I maintain that statements of the form ‘S knows that p ’ in the cases above have certain unique truth conditions, ones which amount to knowing-why, a different form of knowledge. According to the objection, such statements also have unique truth-conditions, but these amount to differences in justification. It is not part of this view that there is a different form of knowledge involved in the Case B’s above. Suppose this is right. Suppose that Lily knows that the Montessori School System provides a better learning environment than others for independent learners because her belief is justified in a particular way: it involves part of the explanation for the truth of p . Such an account, however, will not yield the right intuitions in the cases above. Why not? Once the counter-possibilities are raised in Case B, Lily can say the following: “We know that the Montessori School system is better than others for independent learners, because we believe, though we do not know, that the Montessori School System provides a learning environment x , that x is conducive to the intellectual and social development of children who are y , and children who are independent learners are y ”. Such a defense of her knowledge claim appears suspect and it seems that audience members might be less inclined to grant Lily her ascription. Lily cannot felicitously claim to know that the Montessori School System in fact provides a superior learning environment for independent learners. The problem is that requiring that S’s justification for p include explanatory statements about the truth of p does not

impose the further constraint that S must know these statements. Only by knowing these explanatory statements can Lily truly be ascribed knowledge that p in the B cases discussed above.¹⁴⁹

On my account, the subjects in the B cases knows that p if and only if she knows why p . S knows why p just in case p , S believes that p and S knows a subset of members within a set that explains the truth of p and the interrelationships therein. S does not have a justified belief in a standalone proposition but rather has knowledge of a large informational chunk by knowing each proposition within that group and the explanatory, probabilistic, and other relationships within this group. It is knowledge of this entire body of information that allows S to know that p in the contexts described above. I am not, then, arguing, like Cohen, that the gradability of ‘knows’ piggybacks on one of its constituents. I am suggesting, rather that ‘knowing that p ’ in the types of high-standard scenarios described above is to have knowledge-why p . Knowledge-why can differ with respect to the size of the subset known and the interrelationships grasped. This allows ‘S knows that p ’ to come in different strengths, which vary according to S’s knowledge of why p .

OBJECTION THREE: Stanley’s charge can be applied to my account of Knowledge-why

¹⁴⁹ My objector might further insist that we can construe justification along the same lines. Specifically, S is justified in believing that p in such cases if and only if she knows why p , where she knows a set of integrated statements that contribute to the full explanation of p . I take it that the acceptability of this analysis would hinge on two possibilities: first, justification sufficient for knowledge requires S to know the propositional content of all the reasons that constitute her justification in all contexts or second, that such requirements on justification apply only in certain cases. The claim that justification consists in knowing a certain set of statements *either* within certain contexts or across all contexts is not an uncontroversial one. Since this is not a widely acknowledged feature of justification, the burden of proof will be on my opponent.

My opponent, might, however argue that my characterization of knowing-why is susceptible to Stanley's objection discussed earlier. That is, the degree variability of knowing-why p cannot piggy-back on the gradability of knowing the explanation for p . That is, it is possible to know larger or smaller subsets within the explanation for p . However, in a given context, the size of subset necessary for knowing-why will be fixed to a particular standard, which is not itself variable. Therefore, while knowledge of explanatory statements allows for degree variability, this does not establish the degree-variability of knowledge-why.

I do not think that my account is susceptible to Stanley's charge. This is because even if there is a standard that marks off the size of subset and number of true beliefs in a given context, further advances in knowledge of the explanatory statements carry over to differences in knowing why: If both Joe and Mary know a sufficient explanation for p , but Joe knows a greater number of explanatory statements p and grasps the interrelationships within these statements better than Mary, it seems right to say that Joe's knowledge of why p is better than Mary's, and difficult to deny this. It is right to say that the firefighter's knowledge of why the house burned down is better than mine precisely because she knows a larger number of statements within the set that explains the truth of p , and the explanatory and probabilistic relations within the set and p . Thus, the gradability of knowing why can be based on the subset of explanatory statements known by the subject.

IV. Linguistic Basis for the Context-Sensitivity of 'Knows'

Stanley, Halliday and Blome-Tillmann all reject the idea that 'knows that p ' can come in different degrees. They differ on whether or not this is a devastating objection to

the central thesis of contextualism. Halliday and Blome-Tillmann maintain that the non-gradability of ‘knows’ does not in fact count against contextualism. As discussed above, they maintain that contextualists can appeal to other, non-gradable linguistic categories to strengthen and support their claim, such as modified comparatives and adverbially modifiable verbs. The problem with such accounts is that they neglect the types of cases in which propositional knowledge *can* come in degrees. These are cases in which certain types of counter-possibilities are raised that require the putative knower to know why *p*, which, I argued can come in different strengths. Thus, we must reject the current trend of identifying ‘knows’ within a linguistic-category containing non-gradable expressions.

What then must we say about the linguistic basis for the context-sensitivity of ‘knows’? As Stanley correctly observes, propositional knowledge ascriptions cannot be modified and do not participate in comparative locutions. Statements of the following form are clearly ungrammatical:

8. ??The doctor knows that the patient is suffering from PTSD better than his patient.
9. ?? The doctor more knows that the patient is suffering from PTSD.

However, such awkward locutions do not conclusively show that ‘knows that *p*’ does not come in degrees. Indeed, there are other ways in which we might indicate different degrees of propositional knowledge. In the cases depicted here, the truth conditions of ‘S knows that *p*’ correspond to knowing an explanation for *p* or knowing why *p*.

Comparative locutions, then, will involve reference to such explanations:

10. Lily’s knowledge of why the Montessori school system is suitable for independent learners is better than her mother’s.
11. The doctor’s knowledge of why the patient suffers from PTSD is better than his student’s.

Lily's knowledge that the Montessori school system is better is based on her knowledge of why it is better. Thus, we can felicitously indicate different degrees of propositional knowledge by speaking about the subject's knowledge why. The same is true for the doctor's knowledge of his patient's diagnosis, where we refer to the doctor's knowledge of why the patient has PTSD to indicate the degree to which he knows that the patient has PTSD. This suggests that we identify 'knows' as a gradable, context-sensitive expression such as 'tall' and 'rich'.

One might, however, argue that there remains an important distinction between the latter adjectives and 'knows': while we can make direct comparisons using 'tall' and 'rich' ('very tall'; 'taller than'; 'very rich'; 'richer than'), we cannot do so with propositional knowledge ascriptions. Rather, we must resort to statements about knowing-why. I want to suggest, however, that this objection does not in fact preclude 'knows' as a gradable, context-sensitive expression. Why not? There are other concepts that we readily accept as gradable, but which do not receive direct comparative locutions. For example, take 'understanding' which, undoubtedly comes in degrees. The unit of understanding can either be a proposition or a larger, coherent body of information.¹⁵⁰ Often, when understanding is applied to a proposition, it is intended to serve as a type of hedging maneuver. Jonathan Kvanvig refers to this as the nonfactive use of understanding, and attributes these uses to misspeaking.¹⁵¹ This is what occurs when one believes and has good evidence for believing that something is the case, but the speaker intends to "cushion the force of a bald accusation":

¹⁵⁰ Some epistemologists writing on understanding have distinguished the units of understanding this way, including Catherin Elgin, "Understanding and the Facts," *Philosophical Studies* 132 (January 2007): 33-42.

¹⁵¹ Some theorists allow that the proper use of understanding may be non-factive. See Catherine Elgin, *Considered Judgments* (Princeton: Princeton University Press, 1996). The dispute about the factive versus non-factive nature of understanding will not be relevant here.

1. I understand that you plan to join us this in the spring.
2. I understand that you are angry with me.
3. The trains have resumed their regular schedules, I understand.

In (1)-(3), ‘understand’ is used as a hedging maneuver, to indicate that the speaker is not in fact certain about the relevant proposition. Both Kvanvig and Catherine Elgin note, however, that these ‘moderate’ uses are not theoretically or epistemologically relevant, as they do not indicate cognitive success or achievement.

The second sense of understanding is marked by two features: It involves the grasping of a large body of information and the way in which these propositions relate to one another.¹⁵² Stephen Grimm cites Linda Zagzebski as maintaining that “understanding is fundamentally a matter of grasping how various pieces of information relate to one another; it is a matter of making connections among them, or seeing how they hang together”, as in (5):¹⁵³

5. The students understand the theory of evolution.

The subjects denoted in (5) presumably understand the theses of survival of the fittest and natural selection, facts about gene and environment interaction, the potential for variation in inherited traits, etc. Understanding of this type seems, then, to apply most directly not to single propositions, but large bodies of information. Many have noted, however, that such understanding can also apply to propositions. How so? Elgin writes that propositional understanding in this sense derives from one’s understanding of a larger body of data within which the proposition in question fits. While “understanding is

¹⁵² As such, this form of understanding is importantly similar to what I have called knowledge-why. However, I am not here interested in endorsing the view that the two are identical. I think that there are significant differences between the two notions: see Elgin, “Understanding and the facts” and Jonathan Kvanvig, *The Pursuit of Knowledge and the Value of Understanding* (Cambridge: Cambridge University Press, 2003).

¹⁵³ Stephen Grimm, “Is Understanding a Species of Knowledge,” *British Journal of the Philosophy of Science* 57 (September 2006): 517.

primarily a cognitive relation to a fairly comprehensive, coherent body of information”, “the understanding encapsulated in individual propositions derives from an understanding of larger bodies of information that include those propositions”.¹⁵⁴ It amounts to grasping how a particular “fact fits into, contributes to and is justified by reference to a more comprehensive understanding that embeds it”.¹⁵⁵ To illustrate, she offers the following interpretation of “I understand that Athens defeated Persia in the battle of Marathon”:

My understanding... depends in a suitable way on my overall understanding of such matters as the course of the battle, the strategies and the tactics of two armies, and the history of relations between Athens and Persia... That is, the proposition derives its epistemological status from a suitably unified, integrated, coherent body of information.. [which] is closely connected to explanation¹⁵⁶

The ‘epistemological status’ of understanding that Athens defeated Persia is founded on my grasping a larger, coherent body of information, and how this particular fact fits into it.

Elgin further notes that there are at least three¹⁵⁷ dimensions along which propositional understanding can vary: breadth, depth and significance. A professor’s understanding that the Athenians defeated the Persians might be based on a more comprehensive understanding of Greek history. His understanding might be deeper in that his web of belief is more tightly woven and contains more propositions, etc. We might then speak of the professor having ‘greater understanding’ of the Athenian victory. What is important to note for our purposes, however, is the way in which we indicate different degrees of understanding a proposition in this second sense. While we allow

¹⁵⁴ Elgin, “Understanding and the Facts,” 35.

¹⁵⁵ Ibid.

¹⁵⁶ Ibid.

¹⁵⁷ Elgin in fact cites a further possible dimension along which understanding may be measured, but this commits us to a non-factive account of understanding.

that individuals can have different degrees of propositional understanding it is not clear that we can directly modify or compare ascriptions of propositional understanding.¹⁵⁸

12. ?? The professor more understood that the Athenians beat the Persians than his students.
13. ??The professor's understanding that the Athenians beat the Persians is greater than his students.

We cannot felicitously compare degrees of propositional understanding in this way even though we accept that propositional understanding comes in different degrees. We can however indicate varying degrees of propositional understanding by appealing to the larger body of information:

14. The professor's understanding of the Athenian victory is greater than his students.
- and
15. The professor understood the Athenian victory better than his students.

The moral is that even though propositional understanding can come in degrees, it seems that we can only indirectly indicate different degrees of propositional understanding. This suggests that the fact that we cannot make direct comparisons of propositional knowledge does not show that 'knows' is not a gradable expression. We can then sustain Cohen's original analogy between 'knows' and gradable, context-sensitive expressions, which include 'tall', 'rich', 'expensive' and 'understanding'.¹⁵⁹

¹⁵⁸ More precisely, we cannot directly modify or compare ascriptions of propositional understanding in the more robust, second sense of understanding discussed here.

¹⁵⁹ It might be argued that coming in different strengths is a property unique to adjectives, and that therefore, and on these grounds alone, we must reject the identification of 'knows' with 'tall' and 'rich'. However, the rationale for this view is not obvious. Stanley's linguistic tests display but one way of comparing different degrees of a particular property. Furthermore, linguists have consistently argued that gradability can equally apply to nouns, verbs, etc. Kennedy and McNally, for example, write "gradability is a fundamentally important semantic property, whose influence extends beyond adjectives to other lexical categories" (Christopher Kennedy and Louise McNally, "Scale Structure, Degree Modification, and the Semantics of Gradable Predicates," *Language* 81, No. 2 (2005), 348.

Such an analogy adequately captures the gradable, context-sensitive semantics of ‘knows that p ’.

V. Conclusion

I have argued here that current linguistic analyses of ‘knows’ are not sustainable because they do not accommodate the degree variability of propositional knowledge ascriptions. I argue for this by showing that there are certain cases in which ‘S knows that p ’ is true if and only if S knows why p . Knowing why p involves knowing a subset e within a set x that explains the truth of p , and grasping the interrelationships between e and p . These two features allow knowing-why to come in degrees. Insofar as ‘knows that p ’ is equivalent to knowing why in such cases, propositional knowledge ascriptions can come in degrees. Therefore, it is a mistake to identify ‘knows’ with a non-gradable linguistic category. Rather, ‘knows’ belongs to the category of gradable context-sensitive expressions that include ‘tall’, ‘rich’ and ‘flat’.

Chapter Six Conclusion

My goals in the preceding chapters were two-fold: first, I aimed to show that the general invariantist strategy of invoking pragmatic content to explain ordinary linguistic behavior and judgments with respect to ‘knows’ is ultimately unsuccessful. Without a proper explanation of our use of ‘knows’, one that clarifies why we *believe* that ‘knows’ applies in a given context when it strictly does not, the entire invariantist enterprise fails. Second, I wanted to show *how* ‘knows’ is context-sensitive. In the last chapter, I argued that ‘knows’ is most accurately identified within the category of gradable, context-sensitive expressions such as ‘tall’, ‘rich’ and ‘flat’.

I. Arguments and Implications

How did I argue for these claims and, if I am right, what are their implications? What consequences does this have for the ongoing debate between classical invariantists and epistemic contextualists? I have in a sense, contributed to, and extended, DeRose’s (and others’) war on defective warranted assertability maneuvers (WAMs). How so? DeRose decries the use of WAMs by invariantists by arguing that they fail to meet certain conditions which he believes are necessary for a WAM to be successful. What I show, however, is that despite meeting these conditions, invariantist accounts face more profound problems associated with the practice of identifying unfavorable linguistic intuitions with pragmatic content.

WAMs are employed in two ways: They are used to explain why an assertion seems false even though it is in fact true by showing that the utterance is improper or

conversationally awkward, or they may be used to explain why an assertion seems true even though it is in fact false by showing that it is warranted by previous utterances or the general direction of the conversation. The underlying idea of a WAM, then, is that speakers confuse the propriety/impropriety of an utterance with its truth/falsity. This allows invariantists to explain our allegedly mistaken intuitions about the application of 'knows'. DeRose (1999) presents three conditions that must be fulfilled in order for a WAM to successfully show that speakers are mistaking truth conditions for the appropriateness of an utterance.¹⁶⁰ First, there must be an intuition that an assertion is false, where the opposite assertion also seems false, which suggests that one of these judgments must be incorrect and therefore explained away. Second, he writes that this faulty intuition must be explained away by appeal to the generation of a false implicature, rather than a truth. When something false is conveyed by the utterance in question, DeRose maintains, "it is not surprising that we might mistake that for the assertion's itself being false".¹⁶¹ Lastly, he maintains that the maneuver must appeal to a general conversational rule to explain the generation of the implicature, rather than invoke a new or unique rule that attaches only to the term or phrase in question. He concludes that since WAMs about 'knows' do not typically meet such criteria, they are ineffective. It turns out, however, that there are accounts that satisfy DeRose's conditions, but are nevertheless unviable for other reasons. The discussion here was intended to dissect such accounts and point out their crippling flaws.

In the second chapter, I began by examining Rysiew's defense of moderate invariantism in which he employs a WAM that satisfies DeRose's criteria, but is

¹⁶⁰ DeRose, "Contextualism: An Explanation and Defense".

¹⁶¹ *Ibid.*, 17.

nevertheless unsuccessful. DeRose argues at the outset that invariantist WAMs will fail to meet the first condition because in low standard cases, it seems wrong to deny that the subject has knowledge and right to ascribe knowledge to her, whereas in high standard cases, it seems true that the subject does not have knowledge and false that she does. Hence, it is not the case that we have contradictory intuitions, some of which must be explained away. It is not clear, however, that DeRose's brief remarks decide the case. Just as DeRose allows that we can be ambivalent about our knowledge in radically skeptical contexts¹⁶² one can argue that there are indeed conflicting intuitions about DeRose's bank cases.¹⁶³ On the one hand, it seems wrong for the husband to assert, "I know that the bank will be open tomorrow" when his wife raises the counter-possibility that the bank might have changed its hours. On the other hand, the counter-possibility that the wife has raised does not sound terribly realistic or at least characteristic of financial institutions. Banks do not up and change their hours on a whim. Moreover, if the couple has used this bank in the past and is familiar with its policies, then they should be able to know if this bank, in particular, tends to shift its hours. We might intuit, then on some level that the husband says something false when he asserts that he does not know that the bank will be open the next day. Both these intuitions cannot be right. That is, it cannot be that 'I know that p ' and 'I don't know that p ' are both false. One of these intuitions must be explained away. Hence, DeRose's first condition for a successful WAM—that we judge an assertion and its opposite to be false—is met.

¹⁶² DeRose, "Solving the Skeptical Problem," 208, 209.

¹⁶³ Tim Black articulates precisely how we might entertain opposing intuitions when it comes to the skeptical argument (discussed in an earlier chapter): "You feel each hand with the other, and they feel just as they did on countless occasions when you clearly seemed to know that you have hands. Your fingers smell faintly of the banana you ate for lunch. You touch the desk in front of you..." in "A Warranted-assertability Defense of a Moorean Response to Skepticism," *Acta Analytica* 23 (2008): 194.

Rysiew argues that the first of these intuitions is incorrect. That is, strictly speaking, the speaker would indeed have asserted something true by saying, “I know that the bank will be open tomorrow”. We might nevertheless intuit that this statement is false because it generates the *false implicature* that the husband can rule out the salient counter-possibility that his wife has just mentioned. Rysiew’s WAM then also satisfies DeRose’s second criterion. Further, Rysiew writes that such an implicature is generated due to Grice’s conversational maxim of Relation, specifically, “Be Relevant!” He illustrates how we might apply Grice’s maxim to what the speaker has said:

He has just said ‘I (guess I) don’t know that the bank will be open tomorrow’. And he has said this in response to my raising a doubt as to whether he can really be so sure... Presumably (on the assumption that he’s conforming to [the Cooperative Principle], he wouldn’t have said what he has unless he thought there were possibilities incompatible with the bank’s being open tomorrow—specifically, that it has recently changed its hours—that he could not rule out... On the assumption then, that his conversational contribution is to the point and has been made in light of what *I’ve* just said, *that* must be what he intends to communicate—viz., that he can’t rule out the possibility that the bank has recently changed its hours.¹⁶⁴

If, instead, the speaker were to claim that he did know that the bank would change its hours, we would mistakenly conclude that he *can* rule out the possibility that the bank has recently changed its hours, given Grice’s rule.

Despite meeting DeRose’s three conditions for a successful WAM, Rysiew’s pragmatic explanation faces a different problem. One of the main assumptions underlying Rysiew’s (and other’s) use of a WAM is that what is communicated or imparted by ‘I know (don’t know) that *p*’ diverges, in some instances, from the proposition it literally expresses. This conveyed content then determines whether or not the utterance is warranted or conversationally appropriate. What I tried to show in the second chapter is that it is not clear that the speaker’s communicative intentions can in

¹⁶⁴ Rysiew, 491.

fact be detached from the proposition literally expressed by her statement. I argued for this by showing that when a knowledge denial is followed up by a question about what the speaker actually intended to say, there is no difference between the proposition literally expressed (viz., I don't know that p) and the content that she intended to convey (viz., I don't know that p).

Here is a brief summary of my discussion: Rysiew maintains that a speaker says 'I do not know that p ' not to genuinely deny that she has knowledge of p , but to assert that she cannot eliminate the counter-possibilities that have been raised in the conversation. (This occurs in cases where remote counter-possibilities to p have surfaced.) I show, however, that if we were to press the speaker to determine what it was that she intended to communicate, she would reiterate that she did not know p . Since the speaker's intended meaning is not distinct from the proposition literally expressed by the utterance, Rysiew cannot use the notion of pragmatic (or imparted, or communicated) content to explain away intuitions that are unfavorable to his view. I showed, in other words, that the assumption that he needs to underwrite his WAM is unfounded, thereby extending DeRose's attack on ineffective WAMs by challenging one of the critical presuppositions.

In the third chapter, I addressed Peter Unger's defense of skeptical invariantism. DeRose maintains that Unger's account is actually on better footing than other forms of invariantism as Unger does not treat 'knows' as a unique phenomenon in the language, but rather identifies it with a fairly broad set of expressions called 'absolute terms'. These expressions correspond to semantically stable meanings, but may nevertheless appear to vary from one context to another. In addition to 'knows', the set includes 'certain', 'flat', 'empty', etc. According to Unger, a surface does not literally qualify as

'flat' unless no other object is flatter than it. But we may describe the surface as 'flat' to indicate that it is flat enough for some purpose or the other. In the same way, one knows that p only if she is absolutely certain that p , such that "if someone is certain of something then there never is anything of which he or anyone else is more certain".¹⁶⁵ However, Unger allows that we might ascribe knowledge to a subject who does not meet this condition in order to convey that his/her epistemic position is good enough for the situation at hand. While DeRose commends Unger's attempt to group 'knows' with other terms in the language, DeRose points out other unpleasant implications for his account. DeRose cites both the implausibly stringent truth conditions for common terms, as well as the widespread and systematic falsehoods that results from such excessive demands. These features, according to DeRose, render Unger's skeptical invariantism unattractive.

The problems of unreasonable standards and widespread falsehoods, however, are tempered by Unger's explanation of what we ordinarily communicate, impart or understand by these terms. Presumably, this is the very function of a WAM, which is to explain why we might think that an assertion is true when it is literally false, or alternatively, why we might believe an assertion is false when it is literally true. Admittedly, Unger's account does have the unfortunate consequence of generating false assertions and positing unreasonable standards in some instances, but so does any invariantist account. Low standard and moderate invariantists, for example, appear to demand too little in certain cases (perhaps where the stakes are high) and have to address the issue of systematic falsehoods in these instances. I am inclined to believe, then, that problems of excessive demands and systematic falsehoods do not count uniquely against

¹⁶⁵ Unger, *Ignorance* 67.

Unger's theory. However, his account is not without difficulties. In my discussion, I focused on Unger's distinction between absolute and relative terms and argued that it is not sufficiently supported. Without an accurate and reliable method (or criterion) to distinguish semantically context sensitive terms from expressions that merely appear context sensitive, there is no reason to believe that 'knows' and other terms such as 'flat' fall into the second, rather than the first, category. Unger's paraphrase and emphasis tests, which are intended to distinguish absolute (context-insensitive) from relative (context sensitive) terms do not yield the required results. Unger's account, then, is ultimately unviable not because it fails DeRose's conditions for a successful WAM or results in widespread and systematic falsehoods, but because he has failed to provide a viable means of distinguishing absolute terms, those that are semantically stable but pragmatically variable, from other words. Hence, his pragmatic explanation of our ordinary use of 'knows' cannot be sustained.

In the next chapter, I considered an argument for low-standard invariantism. Again, low standard invariantists have the resources to meet DeRose's conditions for a successful WAM and thus provide an alternative analysis for linguistic data that appears to support contextualism. Tim Black, who defends a neo-Moorean invariantism presents one such example. Black argues that our judgments about knowledge in skeptical contexts can vacillate. We might believe that it is false that we know that we have hands when the skeptic has raised the possibility that we might be BIVs. But at the same time, given the way our hands feel, smell, etc. to us, it seems false to claim that we do not know that we have hands. Again, DeRose's first condition for a successful WAM is met.

Black argues that the intuition that we don't know that we have hands is incorrect. That is, even in skeptical contexts, we know that we have hands. He accounts for this mistaken intuition in a manner different from Rysiew. He rejects Rysiew's contention that ascribing knowledge in skeptical (or any high standard) scenarios flouts Grice's maxim of Relation (Be Relevant!).¹⁶⁶ Contra Rysiew, he maintains that knowledge attributions in high standard contexts are relevant even when they impart falsehoods. Black argues instead that asserting "I know that I have hands" in a radically skeptical context violates the second maxim of Quantity: Do not make your contribution more informative than is required. He calls this a violation of the maxim of strength. He writes that claiming to know that p in a skeptical context generates the false implicature that the speaker's epistemic position is such that she can eliminate the counter-possibility that the skeptic has raised (e.g., she is a BIV programmed to perceive ordinary sensations). This assertion is too strong and is therefore unwarranted. Black proceeds to explain that "I know that I have hands" is stronger than other assertions in the sense that it is more informative. Hence, the statement violates the maxim of strength and this explains why we judge that it would be false in a skeptical context.

Even though Black's explanation satisfies DeRose's conditions for a viable WAM, there are additional worries associated with his reliance on pragmatic content. I was specifically concerned with concessive knowledge attributions, which claim, on the one hand, that a subject has knowledge that p , but on the other, expresses a need for further evidence. Originally articulated by Cohen, the worry is this: if the standards of knowledge are perpetually low, then in situations where it is prudent to gather additional

¹⁶⁶ Black, "AWarranted-assertability Defense of a Moorean Response to Skepticism".

evidence, we would have to utter statements such as, “Well, we know that p , but we have to check further” and “Billy knows that p , but we have to get more evidence” which, unarguably sound awkward. It appears that we cannot cancel the entailment from ‘I know that p ’ to ‘I need not check further’. This suggests that a view that sets the standards of knowledge low across all contexts cannot be correct. Black and Murphy have argued, however, that such statements are elliptical and when fully articulated, their oddity disappears. For example, it is perfectly alright to say, “I know that p , but I need more evidence to publish in journal X”. Given that the implication from ‘I know that p ’ to ‘I need not check further’ *can* be cancelled, Black and Murphy argue that these statements are only *pragmatically related*, as per Grice’s cancellability criterion. Since ‘I know that p ’ only pragmatically implies ‘I need not check further’, their apparent oddity does not count against low standard invariantism.

In the fourth chapter, however, I defended Cohen’s thesis that ‘I know that p ’ semantically entails ‘I need not check further’ by showing that the cases that Black and Murphy use to support their view are actually cases of loose talk. Grice’s criterion of cancellability is not applicable in cases of loose talk. Therefore, the fact that we can cancel the implication from ‘I know that p ’ to ‘I need not check further’ in Black and Murphy’s cases do not conclusively establish that the two clauses are only pragmatically, rather than semantically, related. What this shows is that Black and Murphy’s attempt to defend low standard invariantism by appealing to the pragmatic implication of ‘I know that p ’ is unsuccessful.

If invariantists cannot provide a viable explanation of our variable uses of ‘knows’—that is, why ‘knows’ *appears* to be very demanding in some cases but less so

in others—it seems that we ought to accept the thesis that ‘knows’ is semantically context-sensitive. But these arguments alone cannot support a contextualist analysis of ‘knows’. Rather, we need some understanding of the way in which ‘knows’ depends on the context. Further if no other expression functions in the same way as ‘knows’, this raises doubts about the semantic theory of epistemic contextualism. Contextualists, then, need to identify ‘knows’ within a broader linguistic category to strengthen and support their semantic thesis. To this end, Cohen identified ‘knows’ with gradable context-sensitive expressions such as ‘tall’ and ‘flat’. However, both opponents and proponents of contextualism have taken issue with this analogy and argued that ‘knows’ is not a gradable expression and is therefore more accurately identified with linguistic categories that do not contain gradable expressions. In the fifth chapter, I argue that propositional knowledge ascriptions can come in degrees in certain contexts. This conclusion is based on two premises: First, there are high-standard scenarios in which it is true that the subject knows that p if and only if she knows why p . Second, knowledge-why comes in degrees. Therefore, I conclude that propositional knowledge ascriptions (in these cases) admit of degree variability. For these reasons, I contend that we should sustain Cohen’s original claim that ‘knows’ must be classified as a gradable context-sensitive expression such as ‘tall’ and ‘flat’, though not for the reasons Cohen originally provided.

II. Further Worries

The failure of invariantist accounts to explain our use of ‘knows’ accompanied by a linguistically supported contextualist theory recommends the latter as the correct analysis. However, there remain certain important, but unanswered questions about

epistemic contextualism. Contextualists maintain that the standard of knowledge in any given context depends on factors such as the purposes and interests of conversational participants as well as doubts and challenges that have been raised in the course of conversation. But what happens when participants disagree about the correct standard of knowledge; where their assertions and other conversational factors do not seem to converge to form a single standard? Suppose, for example, that after a speaker has claimed to know an ordinary proposition, her interlocutor raises a fairly remote counter-possibility. Instead of acknowledging that she cannot eliminate this counter-possibility and conceding that she does not know the proposition, the first speaker sticks to her guns, dismissing the skeptic's challenge and insisting that she *does* know the proposition in question. What should the contextualist say in such a situation?

The problem has not gone undetected by theorists. While most contextualists have examined such situations to some degree, Keith DeRose has discussed the topic at length in "Single Scoreboard Semantics".¹⁶⁷ For a framework that both offers some insight and lends helpful concepts, we can begin by turning to David Lewis' "Scorekeeping in a Language Game".¹⁶⁸ Here, Lewis introduces what he calls the 'conversational score' which involves sets of objects that function to determine the truth conditions or acceptability of utterances in the conversation at any given time. These include sets of presuppositions, degrees of vagueness, relevant/ignored possibilities, etc. Each evolves according to the inputs of the conversational participants and is typically governed by the rule of accommodation. According to this rule, the levels or boundaries of each component of the score shifts in order to make an utterance true or otherwise

¹⁶⁷ Keith DeRose, "Single Scoreboard Semantics," *Philosophical Studies* 119, no. 1-2 (May 2004): 1-21.

¹⁶⁸ David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8 (August 1979): 339-359.

acceptable. To illustrate, Lewis offers the following example: The commonsensical epistemologist says, “I know the cat is in the carton - there he is before my eyes - I just can’t be wrong about that!” The skeptic replies, “You might be the victim of a deceiving demon”. Lewis maintains that the latter “brings into consideration possibilities hitherto ignored, else what he says would be false. The boundary shifts outward so that what he says is true. Once the boundary is shifted, the commonsensical epistemologist must concede defeat”.¹⁶⁹ That is, the rule of accommodation dictates that the commonsensical epistemologist must now attend to the possibility that a demon might be misleading him, and since his evidence cannot eliminate this alternative, he must now deny that he knows that the cat is in the carton.

Lewis adds that participants may either conform to the directives or steer components of the shared score in certain directions. While he recognizes that “there may be conflict, as when each of two debaters tries to get his opponent to grant him—to join with him in presupposing—parts of his case, and to give away parts of the contrary case”,¹⁷⁰ Lewis does not in fact discuss such instances. Rather he focuses on cases where the participants cooperate and strictly follow the rule of accommodation. The type of case that I am interested in, however, is one in which this rule is not applied, as one of the participants resists the change in the conversational score that the other is initiating.

In his discussion, DeRose surveys various possible analyses of such a situation before presenting his favored view. He employs two gauges to evaluate each analysis: first, does the analysis show that the two speakers are engaged in a debate? That is, “the skeptic and her opponent do take themselves to be contradicting one another... each is

¹⁶⁹ Lewis, “Scorekeeping in a Language Game,” 355.

¹⁷⁰ Lewis, 345.

publicly indicating that they are (or at least mean to be) contradicting one another”.¹⁷¹

Second, does it sufficiently accommodate the intuition that the truth conditions of a speaker’s knowledge claims should match the epistemic standards that she herself is indicating by her conversational maneuvers? Speakers, it would seem, should be able to assign a content to ‘knows’ that corresponds to their intentions. DeRose begins by considering a view in which there is no shared scoreboard, but only personal, multiple scoreboards. That is, each speaker’s personally indicated content (the epistemic standards that her conversational maneuvers have a tendency to put into place) contributes to a personal scoreboard which lists the truth conditions for that speaker’s uses of ‘knows’. On this view, both the skeptic’s and the nonskeptic’s utterances are true, since the truth conditions for each party’s assertions are determined by her own inputs. The worry with this position is that it does not present the speakers as contradicting one another. Rather, it seems that they are talking at cross-purposes.

DeRose favors a single scoreboard view, in which scorekeeping occurs on one level and the truth conditions of both speakers’ assertions depend on the score registered on this single scoreboard. This view entails that one of three things must happen: (1) the skeptic’s high standards prevail, (2) the non-skeptic’s low standards persist, (3) both speakers contribute to the resulting standards. If we allow that either the skeptic or her opponent ‘wins’, the conversation, then the second of the two intuitions is not accommodated. That is, if it is the skeptic’s standards that govern the conversation, then her opponent’s knowledge claims will not match her own personally indicated content. If, on the other hand, the non-skeptic’s low standards determine the truth conditions, then

¹⁷¹ DeRose “Scoreboard,” 3.

the skeptic's assertions will not match her personally indicated content. The third option can be further subdivided. First, DeRose discusses 'reasonableness' views, according to which the conversational score depends on what would be reasonable given the speakers' interests and purposes. The score may be determined either irrespective of the personally indicated content of the speakers (the 'pure reasonableness' view) or based on the personally indicated content of whoever is being the most reasonable (the 'binding arbitration' model). These views face two major difficulties. First, neither model accommodates the intuition that the truth conditions of a speaker's utterances should match the standards that she has in mind, as either one or both of their personally indicated contents do not contribute to the final truth conditions for assertions about knowledge. Second, there will be tough questions about what counts as 'reasonable'.

There is, however, another option available in this last category, what DeRose calls the Gap view. This model posits the following truth conditions for such a conversation: 'S knows that p ' is true just in case S meets the standards set by both parties and 'S does not know that p ' is true if S fails to meet both participants' standards. However, if S's epistemic status with regard to p meets only one of their standards, then both sentences ('S knows that p ' and 'S does not know that p ') will be truth-value-less. DeRose favors this view over the others for three reasons. First, it seems to capture the sense that the speakers are contradicting one another. How so? DeRose claims that "[the skeptic's] claim is true in the same range of cases where [her opponent's] claim is false", while "the [skeptic's] claim is false in the same range of cases where [her opponent's] claim is true".¹⁷² Second, that in some cases, the speakers' assertions will have no truth

¹⁷² DeRose, "Single Scoreboard Semantics" 16.

value illustrates the fact that this is not an ordinary conversation in which the participants share each other's assumptions and accept each other's challenges. And third, the view does a decent job of respecting the intuition that one's assertions should correspond to the truth conditions she indicates. While in a certain range of cases, the truth conditions of their assertions will not match the speakers' personally indicated contents, there will nevertheless be a 'weak' equivalence such that in all the cases where the speaker's assertion *does* have a truth value, it matches the truth value it would have had if the standards were determined by her own conversational maneuvers. In DeRose's words, "wherever it is both the case that your claim has a truth value, and that it would have had a truth value if it bore the content you personally indicate, the truth value it has is the same as the truth value it would have had if it bore the content you personally indicate".

Examining the Gap view more closely, however, it seems that that DeRose has not offered an analysis that genuinely respects both intuitions. Here is a table of the Gap View:

Speakers	Sentences	S meets skeptic's standards/S meets both standards	S meets Common Sensian's standards	S meets neither party's standards
Common Sensian	S knows that p	T	U	F
Skeptic	S ~know that p	F	U	T

The last three columns present the different epistemic positions that S may occupy with respect to proposition p . In the first instance, the subject has met the skeptic's elevated standards for knowledge. Since the skeptic's standards will be more demanding than the non-skeptic's, this will entail that the subject has also satisfied the non-skeptic's standards. In such a case, it will be true that S knows that p and false that she does not. In the next column, the subject meets the non-skeptic's standards, but not the skeptic's standards. According to the Gap View, neither knowledge ascriptions nor denials will have truth-values in these cases. In the final column, S has not met either party's standards, and thus 'S does not know that p ' is true and 'S knows that p ' is false. DeRose cites the first and third of these columns as evidence that the view accommodates the intuition that the two speakers are contradicting one another. But it appears that the column that actually pertains to our discussion is the second. Why? In the kind of conversation we're considering, a speaker begins by claiming to know a proposition p on the basis of some evidence e . Her interlocutor objects that her evidence does not allow her to rule out some remote counter-possibility—perhaps that an evil omnipotent being has caused her to have false beliefs. The first speaker dismisses this alternative, saying something to the effect of, "Aw, come on!" or "Don't be ridiculous!" In this case, it is clear that we are not assessing a subject who fails to meet the non-skeptic's standards. After all, the non-skeptic begins by ascribing knowledge to herself. Furthermore, it is not the case that the subject has acquired the standards posited by the skeptic. The non-skeptic is, after all, insisting that the subject knows despite not meeting such standards. This is where the disagreement lies; that is, the skeptic and her opponent are contradicting one another in this situation. The Gap View, as is clear from the table, does

not recognize that a contradiction has occurred under these circumstances. Instead of yielding contradictory truth values in these cases, DeRose's preferred analysis has it that neither speaker's claims has truth values. Hence, the Gap view does not properly accommodate the intuition of contradiction.¹⁷³

What, then, are our options? We might step back and re-examine the two intuitions that DeRose attempts to preserve—that the speakers are contradicting one another and that the truth conditions of one's knowledge assertions should match her own personally indicated content. Can we satisfactorily accommodate both? It seems that if we were to favor the latter by allowing each participant to determine the truth conditions of her own statements, then we would have to forfeit the sense that the speakers are contradicting each other. Rather, they would be talking about different things entirely. On the other hand, if we take the intuition of contradiction at face value, we would have to reject the belief that the truth conditions of one's use of 'knows' should match the standards she herself envisions. Verena Gottschling has proposed that we abandon the intuition that the two speakers are in fact contradicting each other in such a conversation. He notes that for a genuine contradiction 'S knows that *p*' uttered by either the skeptic or the non-skeptic must have identical content, but different truth values. Gottschling further writes that contextualists claim that the truth values of knowledge ascriptions vary because the *content* of such statements differ. He then applies this core contextualist thesis to the situation we are considering and concludes "there is no contradiction...."

¹⁷³ Verena Gottschling (in "Keeping the Conversational Score: Constraints for an Optimal Contextualist Answer," *Erkenntnis* 61, no. 2-3 (November 2004): 303.) maintains that this view is not unlike a view that DeRose rejects, which he calls the 'exploding' scoreboard, where there is no correct score. Such a view seems to clearly violate the intuition of contradiction and Gottschling argues, and for this reason, so does the Gap View. Gottschling writes that "a shared scoreboard without entries in the relevant arrays, seems only a marginal improvement to no scoreboards in these cases"

Apart from containing the same words, the proposition semantically expressed by [the non-skeptic's] utterance is not the denial of the proposition semantically expressed by [the skeptic's] utterance".¹⁷⁴ The intuition that we should be trying to preserve is that each speaker *believes* that the other is contradicting her statement.

Gottschling's understanding and application of the contextualist semantic thesis, I believe, is incomplete not helpful to the current discussion. While contextualists maintain that 'S knows that *p*' when uttered in a low standard scenario does not contradict 'S does not know that *p*' when uttered in a high standard context (as each assertion bears content that pertains to the speaker's respective interests, purposes and conversations) that model does not apply in a situation where two people who are party to the same conversation and context disagree about the standards of knowledge. In this case, I believe that there is a genuine dispute about what qualifies as 'knowing'. In other words, the clash is not centered on whether the subject has met a particular standard for knowledge, but rather, which standard applies. Foregoing the intuition of contradiction neither advances the discussion nor does it properly apply the thesis of contextualism. What, then, should the contextualist say?

I think a better approach lies in explaining away the intuition that the truth conditions of one's assertions about knowledge should always match her personally indicated content. I suspect that while we might share this intuition to some degree, it does not seem like it will persist in all conversations that feature disagreement. For example, when I tell my husband, "I know that I will be home by 5 if I leave at 4.45", this statement corresponds directly to the standards that I envision for knowledge. Perhaps

¹⁷⁴ Gottschling, 309.

evidence for this claim comes from my past experience of making the same trip on several different occasions at a similar time of day. I take it that such evidence is sufficient for my belief that I will be home by 5 to qualify as knowledge. This much seems right. But the connection between the actual truth conditions (standards) for my knowledge ascription and the truth conditions that I might have in mind, I believe, is disrupted when reasonable doubts or challenges have emerged that serve to elevate the standard for knowledge.

Suppose my husband responds, “Are you sure that the parade isn’t today? If it is, the road will be closed by 4, so you might want to leave earlier to get home by 5”. Such counter-possibilities raise the standard of knowledge by requiring a greater degree of evidence in order to know that I will be home by 5. However, suppose I do not take this possibility into consideration and simply persist: “I’m not sure about the parade, but I know I’ll be home by 5”. Such a response is awkward. It does not seem like I can felicitously maintain my knowledge ascription, and the second conjunct appears false because I have not acknowledged that the standard of knowledge has been elevated. There is a detachment, then, between my personally indicated standards and the actual standards of the knowledge-ascription. That is, without changing something in my arsenal of evidence to address my husband’s concern (something such as- I read in the newspaper that they rescheduled the parade or I know another route that will get me there by 5, etc..) it does not seem that my knowledge claims should continue to match the truth conditions I originally had in mind. What this suggests is that although it seems right that one’s knowledge claims should correspond to the standards that she has in mind, truth

conditions should also be sensitive to the direction of the conversation as well as the purposes and interests of the conversational participants.

Where does this leave us? I believe that debunking, or rather, fine-tuning, this intuition opens the way for the correct analysis of such cases. The correct view, then, must show that the speakers are disagreeing about the correct application of 'know' and must allow the truth conditions of their knowledge ascriptions and denials not only to correspond to the utterances of each speaker, but also to the shared interests of the group and the general direction of the conversation. Keeping this in mind, we can appreciate the value of the reasonableness views that DeRose discusses. On the binding arbitration model, the epistemic standards are put in place by the speaker whose personally indicated content appears to correspond to the most reasonable standards. On the pure reasonableness view, the epistemic standards depend only on what is actually reasonable and does not refer to the personally indicated contents of either speaker.

I am not sure which option is preferable to the other. One problem with the arbitration model is that it seems that in some cases, both participants might unreasonably be employing standards that are too relaxed for the purpose at hand. Imagine, for instance, a doctor who prescribes medicines based solely on his knowledge of studies and experiments conducted fifty years ago (perhaps during his schooling), whereas his colleague insists that they look at studies conducted forty-five years ago. It seems that both the doctor and his colleague have failed to posit epistemic standards that are sufficiently stringent for their task. The worry with the pure reasonableness view is that it seems somewhat strange for the contextualist to maintain that neither speaker's personally indicated content bears relevance to the truth conditions of knowledge

assertions in that conversation. And, perhaps most daunting, there remains the task of settling what counts as 'reasonable'.

One possibility is to identify 'reasonable' with what normal people would agree is reasonable in the circumstances presented. This would place some constraints on what counts as 'reasonable', but perhaps the worry would now be transferred to what counts as 'normal'. However, I suspect that we might be able to rely on some intuitive usage of 'normal', such as someone void of mental illness, etc. Such an account, I believe would (1)allow the contextualist to maintain their thesis that the truth conditions of knowledge claims depend on the interests and purposes of individuals in such a context and (2)yield a verdict that is intuitively plausible. While I suspect that the solution is within reach, additional work must be done to settle the issue.

Works Cited

- Bach, Kent. "The Semantics-Pragmatics Distinction: What it is and Why it Matters." In *The Semantics-Pragmatics Interface from Different Points of View*, ed. Ken Turner, 65-84. Oxford: Elsevier, 1999.
- _____. "Speaking Loosely: Sentence Non-literality." *Midwest Studies in Philosophy* 25 (2001): 249-263.
- _____. 2001. 'You Don't Say?' *Synthese* 127:11-31.
- _____. "The Emperor's New 'Knows.'" In *Contextualism in Philosophy: Knowledge, Meaning and Truth*, ed. Gerhard Preyer and Georg Peter, 51-89. Oxford: Clarendon Press, 2005.
- Black, Tim. "Classic Invariantism, Relevance and Warranted Assertability Maneuvers." *Philosophical Quarterly* 55, no. 219 (April 2005): 328-336.
- _____. "A Warranted Assertability Defense of a Moorean Response to Skepticism." *Acta Analytica* 23 (2008): 187-205.
- Black, Tim, and Peter Murphy. "Avoiding the Dogmatic Commitments of Contextualism." *Grazer Philosophische Studien* 69 (2005): 165-182
- _____. "Defending a Sensitive Neo-Moorean Invariantism." In *New Waves in Epistemology*, ed. Vincent Hendricks and Duncan Pritchard. 8-27. Ashgate: Palgrave Macmillan, 2005.
- Blome-Tillmann, Michael. "The Indexicality of Knowledge." *Philosophical Studies* 138 (2008): 29-53.
- Brogaard, Berit. "What Mary did Yesterday." *Philosophy and Phenomenological Research* 78 (March 2009): 439-467.
- _____. "I know, therefore I Understand." *Typescript*.
- Brown, Jessica. "Subject Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Nous* 42, no. 2 (2008): 167-189.
- Cappelen Herman and Ernie Lepore. *Insensitive Semantics*. Oxford: Blackwell Publishing, 2005.
- Cohen, Stewart. "How to be a Fallibilist." *Philosophical Perspectives* 2 (1998): 102.
- _____. "Contextualism, Skepticism and the Structure of Reasons." In *Philosophical*

Perspectives 13: Epistemology, ed. James E. Tomberlin, 57-89. Atascadero: Ridgeway, 1999.

- _____. "Contextualism and Unhappy Face Solutions: Reply to Schiffer." *Philosophical Studies* 119 (2004): 185-197.
- _____. "Knowledge, Speaker and Subject." *The Philosophical Quarterly* 55, no. 219 (April 2005): 199-212.
- Conee, Earl and Richard Feldman. "The Generality Problem for Reliabilism." In *Epistemology: An Anthology*, ed. Ernest Sosa and Jaegwon Kim, 372 (Malden, MA:BlackwellPublishing), 372-386.
- DeRose, Keith. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52, no.4 (December 1992): 913-929.
- _____. "Solving the Skeptical Problem." In *Skepticism*, ed. Keith DeRose and Ted A. Warfield, 183-219. New York: Oxford University Press, 1999.
- _____. "Contextualism: An Explanation and Defense." In *The Blackwell Guide to Epistemology*, ed. John Greco and Ernest Sosa, 187-205. Malden,Mass.:Blackwell Publishers, 1999.
- _____. "Assertion, Knowledge and Context." *The Philosophical Review* 111, no. 2 (April 2002): 167-203.
- _____. "Bamboozled by our own words by Our Own Words': Semantic Blindness and Some Objections Against Contextualism." *Philosophy and Phenomenological Research* 73, no. 2 (September 2002): 316-338.
- _____. "Single Scoreboard Semantics." *Philosophical Studies* 119, no. 1-2 (May 2004): 1-21.
- Dretske, Fred. "Contrastive Statements" *Philosophical Review* 81, no. 4 (Oct. 1972):411-437.
- _____. "The Pragmatic Dimension of Knowledge" *Philosophical Studies* 40, no. 3 (Nov. 1981): 364.
- Elgin, Catherine. "Understanding and the Facts." *Philosophical Studies* 132 (January 2007): 33-42.
- _____. *Considered Judgment*. Princeton: Princeton University Press, 1996.
- Gottschling, Verena. "Keeping the Conversational Score: Constraints for an Optimal Contextualist Answer." *Erkenntnis* 61, no. 2-3 (November 2004): 303.

- Grice, Paul. "The Causal Theory of Perception." *Proceedings of the Aristotelian Society, Supplementary Volumes* 35 (1961): 121-168.
- _____. *Studies in the Way of Words*. Cambridge: Harvard University Press, 1989.
- Grimm, Stephen. "Is Understanding a Species of Knowledge?" *British Journal of the Philosophy of Science* 57 (September 2006): 515-535.
- Halliday, Daniel. "Contextualism, Comparatives and Gradability." *Philosophical Studies* 132, no. 2 (January 2007): 381- 393.
- Hawthorne, John. *Knowledge and Lotteries*. Oxford, Clarendon Press, 2004.
- Heller, Mark. "The Simple Solution to the Problem of Generality." *Nous* 29, no. 4 (December 1995): 501-515.
- _____. "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives* 13 (1999): 115.
- Hetherington, Stephen. "Knowing (How it is) that *p*: Degrees and Qualities of Knowledge." *Veritas* 50 (December 2005): 129-162.
- Higginbotham, James. "The Semantics of Questions." In *The Handbook of Contemporary Semantic Theory*, ed. Shalom Lappin, 361-383. Oxford: Oxford University Press, 1996.
- Kennedy, Christopher. "Vagueness and grammar: the semantics of relative and absolute gradable adjectives." *Linguist Philos* 30 (2007): 1-45.
- Kennedy, Christopher, and Louise McNally. "Scale Structure, Degree Modification and the Semantics of Gradable Predicates." *Language* 81, No. 2 (2005): 345-381.
- Kvanvig, Jonathan. *The Pursuit of Knowledge and the Value of Understanding*. Cambridge, Cambridge University Press, 2003.
- Lewis, David. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8 (August 1979): 339-359.
- _____. "Whether Report." In *Papers in Philosophical Logic*. Cambridge: Cambridge University Press, 1998.
- _____. "Elusive Knowledge." In *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted Warfield, 220-238. Oxford: Oxford University Press, 1999.
- Malcolm, Norman. "Knowledge and Belief." *Mind* 61, no. 242 (April 1952): 178-189.

- McGrath, Matt. "Evidence, Pragmatics and Justification." *Philosophical Review* 111, no. 1 (2002): 67-94.
- Rysiew, Patrick. "The Context-Sensitivity of Knowledge Attributions." *Nous* 35, no. 4 (2001): 477-514.
- _____. "Speaking of Knowing." *Nous* 41, no. 4 (December 2007): 627-662.
- Schaffer, Jonathan. "Knowing the Answer." *Philosophy and Phenomenological Research* 75, No. 2 (September 2007): 383-403.
- Schiffer, Stephen. "Contextualist Solutions to Skepticism." *Proceedings of the Aristotelian Society* 96 (1996): 317-333.
- Sperber, Dan and Deirdre Wilson. "Loose Talk." *Proceedings of the Aristotelian Society* 96 (1985-1986): 153-171.
- Stanely, Jason. "On the Linguistic Basis for Contextualism." *Philosophical Studies* 119, no. 1 (May 2004): 124.
- _____. *Knowledge and Practical Interests*. Oxford: Oxford University Press, 2005.
- _____. "Knowledge and Certainty." *Philosophical Issues* 18 (2008): 41.
- Stanley, Jason and Timothy Williamson. "Knowing How." *Journal of Philosophy* 98 (August 2001): 411-444.
- Stine, Gail. "Skepticism, Relevant Alternatives and Deductive Closure." *Philosophical Studies* 29, no. 4 (April 1976): 249-261.
- Unger, Peter. *Ignorance*. Oxford: Oxford University Press, 2002.
- Williamson, Timothy. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.