

KNOWLEDGE CLAIMS AND CONTEXT: BELIEF DIFFERENCES

Wayne A. Davis
Georgetown University
April 2007

The use of “S knows p” varies from context to context. Contextualist theories explain this variation in terms of semantic hypotheses: “S knows p” is relative or indexical in meaning, referring to practical features of either the subject’s context (Hawthorne and Stanley) or the ascriber’s context (Lewis, Cohen, and DeRose). The linguistic evidence against both forms of contextualism is extensive. I maintain that the contextual variation of knowledge claims results from common pragmatic factors. One is variable strictness (Davis 2007). “S knows p” is commonly used loosely to implicate “S is close enough to knowing p for contextually indicated purposes.” Here I will show that there is another: differences in belief about what is known.¹ This multivariate theory provides a much better explanation of the available linguistic data.

I. CONTEXTUALIST THEORIES.

Whether people say they know sometimes varies from context to context without any evident change in the content, truth, or justification of what they believe. Similar variation occurs when knowledge is ascribed to others. Here is a well-known and representative case, each of which involves two contexts, A and B.²

The Bank Case

A. Hannah and Bob are driving home on Friday. They had planned to stop at the bank, but notice long lines. Asked whether she knows if the bank will be open tomorrow, Hannah recalls going to the bank on Saturday, and says “Yes, let’s go then instead.” B. Bob reminds her that the funds must be in the bank by Monday morning, otherwise some checks will bounce. Noting that banks sometimes change their hours, he asks whether she really knows that it will be open. Hannah pauses, and says “No. We’d better go now.”

The Parking Case

A. Before lunch, Dick and Jane were discussing how often they forget where they parked their cars, and have to hunt for them. Jane asks, “Does Alan know his car is in section 5?” Dick answers “Yes. He just told me that this morning.” B. After lunch, Dick and Jane begin discussing car thefts, noting that cars have been stolen from even the safest neighborhoods. Jane again asks, “Does Alan know his car is in section 5?” Dick answered “No. He hasn’t checked.”

Hannah believes that the bank will be open on Saturday both before and after her

1. I am again indebted to Emily Evans, Paul Naquin, Nate Olson, David Pierce, Mark Pitlyk, Diana Puglisi, Dan Quattrone, Gerhard Ernst, Paul Portner, Elena Herburger, and Stephen Gross. My greatest thanks go to Stewart Cohen, who continues to provide a steady stream of thought-provoking objections and replies.

2. DeRose 1992: 913; Vogel 1999: 161. Cohen’s (2000a: 95) airport and Hawthorne’s (2004: 69) capital cases have the same features.

husband's reminder, and we may imagine that her belief is true and justified throughout. Her reason for believing that the bank will be open, moreover, does not change in the example. Nor does the likelihood that the bank will be open on Saturday given the evidence she has. As a result, her belief that the bank will be open is just as justified after her reminder. Nevertheless, Hannah claimed knowledge in A and denied it in B. And her affirmation in A seems as natural and appropriate as her denial in B.

An account of this variation in terms of *ambiguity* is implausible (contrast Malcolm 1952). "Know" has a markedly different meaning, for example, in "I don't know Vladimir Putin" and "I do know how to play the piano" than it does in "I know that the bank will be open on Saturday." But no meaning shift is apparent in any of our examples. If "know" had a different meaning after the original knowledge claim was questioned, then we would not hear the speaker as denying what he or she had earlier affirmed. We would not hear Hannah's subsequent "No" as a concession or retraction if Bob had changed the subject. What needs to be explained is how the use of "know" can vary from context to context *even when it is used with the same sense*.

What changes in the bank and parking examples is the level of justification necessary for claiming knowledge. This change occurs in the bank case because Hannah's husband reminded her of the importance of getting money in the bank before Monday.³ *Contextualism* maintains that the level of justification necessary for claiming knowledge changes in these cases because the truth conditions of the knowledge claims change. In general, contextualism hypothesizes that the truth conditions of "S knows p" are relative to a standard of justification that varies with contextual factors such as interests, stakes, and salience. These factors are independent of whether S has a true, justified, belief that p; of what reasons S has for believing p, and whether S's reasoning is sound; of how likely it is that S's belief is true given S's evidence; of how justified S's belief is; and of whether S's justification is defective in the ways that produce Gettier cases. *Invariantism* denies that the truth conditions of "S knows p" are relative to a standard that varies with such independent factors. Invariantism allows as well as contextualism that some facts about S's context affect the truth conditions of "S knows p." For example, the fact that two objects look the same may enable a subject to know that they are the same color when the lighting is good but not when the lighting is poor. But this contextual factor determines how likely it is that the objects are the same given that they look the same. So it is not the sort of independent factor the contextualist postulates. No brief characterization is completely apt, but we can roughly describe contextualism as saying that the truth of a knowledge claim depends on "practical" facts that are not "truth-related," while invariantism insists that only "intellectual," "truth-related" facts are relevant (Stanley 2005: 1-3).

Contextualist accounts can be *subject relative* or *ascriber relative*, depending on whether the independent factors determining the truth conditions of "S knows p"

3. Since the mere difference in time does not explain the variation in our cases, will ignore the indexical temporal reference implicit in the present tense.

are features of the subject's context or those of the ascriber's context. In the bank case, the subject term was the first person pronoun "I," so the subject and the ascriber were one and the same. But in general, they may differ, as in the parking case, where the subject is Alan and the ascriber is Dick. The most widely accepted form of contextualism is the ascriber relative theory of Lewis, Cohen, and DeRose.⁴ On this theory, "know" has a syntactically unmarked element of *indexicality* in addition to its tense. The standard of justification S needs to meet for "S knows p" to be true is determined by what is salient in the context of the speaker using "S knows p," which varies with the interests of the speaker. When possibilities of error, or negative consequences of error, become salient, a higher level of justification is needed for "S knows p" to be true. Hawthorne (2004) and Stanley (2005) have more recently championed subject contextualist views. Their general idea is that the more important being right about "p" is to decisions S needs to make (the greater the cost to S of being wrong; the more S has at stake), the higher the level of justification S needs to count as knowing p.

Since the ascriber was the subject in the bank case, it supports both forms of contextualism equally well. When the subject and ascriber differ, however, the variation in knowledge claims tracks variation in the ascriber's context, not that in the subject's context. In the parking case, for example, the subject of the knowledge attribution is Alan. There is no change at all in his context. The change from A to B is a change in the context of ascription. Subject contextualist views also imply, most implausibly, that subjects can acquire knowledge by moving to contexts in which the stakes are lower, and that it is correct to say of Hannah in context B that she *used to know* that the bank would be open.

Ascriber contextualism has its own counterintuitive consequences, however. Like subject contextualism, it predicts that Hannah's and Dick's second knowledge claims do not contradict their first, and are not denials, retractions, or corrections thereof. Hannah would not defend both claims the way she might defend both "Michael Jordan is tall" (said when comparing him to children) and "Michael Jordan is not tall" (said when comparing him to basketball players).

If contextualism were true, "know" should behave like "flat," "tall," and other relative terms (Cohen 1999: 60). It does not (Davis 2007: §VIII). In addition to having no comparative and superlative forms, we never say things like "*He knows it's true, but not perfectly,*" "*He knows by low standards, but not high standards,*" "*He knows compared to Mary, but not Jane.*"

Furthermore, if ascriber contextualism were true, argument form (1) would be invalid in contexts in which the speaker is in a high-standard context and S a low-standard context.⁵

4. See Lewis 1979; 1996; Cohen 1986; 1987; 1988; 1999; 2000a; 2000b; 2001; 2004; 2005; DeRose 1992; 1995, 1996a; 1996b; 1999; 2002. See also Unger 1975; 1984:46-54; 1986; Stine 1976; Goldman 1976: 776; Dretske 1981a: 367; Heller 1999; Hofweber 1999; Klein 2000; Fogelin 2000; Hawthorne 2004: Ch. 2; Ludlow 2005; Schaffer 2006.

5. Cf. Hawthorne 2004: 107, fn. 125. Contrast Cohen 2005: 205-6.

- (1) S speaks truly when saying “I know p.”
Therefore, S knows p.

In fact, we understand (1) to be valid in all contexts. Similarly, ascriber contextualism predicts that in context B, Hannah should be willing to say the following:

- (2) In A, I did not know that the bank would be open on Saturday, but what I said when I uttered “I know that the bank will be open” was true.

For according to ascriber contextualism, what Hannah said when using “know” in A was determined by the low standards prevailing in A. So what she said would have been true assuming that the bank does open on Saturday. Yet (2) sounds completely contradictory.⁶

Knowing p is not in general sufficient for properly asserting p. Asserting p may be improper simply because doing so would be misleading, insulting, or question-begging. It does seem obvious, however, that if S knows p, then S’s assertion can not be criticized on the grounds that S does not have enough evidence to support it. A statement like (3) seems absurd. The two conjuncts appear incompatible.

- (3) Mary knows that Tom was home last night, but the evidence she has is not enough for her to assert that in some contexts.

Contextualism allows (3) to be true, however, when the speaker (or subject) is in a low standards context.

V. VARIATION IN STRICTNESS

Contextualist theories seek to account for variability in the use of “know” by proposing a *semantic* theory, according to which the truth conditions of “S knows p” refer to a contextual factor. Since these theories fit the linguistic data so poorly, I have sought a *pragmatic* explanation of the observed variability. Variation in use is often due to factors that have no effect on truth conditions. For example, logicians and linguists have long recognized that the use of “Some S is P” varies with contextual factors. Sometimes speakers use it to imply “Not all S are P.” Sometimes they do not. The general consensus is that “Some S is P” unambiguously means “At least one S is P,” and that the “Not all” implication is a *conversational implicature*.⁷ Implicature is the phenomenon of meaning or implying one thing by saying something else. Conversational implicatures are those that depend on the

6. Cf. Kompa 2002; Cappelen & Lepore 2003: 29, 31; Davis 2005b: 39; Bach 2005: 60; Stanley 2005: 54ff, 115, 119ff. Contrast DeRose 1992: 925-6; Cohen 2005: 205-6.

7. The theory of conversational implicature was developed by Grice 1975; 1978; 1989. See Davis 2005c for an introduction and references.

conversational context, and are not part of the meaning of the sentence used. No one, to my knowledge, has tried to account for the variable usage of “some” by suggesting that it has an indexical meaning sensitive to whether or not the context is inclusive or exclusive. Such a theory would be implausible for some of the same reasons we have offered against an indexical theory of knowledge. For example, S’s claim “Some metals are radioactive” contradicts T’s claim “It is not the case that some metals are radioactive,” even if T is in a context in which uttering “some” would exclude “all” and S is in a context in which it would not.⁸

I attribute the differences in the bank and parking cases to the difference between *loose* and *strict* usage. These cases are very similar to the following.

The Coffee Case

A. When the scoop comes up empty in the coffee jar, I yell to my wife, “The coffee is all gone.” B. When my son comes down for breakfast a few minutes later, he announces that he needs a few coffee grounds for his science project, and then asks “Is the coffee really all gone?” I say with no embarrassment, “No, there may be enough for you.”

What I *say* in A contradicts what I say in B. But what I *mean* in A is that *the coffee is close enough to being all gone for the purpose at hand* – making coffee. That is, the coffee is *close enough to being all gone to count as all gone for the indicated purpose*. This is not what the sentence I use means. But by saying that the coffee is all gone, I convey the less precise thought indicated. By ignoring irrelevant detail, I make my point more effectively. When my conversational purposes change in context B, I use the term more strictly. Put another way, when we use a term loosely in a sentence, we imply that the proposition it expresses is a good enough *approximation* to the truth so that the difference does not matter for current purposes. My use of “all gone” displays the same kind of variability we have observed with “know” even though “all gone” does not have a relative or indexical sense.

Some types of sentences are almost always used with some looseness, such as those giving measurements on a continuous scale. If I said “My son is 5’8” tall,” I would be unlikely to mean that precisely. In other cases, loose use is less common but often helpful. Thus if people ask me where I was born, I will generally answer “In Detroit Michigan,” even though strictly speaking I was born in Highland Park, a little city completely inside Detroit. Most people have no idea where Highland Park is, and so would get nothing out of my answer if I spoke strictly. Speaking loosely, however, generally gives people the information they really want. In the rare case in which someone familiar with Highland Park asked me where I was born, I would speak strictly.

I noted above that “know” is not a relative term like “flat.” What do behave like “know” are the absolute terms “perfectly/completely flat.” Strictly speaking, something is perfectly or completely flat only if its surface has no bumps of any size (no matter how small) and no curvature of any degree. The terms literally apply only

8. Be careful not to give “some” in the negation the fall-rise intonation indicative of a “metalinguistic” negation (Horn 1989: 368; 374-5; 392-413).

to geometric objects. No material object, no matter how precisely it has been planed or polished, is perfectly flat. Nevertheless, we often describe surfaces as “perfectly flat,” and the standards by which we do so vary from context to context. Fine woodworkers apply the term much more strictly than road builders do. It seems clear that we commonly use “perfectly flat” to implicate *close enough to perfectly flat for contextually indicated purposes*. Since the woodworker’s purposes require more precision than the road builder’s, the word worker uses the term more strictly, and the road builder’s usage is looser. Both should agree that when they describe something as perfectly flat, what they say is strictly speaking false. This does not matter because it is understood that they are using the term loosely, to imply something that is true.

I suggest that the contextual variability illustrated by our bank and parking cases is the same phenomenon found in these cases of variably strict usage. “Know” is used loosely in the A-contexts of our two cases, and more strictly in the B-contexts. Thus at first, Hannah used “I know the bank will be open” loosely to mean or imply that *she is close enough to knowing for their purposes*. Hannah did not care whether strictly speaking she knows. The purposes she and her husband were focused on in A did not require precision. The difference between knowledge and highly justified true belief was immaterial. After her husband made greater precision important, however, she used “know” more strictly, denying that she knows. On this account, we would not take “Hannah used to know” to be true in the B-context any more than we would take “The coffee used to be all gone” to be true in the B context of the coffee case. The variable strictness account also allows that what Hannah *said* in the A-context contradicts what she said in the B-context. But what she *implied* in the two contexts is consistent. The difference in usage is due to a difference in the perceived truth value of an *implicature*, which referred to different purposes in the two cases. Finally, because the purposes at hand may require more or less justification, loose uses of “know” may be more or less strict, as emphasized by the contextualist. What is close enough to knowing will vary from context to context.⁹

What purposes require knowledge, or something close to knowledge? Our cases illustrate two: *making a decision*, and *answering a question*. Hannah and Bob are trying to decide what to do. What Hannah claimed to know was a reason for making one decision rather than another. The more important that decision is, the closer to knowledge their reasons for action need to be. The raising of the standards in the bank case led to a change in what the subjects said by requiring a closer approximation to knowledge. In the parking case, Dick’s purpose is to answer Jane’s question. His answers change when his view of what information Jane wants changes.

Some reject a pragmatic account of the contextual variability of knowledge

9. Compare and contrast Unger 1975: 68-9, 83-7; 1984: 6ff; Rysiew 2001: 487-8. Cohen has suggested treating loose use semantically, in terms of indexical quantifier restriction or parameter specification. I show that that does not work in Davis 2007: 415-7.

claims on the grounds that it endorses skepticism.¹⁰ Hawthorne (2004: 116-8), for example, briefly considers the theory that all positive knowledge claims are “harmless exaggerations”: literally false overstatements that pragmatically communicate weaker truths. He recounts Grice’s treatment of exaggeration, according to which we communicate something true by flouting the Maxim of Quality at the level of what is said. Thus Jane might say “I know I got everything wrong,” which is obviously much too strong, in order to communicate the weaker truth that she did poorly.

The thesis that variation in knowledge claims is due to variations in strictness does not similarly imply skepticism. Loose use may but need not be *hyperbole*. We engage in hyperbole when we know our statement is literally false, and wish to overstate the situation for effect. We may engage in loose use when we do not know or care whether our statement is literally false. In hyperbole we want hearers to notice the contrast between what was said and meant. In other cases of loose use, the focus is on what is meant. Thus I did not look carefully at the coffee jar until our son’s need became apparent. The conversations typified by the bank and parking cases do not involve hyperbole. Conversely, not all hyperbole is loose use. Suppose I see that the coffee is half gone, and say “The coffee is all gone.” I might be implying that someone has been drinking a lot of coffee, not that the coffee is nearly gone. In that case, I would be exaggerating rather than speaking loosely.

Hyperbole is typically conveyed in a special tone of voice. No special intonation marks loose usage that is not hyperbole. We generally use “literally” when we wish to contrast with hyperbole. Thus if say “I *know* he’s going to screw up,” and then worry that someone could misinterpret me, I might say “I don’t literally know it, of course, but it is more than likely.” We generally use “strictly,” “exactly,” or “precisely” instead when we wish to contrast with loose use that is not hyperbole. Thus even before her husband weighed in, Hannah might have clarified her first answer by saying “Strictly speaking, of course, I don’t *know* that the bank will be open tomorrow, but I have every reason to believe it.” We use “really” with emphasis when contrasting with hyperbole, and without emphasis when contrasting with loose use. Compare “I don’t *really* know he’s going to screw up” with “I don’t really *know* that the bank will be open.”¹¹ Hyperbole is a figure of speech. Loose use is approximation.

Grice observed that implicatures are typically *cancelable*. For example, the “not all” implicature of “some” is canceled when we say “Some passengers died, indeed all did.” Loose use provides many exceptions to cancelability, because what is said entails what is implicated. “The coffee is all gone” entails “The coffee is close enough to being all gone for current purposes.” Grice (1975: 39) considered it impossible, but implicatures are often entailed (Davis 1998: 6). For example, if someone asks “Did anyone die?” we could answer by saying “Everyone did”; we would then implicate “Someone died.” Thus the loose-use implicature of “S knows

10. Mylan Engel raised this objection at the Mainz conference on contextualism (September 2003). Cf. DeRose 1999: 202; Cohen 1999: 83; Stanley 2004: 141ff; 2005: 84.

11. See Partee 2004: 155-7. Compare and contrast Hawthorne 2004: 120.

p” is uncancelable because it is entailed.¹² “I know, but am not close enough to knowing” is inconsistent. When terms are used loosely in negative sentences, on the other hand, what is implicated entails what is said. So the implicature should be cancelable. “S doesn’t know that the bank is open on Saturday, but he is close enough to knowing for current purposes” is consistent and natural.

The view that the variation of knowledge claims in the bank and parking cases can be accounted for in terms of implicature does not entail that the loose-use implicature can be “calculated” or “worked out” the way Grice (1975) proposed.¹³ Few if any implicatures involve such reasoning (Davis 1998), including the paradigm case of quantity implicature. Purported derivations inevitably fail to show that the conversational principles require the observed implicature, and would just as well show that other implicatures exist that we do not observe.¹⁴ This is true even of the paradigm case of quantity implicatures, illustrated by “some” implicating “not all.” It is also true of hyperbole and other figures of speech. It is nonetheless easy to see how loose use enables us to be cooperative, and to contribute a sufficient quantity of true and relevant information to the conversation in an efficient way. We use “know” loosely when the purposes of the conversation do not require us to commit ourselves to the proposition that the subject strictly speaking knows – when communicating that the subject is close to knowing suffices. The difference between justification and complete justification is often irrelevant. Loose use enables us to achieve brevity (honoring Manner) and relevance while being truthful (satisfying Quality) at the level of what is implicated if not what is said.¹⁵

Hawthorne (2004: 23-4) has argued for subject contextualism on the basis of Williamson’s (2000: §11.3) thesis that *S may properly assert p only if S knows p*. Williamson observes, for example, that we are as disinclined to flat-out assert that a man will lose a lottery as we are to assert that we know he will, even though we have exceedingly strong evidence that he will lose. Williamson also observed that a sentence of the form “*p, but I do not know that p*” is always unassertible, even though it may well be true (Moore 1962: 277). Williamson’s final piece of evidence is that “*How do you know?*” is always an appropriate response to an assertion. That question’s presupposition that the addressee knows must therefore always be appropriate. Whether S may properly assert p, however, depends on S’s context. Thus it is just as appropriate for Hannah to say in context A is “*The bank will be open*” as it is for her to say “*I know the bank will be open.*” And in the B context, the simple assertion would be just as inappropriate as the knowledge claim.

Given the contextual variation in assertibility, Williamson’s assertion constraint

12. Contrast Cohen 1999: 60; 2000b: 138.

13. Contrast DeRose 2002: 176; Bach 2005: 73; and Stanley 2005: 15

14. Rysiew’s (2001: 491) attempt to derive the implicature “S’s epistemic position is good enough” from Grice’s Maxim of Relevance and the Cooperative Principle suffers from the same defects, as does DeRose’s (1999: 197-200) attempt to derive the “don’t know” implicature of “possible” from a variant of the Maxim of Quantity.

15. Contrast DeRose (2002: 193), who assumes that cases of “misdirection, like irony and hyperbole” are the only cases in which cooperative speakers implicate something true by saying something false.

is tenable only if the truth conditions of “S knows p” also varies with the subject S’s context. As we have seen, though, they do not. Suppose that we would be in big trouble if we counted on the bank being open on Saturday and it is not. Then with reference to context A, we will deny that “Hannah knows the bank will be open” is true. We would not, however, fault Hannah for having asserted that the bank will be open. For we recognize that she had sufficient evidence for making it given her and her husbands purposes. S may properly assert p without knowing p as long as S is close enough to knowing for the purposes relevant in S’s context. That is, the condition for the assertibility of “p” is that the implicature of “I know p” is true.

VI. AN INVARIANTIST SEMANTICS

Contrary to what DeRose (2002: 176) and Stanley (2005: 15) have suggested, explaining variation in the use of “know” by reference to pragmatic factors does not require providing a semantic analysis. There is no requirement that meaningful terms be definable whether they have implicatures or not. We may observe, for example, that “impossible” can be used loosely, to implicate “close enough to impossible for current purposes,” without being able to define modal notions like “impossible.” We know well from the Gettier (1963) problem, moreover, that the analysis of “S knows p” is very difficult, and may have no solution.

Nevertheless, we can give a rough explanation of what “know” means that is accurate enough to account for the variation observed in the bank and parking cases: *completely and non-defectively justified true belief*. I would characterize a completely (or *fully*) justified belief as one that is either self-evident or based on evidence sufficient to establish its truth. Defective justification is the sort found in Gettier cases. Whether either of these concepts can be clarified further is an open question. This definition attempts to describe what speakers mean when “know” is used strictly. It is used loosely to implicate *close enough to completely and nondefectively justified true belief for contextually indicated purposes*.

The requirement that the justification be complete can be defended by observing that “S knows p, but is not completely justified in believing p” sounds self-contradictory, as does “S knows p, but S’s evidence does not suffice to establish p.” The complete justification requirement also explains why no one knows the results of a lottery based solely on the odds. In my view, furthermore, the claim that Hannah literally knew on Friday night that the bank would be open on Saturday is untenable. There are too many very real possibilities that Hannah cannot rule out: that the bank changed its hours since she last checked, that an unexpected thunderstorm would lead to a massive power outage, and so on. These things happen with some regularity. Hence her evidence does not suffice to establish that the bank is open on Saturday. While Hannah is justified in believing that the bank will be open, she is not completely justified. So in my view, what Hannah *says* in the A-context is false, even though what she *means* (that she is close enough to knowing for current purposes) is true. In the B-context, what she says and means are both true.

My purpose here, however, is not to defend the particular invariantist

semantics proposed, but to show that pragmatic factors such as variable strictness combined with a suitable invariantist semantics provides a better account of the contextual variation of knowledge claims than a contextualist semantics. The bank and parking cases could be explained in terms of variable strictness even if something weaker replaced “completely,” such as “nearly completely” or “highly” in the semantics. Indeed, a pragmatic account could use an invariantist semantics so weak that the affirmations in the A-contexts of the bank and parking cases were true while the denials in the B-contexts were false (cf. Bach 2005: 75). The change in usage could still be due to a change in strictness: the speakers might have been wrong in thinking more evidence was needed in the B contexts. We could even relax my assumption that what counts as complete justification does not vary from context to context.¹⁶ We could account for the variation in the bank and parking cases in terms of variably strict usage as long as the range of variation displayed in the bank and parking cases is outside that allowed by complete justification. I would expect, though, that even limited context-sensitivity would lead to the sorts of problems we noted in §I.

Cohen has objected that if all sentences ascribing knowledge of precise measurements are strictly speaking false, then no one ever knows what time it is, or when anything occurs, or how long anything is (cf. §XI). That conclusion would be absurd indeed, but it does not follow. I can know that John is *about* 5'8" tall, and that suffices to know how tall he is. What I cannot know is *precisely* how tall he is, with absolutely no margin of error. But knowing that is not necessary for knowing how tall he is. Furthermore, even though “I know John is 5'9" tall” is strictly speaking false, I can nonetheless use it to imply that I know how tall he is by using it loosely. The semantic thesis that knowledge requires complete justification does not entail skepticism without the supplementary epistemological claim that we are never (or not ordinarily) completely justified in believing anything. I deny that additional claim. While Hannah did not have enough evidence in either context to know that the bank would be open on Saturday, she was completely justified in believing many things, such as “I am speaking” and “It is Friday.”

DeRose (1995: 46) has suggested that on *linguistic* grounds, we should prefer a semantics for “S knows p” that succeeds in “crediting most of our attributions of knowledge with truth.” But on purely linguistic grounds, why should attributions be privileged over denials? Why should truth in what is said be favored over truth in what is implicated? That such a preference has no basis in linguistics is shown by the fact that terms like “witch,” “extraterrestrial,” and “Santa Claus” are perfectly meaningful even though all attributions of these terms have been false. We can no more require a semantics for “know” to yield the result that most attributions of knowledge are true than we could require a semantics for “god” to yield the result that most statements about gods are true. To argue for his principle, DeRose asked why we would reject a definition of “physician” that made “able to cure any conceivable illness” part of its meaning. His answer was that we apply the term “physician” to many people we do not take to satisfy that requirement. But this

16. Contrast Dretske 1981b: 49, 51-2, 56.

relies on the principle that the meaning of a term is determined by the *beliefs* we use it to express (see Davis 2003). The meaning of a term does not depend on the *truth* of those beliefs.

Similarly, if someone proposes that part of what “S knows p” means is “S can rule out all conceivable alternatives,” we can gather evidence against this semantic hypothesis by observing that even when people use “S knows p” strictly and literally, they commonly do not have such a belief, and are not trying to indicate that they do. If we were to find that speakers generally intended to communicate the idea that S can rule out all conceivable alternatives, then we would have to take that as part of the meaning “S knows p” even if that meant that its instances are generally false.

Contextualists seek a semantic analysis that makes the knowledge affirmations in the low-standards A-contexts above true as well as the denials in the B-contexts. We have seen that this assumption cannot be justified on the basis of general “Moorean” principles. A more direct way to defend the assumption is by appeal to *linguistic intuition*. People do have the intuition that the knowledge claims in both the A and the B contexts are *correct*, whether this means *true* or *proper* or both. The intuitions of propriety support the pragmatic theory I have sketched as well as the contextualist theories. Neither semantic theory accounts for all of our semantic intuitions. The invariantist is forced to conclude that one of these intuitions of truth is incorrect. The contextualist must give up the even stronger intuition that the knowledge claims in the A and the B contexts are incompatible. Since our intuitions about these cases are not completely consistent, no consistent theory will fit the intuitive data perfectly.

Implicatures often account for conflicting linguistic intuitions. In the case of quantity implicatures, for example, “Some S are P” intuitively implies “Not all S are P”; but “All S are P” even more clearly entails “Some S are P.” We resolve the conflicting intuitions by concluding that “Some S are P” implicates, but does not entail, “Not all S are P.” Loose use does the same. My use of “all gone” in the A context of the coffee case seems intuitively correct, as does my use in the B context. But the B claim seems clearly to contradict the A claim. We resolve the conflict by concluding that even though what I said was strictly speaking false in the A context, what I meant was true. The same resolution can be given for our conflicting intuitions about the bank and parking cases.

The support for contextualist theories provided by linguistic intuition is fragile, furthermore, because the intuitions of truth in the A-contexts do not survive *careful reflection*. As soon as we ask, “Does Hannah really know that the bank will be open on Saturday?” and reflect for a while, we lose the intuition that her statement is true and acquire the intuition that it is false. This happens because we reflect on the fact that banks often change their schedule without notice, that various natural or man-made disasters might close the bank, and so on. Careful reflection is an everyday activity, not something confined to epistemology classes or philosophical tomes. Unlike subject contextualism, ascriber contextualism can accommodate the change in intuition after careful reflection, because it maintains that the act of reflection constitutes a change in our context that raises the standard of justification.

I explain the change due to critical reflection in the bank case by hypothesizing that at first we interpret “Hannah knows” as a loose use, and that the

critical question makes us reinterpret it as a strict use. That is, at first we focus on what the speaker meant, while afterwards we focus on what the speaker strictly speaking said. Consequently, at first our intuition is that what was conveyed (i.e., implicated) is true, and afterwards our intuition is that what was conveyed (i.e., said) is false. Where contextualism postulates a shift in context that changes truth conditions by selecting a higher level of justification, my account postulates a shift in context that changes the level of intended strictness without affecting truth conditions.

A problem for contextualist theory is that there is no context in which “S knows that he will lose” seems true when the subordinate clause refers to the outcome of a lottery and S’s only evidence is the extremely low probability of winning.¹⁷ Given that the evidence about the lottery may provide extremely strong justification for believing that S will lose, and assuming that the belief is in fact true, why doesn’t “S knows that he will lose” seem true in low-standard contexts? Why wouldn’t the lottery case below be treated just like the parking case?

The First Lottery Case

A. Harry, the career dishwasher, announces that he is going to buy a Ferrari Testarossa. Incredulous, Joe asks how he plans to pay for it. Harry responds, “I’m going to win the lottery.” When he realizes that Harry is serious, Joe blurts out “That’s ridiculous. You know you’re not going to win any lottery.” B. Joe is later talking to Tom, who wonders whether Joe has some inside information about the lottery. “Does Harry really know that he is going to lose?” Tom asks. “Of course not,” Joe says. “But the chances of his winning are infinitesimal.”

Joe’s use of “know” in A is perfectly natural. And it would be inappropriate for people therein to dispute what Joe says. But this is not plausibly taken as evidence that the truth conditions for “S know p” are relative to the standards prevailing in either the ascriber’s or the subject’s contexts. We count what Joe said as strictly speaking false, and take him to be engaging in mild hyperbole. We accept Joe’s statement as permissible – indeed highly appropriate – because his point is that Harry’s plan is completely irrational given that he has more than adequate reason to believe that he will lose the lottery. Harry is close enough to knowing for the contextually indicated purpose. Cases like this thus undermine Williamson’s strongest argument for the strict assertion constraint (§II).

Cohen (1988: 107; 2000b: 135) observes a special feature of this case: it is not possible to think about a lottery without reflecting on the chance of winning. Thus whenever we think that someone has good reason to believe he will lose, it also occurs to us that the belief might be false. We engage in critical reflection immediately. In contrast, the chances of Hannah and Dick being wrong were probably not salient until the context shifted in the bank and parking cases. Because of this special feature of lotteries, Joe surely intended Harry to recognize that what he said was not strictly speaking true. That is why we take Harry to be engaging in

17. Contrast Hawthorne (2004: 161) and Cohen (2000b: 135-6; 2005: 200). Hawthorne (2004: 8) nevertheless takes it to be a datum that “there is a strong inclination to claim that lottery propositions are not true,” as he did in Hawthorne (2000: 117ff) when criticizing Cohen.

hyperbole, but not Hannah or Dick. If contextualism were true, on the other hand, it should seem evident that Harry's statement is true in the A context. But that is not our intuitive judgement.

IV. VARIATION IN BELIEF

I have accounted for much of the variability in the use of "know" pragmatically in terms of loose and strict usage. "Know" can be used strictly to mean that the subject has a fully and non-defectively justified true belief. "Know" can also be used loosely, to implicate that the subject is close enough to knowing for purposes of the context. This is not the only source of variability, however.

We distinguished loose use above from hyperbole. Both must be distinguished from *sloppy* or *careless* use. We use a term sloppily if we use it strictly without taking proper care to verify that the term does strictly apply. For an example, imagine a variant of the coffee case in which I never looked carefully in the jar; yet when my son asked whether the coffee was really all gone, needing just a few grounds, I answered "Yes, it is all gone." In that case, I was being careless. I was claiming that the coffee was literally all gone without having made the observation necessary to justify the claim. In this case, careful reflection does not lead me to use the term more strictly, because I am already using it strictly. Instead, critical reflection leads me to ask whether I have adequate evidence for my claim. When I determine that I did not look carefully, I stop believing that the jar is empty. When I observe that there are a few grounds in the jar, I form the belief that the jar is not empty and as a result say "I was wrong: the coffee is not all gone." In this case, a change in usage resulted from a change in belief. Similarly, if someone asks "Do you know what instrument that is?" while listening to Berlioz's *Herold in Italy*, I might quickly answer "Of course: a violin." If she responds "Are you sure? I think it might be a viola," I might then listen more carefully and say "I was wrong." My first answer was careless, not loose.

Unenlightened usage arises from lack of education rather than lack of care. Someone who claims that a yardstick is precisely three feet long may mean exactly that because he has never learned about the ineliminable imprecision of all measurement. If he carefully examined the markings on the yardstick – perhaps measuring it with the best computerized apparatus – his usage is unenlightened rather than sloppy. Once he learns about the limits of precision, he will, if candid, admit that he was mistaken. Before the change, he may have said and meant "I know how long the stick is: exactly three feet" because the instrument read "3.000." He was using "know" strictly, not loosely. After learning the limits of precision, he will retract his knowledge claim.

The principal difference between loose use and both careless and unenlightened usage is a difference in what the speaker means or implies. The latter involve using "p" strictly to say and mean that p. Using "p" loosely involves saying that p while meaning only that it is close enough to being the case that p. In loose usage, the speaker does not intend to commit himself to what he strictly speaking said. Hence he cannot be criticized for being sloppy or ignorant. In the bank case, it is unfair to interpret Hannah as meaning that she strictly speaking knows that the

bank will be open. For then she would be subject to criticism that seems unwarranted. Either she failed to realize that schedules sometimes change, or failed to confirm that the previous schedule was still in effect.

Variation in belief is an independent pragmatic factor accounting for variation in usage. Whether we say “p” generally depends on whether we believe that p. This holds in particular when “p” is a sentence of the form “S knows that q.” Of course, the change in the bank and parking changes does involve *a* change in belief. At first Hannah believes that she is close enough to knowing for their purposes; later she does not have that belief. But there is no indication that she ever believed what she strictly speaking said in the first context. So we cannot say that she stopped believing that she knows that the bank will be open. In the Berlioz example, in contrast, I no longer believed that I knew the instrument was a violin after my companion suggested it might be a viola. When I use “variation in belief” to refer to this additional factor accounting for whether or not we say “p,” I mean specifically the difference between believing that p and not believing that p. I am referring to the difference between holding and not holding the belief that is expressed by “p,” not beliefs that might be implied by “p.”

Whereas the speaker’s believing that p is a general factor influencing whether the speaker uses “p,” it is only rarely a factor relevant to the meaning or truth conditions of “p.” “Hannah knows that the bank is open tomorrow” obviously does not mean “The speaker believes that Hannah knows that the bank is open tomorrow.” And only for special cases like “I have a belief” will the speaker’s believing something be relevant to the truth conditions of what he said. So variation in belief is a “pragmatic” factor.

Because I attribute the contextual variation in knowledge claims to two independent pragmatic factors, strictness of usage and belief in what is expressed, I am offering a *multivariate pragmatic theory*. Other well-known factors are figures of speech, such as irony and metaphor as well as hyperbole. Like strictness of usage, the difference between using “S knows p” literally or ironically does not result from the speaker’s believing that S knows p in one case and not believing it in another. While variation in strictness accounts for the cases cited as the clearest evidence for contextualist theories, it does not account for the most interesting cases.

V. THE SECOND LOTTERY CASE

Consider a puzzling difference in our commonsense knowledge attributions that Cohen (1988: 106-9) and DeRose (1996b) have tried to explain.

The Second Lottery Case

A. Despite knowing the astronomical odds against Sheila winning the lottery, Ed says “There is a chance her ticket was drawn, so I don’t yet know that Sheila lost.” B. After reading in the newspaper that someone else won, Bill says “I now know that Sheila lost.”

We intuitively judge that Ed’s denial and Bill’s affirmation are both true. We have

these intuitions even though we know that the likelihood of an error in the newspaper may be much greater than that of Sheila's ticket being drawn. Why the difference? Cohen's explanation invokes Lewis's theory of salience as a determinant of indexical reference.

[T]he contexts in which we evaluate the two cases are different, i.e., the standards of relevance that operate are different. The explanation for why the standards differ lies in the fact that in the statistical case, unlike the other cases, the nature of the reasons make the chance of error salient. (Cohen 1988: 107)

But *we* are evaluating Ed's denial and Bill's affirmation in a single context. *We* are inclined to say that Bill knows Sheila lost even though *we* are in a context in which, if Cohen's hypothesis is correct, the standard of evidence has been raised by the salience of the chance of error. We can also imagine that Bill made his statement right after Ed did, in which case the chance of error should be as salient in B as it is in A. Differences in practical interests, levels of justification, or strictness of usage seem to play no role in this case. It would be particularly implausible to claim that Bill spoke loosely after Ed spoke strictly. Once terms have been used strictly, it is hard to use them loosely without being misunderstood (Davis 2007: 26). The level of strictness can easily shift up. Thus if Bill carefully reflects on the fact that newspapers sometimes make mistakes, he is likely to move into a context C in which he says "Of course, I don't really *know* that Sheila lost; I need to confirm the newspaper report" (cf. Cohen 1988: 107ff). The change from B to C is just like the change from A to B in the bank and parking cases. The change from A to B in the second lottery case is different, and remains to be explained.

We can attribute the difference in usage of "know" between A and B in the second lottery case to a *difference in belief*. Roughly speaking, people believe they can know things on the basis of testimony, but not on the basis of odds. People often believe that testimony that p is sufficient evidence to establish that p, but never believe that a high probability that p is sufficient to establish that p. Hence Bill believed he knew, Ed did not. As a result, Ed said "I don't know" while Bill said "I know." More precisely, Bill believed he was close enough to knowing in context B, but Ed did not believe he was close enough to knowing in context A.

Explaining why people believe they can know things on the basis of testimony but not on the basis of odds is a further problem, one for psychology rather than linguistics. One difference that may be relevant to this cognitive question is that in typical lotteries, we are certain that someone using exactly the same reason will end up with a false belief. We have no such certainty in the testimony case. Bill's reason for believing that Sheila lost is not – or at least not simply – that the newspaper (or newspapers in general, or testimony in general) is reliable. It is that a particular newspaper at a particular time reported that she lost. This suggests that there may be a contextually invariant requirement at work. But whether such a difference in belief is really justified, and if so, what that might tell us about the truth conditions of knowledge claims, has to be considered an open question.

VI. THE EPISTEMOLOGY CASE

The epistemology case that motivates much contextualist thinking does not plausibly involve switching from loose to strict usage either.

The Epistemology Case

A. Like most people, David gives little thought to philosophical questions in everyday life. If asked, “Do you know how many fingers you have?” he would answer “Yes of course, ten.” B. In college, he takes an epistemology course. The professor patiently explains Putnam’s brain in the vat hypothesis to him, and then asks him whether he *knows* he is not a brain in a vat. As expected, David reflects that things would look and feel exactly the same to him if he were; and so, after some resistance, he answers “No.” The professor then asks him again whether he knows he has ten fingers, noting that if he knows that, then he must know that he is not just a brain in a vat. David hesitates, and eventually concedes “No, I guess not.”

It is not plausible that David is using the term “know” loosely in context A, assuming he is typical of the ordinary person. In ordinary contexts, David could hardly imagine being any closer to knowing that he has fingers. He would regard himself as fully justified in believing such a thing. Nothing could be more certain, he might say. He would thus take his statement to be strictly speaking true. As a result, ordinary careful reflection would not prompt a retraction or correction. In context B, moreover, David would say that he *used to believe he knew* such things, but now he does not. David may have *disagreed* with his professor at first, and *agreed* with him later. David is different from Hannah and Dick in having *changed his mind*. David would regard himself as *enlightened*. Before that happened, he probably went through a period of mental conflict and uncertainty, in which he did not know what to believe. One of the purposes of a liberal arts education is get students to grow intellectually by making them question their established beliefs, and either replace them with better justified beliefs or get clear on why they should retain them. The process in epistemology is little different from what happens to students who have their belief in God challenged in a philosophy of religion class, or their belief in free will tested in metaphysics. Contextualism maintains, most implausibly, that all we are doing in epistemology is exposing our students to new contexts, to which their language responds sensitively.

Another distinctive feature of the epistemology case is that the subject considers additional possibilities in the B-context. Even in the A-context, Hannah was fully aware that her memory might not be perfect, and that bank schedules change from time to time, or get upset by uncontrollable events. Everyone knows that such things happen. It is not even plausible that Hannah had forgotten about such possibilities until Bob reminded her. In the epistemology case, in contrast, David did not in the A-context believe that he might be a brain in a vat; it is not common knowledge that such things are possible. The professor got David to acquire that belief, and to reflect on its implications. When David reverts to his customary beliefs, as nearly everyone does after encountering skepticism, he either abandons the belief that he might be a brain in a vat, or concludes that he can rule out or ignore that possibility. It is particularly clear that the parking case, as I have

formulated it, does not involve any change in the range of possibilities the subject considers. Alan is the subject of the conversations in the A- and the B-contexts, but is unaffected by them. Note too that only David would feel the sort of embarrassment that results when we realize we were wrong about something.

Accounting for the epistemology case was one of the prime motives of contextualism. The Hawthorne-Stanley theory fails because there is no change in stakes. The consequences of David being right or wrong about whether he has ten fingers are the same both before and after he takes seriously the possibility that he is a brain in a vat.¹⁸ A subject contextualist theory more like the Cohen-DeRose theory in making knowledge claims directly relative to the standards of justification prevailing in the subject's context would work better. On this version, practical interests are relevant only if they affect standards of justification. What neither version of subject contextualism explains is why David in the B-context would deny that he or anyone else ever knew that he has ten fingers (Cohen 2005: 207-8).

The Cohen-DeRose theory appears to account for the change of knowledge claims because the epistemology case appears to involve a raising of the standard of justification required for knowledge. In context A, Cohen and DeRose maintain, David does not have to rule out the possibility of being a brain in a vat. In context B he does. It is plausible that B requires David to have more evidence than A. So the epistemology case appears to resemble the bank and parking cases. Ascriber contextualism correctly predicts that David would deny that he or anyone else ever knew, even in the A-context. Yet without *ad hoc* supplements, it also falsely predicts that David will count his previous utterances of "I know" as true. Neither form of contextualism, therefore, accounts for David's concern that his previous claims were *mistaken*.

A further problem is that contextualists make an epistemological assumption about context B that is both debatable and detachable: that David has to rule out the possibility of being a brain in a vat in order to know that he has a hand, but cannot. Contextualists therefore assume that David's evidence does not suffice to establish that he has fingers. However, it is a controversial philosophical issue whether David needs any more evidence than he already has in the B-context to know that he has fingers. Skeptics think he does. Non-skeptics like me do not. Despite obedience to the Moorean constraint, contextualists are here siding with the skeptic against Moore.¹⁹ This issue cannot be resolved by claiming that the skeptics and non-skeptics are in different contexts. They are both in the same context discussing David in the B-context.

The contextualist's skeptical assumption is detachable in the sense that we could reject David's denial of knowledge in the B-context as false without giving up the hypothesis that "know" is context-sensitive. Even if this particular contextual

18. Cf. Stanley 2005: 126. Stanley (2005: 125ff) also suggested that perhaps the epistemologist has a practical interest that ordinary people lack, namely, learning and espousing the truth about skepticism. However, what is at stake there are the consequences of being right or wrong about whether *he knows that he has fingers* rather than the consequences of being right or wrong about whether *he has fingers*.

19. See Feldman 1999, esp. 104-7, 111; 2001: 62, 68; Klein 2000: 113; Rysiew 2001: 483; Davis 2004: 270; 2005b: 35-40; Bach 2005: 68-71.

difference is not relevant to the truth conditions of “S knows that p,” they could still be sensitive to the differences in the bank and parking cases.

VII. THE SENSE OF PARADOX

DeRose and Cohen also sought to explain the sense of paradox we feel when we do epistemology:²⁰ we generally think we know that we have fingers, and that if we have fingers, then we are not bodiless brains in a vat; yet we sometimes think we cannot know that we aren’t bodiless brains in a vat. We are pulled by the evidence of our senses to believe that we know we have fingers. We are pushed by the recognition that our immediate sensory evidence is logically compatible with being a bodiless brain in a vat to believe that we do not. The fact that the following three propositions are independently plausible but mutually incompatible thus presents a paradox.

- (4) (a) I know that I have fingers.
 (b) If I know that I have fingers, then I know that I am not a brain in a vat.
 (c) I do not know that I am not a brain in a vat.

The contextualist explanation is that we see a paradox here because we overlook the context sensitivity of knowledge claims. We find “I know I have fingers” acceptable, they maintain, because in ordinary circumstances the standard of evidence is low. We may ignore relevant alternatives. Introduction of the skeptical possibility by (4)(b) moves us into a new context in which the standard of evidence is high. We can no longer ignore the possibility that we are a brain in a vat. As a result, we accept (4)(c). Thus “I know that I have fingers” is true in the A context, false in the B context. Since the ascriber and the subject coincide in this case, subject contextualist theories can offer the same solution.

I have argued that contextualism fails to account for the sense of paradox and conflict (Davis 2004: §3, 2005: 35ff). For if contextualism were true, we should not be uncertain whether “I know that I have fingers” is true or false. It should strike us as clearly true when we are in A-contexts, clearly false when in B-contexts. In neither context should we feel any paradox arising from conflicting beliefs. DeRose does raise this question, but his answer misses the point.

[W]hy do we find these claims to know plausible even when we’re in a context in which the skeptic has raised the standards to such a level that these claims are false? ... [W]e realize that as soon as we find ourselves in more ordinary conversational contexts, it will ... be true for us to claim to know these very Os that the skeptic now denies we know.... (DeRose 1995: 41)²¹

This cannot be right. If we recognize that (4)(a) is true in ordinary contexts, then we

20. Cf. Lewis 1979: 355; Cohen 1999: 64-7; DeRose 1992: 917; 1999: §6; 2002: 168. Contrast Feldman 1999: 104-7, 111.

21. Cf. Lewis 1996: 441-4; Cohen 1999: 67; 2000a: 102.

must also recognize that “I know that I am not a brain in a vat” is true there too. We are as unlikely to accept that when we are in a high standard context as (4)(c). If we did accept both things, then by the facticity of knowledge we could infer “I am not a brain in a vat, but I do not know that I am,” a Moorean paradox.²² More importantly, when we are in the B-contexts, the fact that the sentence is true in an A-context is completely irrelevant and explains nothing about our cognitive state. The same holds, *mutatis mutandis*, when we are in an A-context. Contextualism itself provides no more reason to expect a sense of paradox in the epistemology case than it does in the bank case. There is no sense of paradox there.

As noted, contextualism wrongly likens the change in David to the change in a person who affirms “That is a foreign car” when in the U.S., while denying it in Germany, rather than to the change in a person who loses his belief in God as a result of a course in philosophy of religion. If contextualism accounts for why (4) presents a paradox, then it should also predict that (5) presents one.

- (5) (a) George’s statement that the Ferrari is a foreign car was true.
 (b) If George’s statement that the Ferrari is a foreign car is true, then Luciano’s statement that it is a foreign car was false.
 (b) Luciano’s statement that the Ferrari is a foreign car was false.

These might seem contradictory at first. But that appearance dissolves as soon as we learn that George was in America, Luciano in Italy. There is no paradox here. But if knowledge claims were context sensitive as contextualists maintain, (4) should present no more of a paradox than (5).

Contextualism does provide a coherent answer to the question of why we affirm “I know I have fingers” in one context and deny it in another. But that psycho-linguistic puzzle is not the psychological problem of accounting for the difficulty we have making up our minds about a non-linguistic matter.

VIII. ERROR THEORIES

One problem for the contextualist is to explain why we generally find ourselves impelled to retain the belief that we know we have fingers when we are in the B context. We may yield to the skeptic, but we do not do so easily, and continue to be pulled by the evidence of our senses. It is not generally true of speakers here, as it was in the bank and parking cases, that “when error possibilities become salient to them, their ascriptions track the attendant contextual shift” (Cohen 2005: 209). Since it takes a while for the skeptical arguments to sink in, people like David typically continue to believe they know for some time after the error possibilities become salient to them. I myself still do now. One suggestion is that while salience is the trigger, the context does not actually shift unless and until the speaker’s ascriptions change. But without some independent criterion of context-shift, this move is *ad hoc* and robs the contextualist theory of its ability to explain the change in

22. Elke Brendel and Mark Lance pointed this out at the Mainz conference on contextualism.

ascriptions. Any plausible theory must allow people in any given linguistic context to change their minds about what they know. This seems to be a paradigm case.

One proposal is that we are *inattentive to contextual shifts* (Cohen 1988: 111; 2005: 206, 208). While contextualism alone does not account for the data, the conjunction of contextualism with the inattention hypothesis might. As Rysiew (2001: 484ff) observes, however, when people appear to be disagreeing because they fail to notice a relevant difference in context, pointing that out normally makes the appearance disappear. This does not happen in the epistemology case. The inattention hypothesis causes internal problems for the contextualist, moreover. If David did not notice the change from the A context to the B, then contextualism has no explanation of why David eventually said “No, I guess I don’t” in the B context. If he does not notice the shift back to the everyday context, then contextualism cannot account for why he starts saying “I know that I have fingers” again. In general, if speakers are inattentive to contextual shifts, then the contextualist has no account of the variation in usage of “know” from context to context. The auxiliary hypothesis that enables contextualism to cope with the skeptical paradox undermines its ability to account for the contextual variation of knowledge claims.

Another problematic proposal is that speakers are *unaware of the context sensitivity of “know”* (Cohen 1999: 78-9).²³ This is the “error theory” Schiffer (1996: 325-8) criticizes. Cohen (2005: 208-9) correctly observes that speakers occasionally overlook the context-sensitivity of terms like “flat.” His example is Unger (1975: 49ff, 54-62), who fixated on the geometric standard of flatness, and erroneously concluded that when a Coloradan says that Iowa is flat, what he says is false. We can classify “flat” as indexical despite Unger because speakers abound who treat “flat” as an indexical, and contrast it with “perfectly flat.” English speakers commonly say things like “A person from Colorado might call Iowa ‘flat,’ but those of us from Florida would call it ‘hilly.’” What the semantic blindness theory postulates, however, is that speakers are *generally* unaware of the context-sensitivity of “know.” The idea is that speakers automatically vary their knowledge claims with changes in context, but do not recognize that they or other speakers do this. As a result, they misinterpret what people in different contexts, including themselves, have said. This supplemental hypothesis is *ad hoc*. Other context sensitive terms do not have this feature, and contextualism provides no explanation for why “know” should be special.

Imagine that a foreigner is learning English. He reliably uses “here” to refer to his present location, wherever that happens to be. However, he also interprets “here” as used by everyone else, and himself on prior occasions, as referring to his own present location. For example, when his wife calls from overseas and says “I’m with your mother here,” he responds “You can’t be: I’m all alone here.” This foreigner has not yet grasped the meaning of “here,” and is not competent in its use. Consequently, we cannot explain why his use of “here” varies with his location by observing that “here” is indexical. And since this foreigner’s mistake is exactly the

23. See also Cohen 1999: 79; 2001: 89ff, 95ff; 2005: 206, 208; Hawthorne 2004: 107-11, 114-5; Stanley 2005: 29-30. Contrast Schiffer 1996: 328; Hofweber 1999; Davis 2004: 263ff; Stanley 2005: 116; Bach 2005: 67.

error that contextualists postulate we all make with “know,” the semantic blindness theory implies that even native speakers have not grasped the meaning of “know,” and are not competent in its use. This consequence is absurd.

Hofweber (1999: 99-100) uses “summer” and “winter” to argue that the contextualist’s semantic blindness hypothesis is not as absurd as Schiffer thought. He observes that people have to learn that August is summer in the northern hemisphere and winter in the southern hemisphere. Before they learn this, they will interpret “It is winter” as uttered by an Australian in August as expressing a falsehood. However, such speakers have not yet fully understood the meaning of “It is winter” *in English*, and are not fully competent speakers of English. As a result, *their* usage of “It is winter” will *not* vary from context to context. Even if they travel to Australia and talk to Australians, they will continue to say “It is summer” in August – and not just because of a habit-induced verbal slip. If speakers are blind to the indexicality of “know,” as Cohen says, then there is again no reason to expect their usage of the term to vary from context to context in the way that would be appropriate if the contextualist theory were true.

Moreover, there is a fundamental incoherence in the error hypothesis. If it is true, then speakers of English do not reliably communicate using the term “know.” Speakers and hearers are commonly in different contexts with different standards. When they are, what speakers use “know” to say will not be what others understand them to be saying. It is only when the speaker and the hearer happen to be in the same context that what the speaker means ever coincides with what the hearer takes him to mean. This coincidence is purely accidental, given that it is not due to either speaker or hearer understanding the indexicality of “know.” If these things are true, how could “know” have any meaning in English at all? How could it acquire a meaning that even native speakers do not grasp? The foundational theory of meaning I have developed (Davis 2003, 2005a) would rule that “know” does not have a meaning because no interpretation is conventional. I am not aware of any other theory of meaning that would rule otherwise.

IX. FOUNDATIONAL NORMATIVE QUESTIONS.

Perhaps the most important feature of “S knows p” that the indexical theory ignores can be seen by reflecting further on the explanatory failure of the indexical theory. If the indexical theory were correct, we should not have difficulty figuring out which proposition in the incompatible epistemic triads to reject, any more than we have difficulty in the triad concerning the foreign car. I submit that we have trouble because consideration of the skeptical hypothesis raises a foundational question about the evidence required for knowledge. It is not an empirical question about which standard is in play in any given context, but a normative question about the proper standard.²⁴ I believe firmly, for example, that the evidence of my senses is not only the best evidence I can have that I have a hand, but all the evidence I need

24. Cf. Feldman 1999, esp. 104-7, 111; Klein 2000: 113; Kornblith 2000: 29; M. J. Williams 2000: 83. I discuss a normative version of contextualism according to which the proper standard is what varies from context to context in Davis 2005b: §IV.

to know this. Yet the fact that I would have all the same evidence if I were a brain in a vat makes me wonder whether I do not need some more evidence to rule that possibility out. The skeptical hypothesis makes me question, that is, whether the evidence of my senses is sufficient to establish that I have a hand. I convince myself that more evidence is unnecessary by reflecting on the fact that the brain in the vat hypothesis is completely ad hoc; that we have never observed anything like a disembodied brain getting all the same sensory stimulation as a real human, and cannot realistically see how anything like it could be arranged; that it is more complex than the simple hypothesis that I have a hand; and so on. Still, the hypothesis is a logical possibility. I recognize that others are unable to convince themselves that more evidence is unnecessary, and as a result have a skeptical moment. I remember having had one myself when first studying epistemology, and still wonder occasionally about the strength of my own reasons.

I believe we vacillate over which proposition in the inconsistent epistemic triad to give up because we have no obvious way of *settling* the question as to whether the evidence of our senses completely justifies believing that we have a hand. It is not clear how we can proceed without begging the question. The skeptical paradoxes raise foundational issues. The most serious flaw in the indexical theory of knowledge is that it makes any such normative deliberation as otiose as a debate over whether the .01 significance level is more proper than the .05 or .10 level. Within limits, the choice of a precise significance level is up to the investigator. The choice of epistemic standards is not. Similarly, it makes no sense to deliberate about whether the Floridian's standard or the Coloradan's is more appropriate for judgments of flatness. Since no standard of comparison is intrinsically correct, we are free to choose. No deep issues are raised. When the skeptic questions whether we really know the things we think we know, suggesting that the evidence of our senses is not good enough, it is not adequate to say that we have a different standard of evidence in mind. The question is whether the standard we have in mind is correct.

I have suggested two alternative pragmatic accounts of why David affirms "I know I have fingers" in one context and denies it in another, and argued for the superiority of the belief change account over the strictness change account. Neither account helps us figure out whether or not we really know that we are beings with fingers rather than bodiless brains who only appear to have fingers. The only way to resolve that question is to answer the normative question of what we are fully justified in believing. To answer that question, we have to determine whether we need anything more than the evidence of our senses to establish that we have fingers, and if so, what. We have to determine whether we need to rule out even far-fetched logical possibilities, and if so, how. These are issues about which rational people can and do disagree. They are difficult philosophical questions because they are foundational. Contextualism holds out the promise of dissolving this problem – of showing that once the meaning of "know" is properly understood, the problem goes away. The linguistic evidence, however, disconfirms the contextualist thesis. That evidence is better explained by a multivariate pragmatic theory.

REFERENCES

- Bach, K. (2005) The emperor's new 'knows'. In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, ed. G. Preyer & G. Peter, pp. 51-90. Oxford: Clarendon Press
- Cappelen, H. & Lepore, E. (2003) Context shifting arguments. *Philosophical Perspectives*, **17**, 25-50
- Cohen, S. (1986) Knowledge and context. *Journal of Philosophy*, **83**, 574-83
- Cohen, S. (1987) Knowledge, context, and social standards. *Synthese*, **73**, 3-26
- Cohen, S. (1988) How to be a fallibilist. *Philosophical Perspectives*, **2**, 91-123
- Cohen, S. (1999) Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, **13**(Epistemology), 57-89
- Cohen, S. (2000a) Contextualism and skepticism. *Philosophical Issues*, **10**(Skepticism), 94-107
- Cohen, S. (2000b) Replies. *Philosophical Issues*, **10**(Skepticism), 132-9
- Cohen, S. (2001) Contextualism defended: Comments on Richard Feldman's skeptical problems, contextualist solutions. *Philosophical Studies*, **103**, 87-98
- Cohen, S. (2004) Knowledge, assertion, and practical reasoning. *Philosophical Issues*, **14**(Epistemology), 482-91
- Cohen, S. (2005) Knowledge, speaker and subject. *Philosophical Quarterly*, **55**, 199-212
- Davis, W. A. (1998) *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*. Cambridge: Cambridge University Press
- Davis, W. A. (2003) *Meaning, Expression, and Thought*. New York: Cambridge University Press
- Davis, W. A. (2004) Are knowledge claims indexical? *Erkenntnis*, **61**, 267-8
- Davis, W. A. (2005) Contextualist theories of knowledge. *Acta Analytica*, **20**, 29-42
- Davis, W. A. (2005c) Implicature, ed. E. N. Zalta. Palo Alto, CA. [Http://plato.stanford.edu/entries/implicature](http://plato.stanford.edu/entries/implicature)
- Davis, W. A. (2007) Knowledge claims and context: Loose use. *Philosophical Studies*, **132**(3), 395-438
- DeRose, K. (1992) Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, **52**, 913-29
- DeRose, K. (1995) Solving the skeptical problem. *Philosophical Review*, **104**, 1-52
- DeRose, K. (1996a) Relevant alternatives and the content of knowledge attributions. *Philosophy and Phenomenological Research*, **56**, 193-7
- DeRose, K. (1996b) Knowledge, assertion, and lotteries. *Australasian Journal of Philosophy*, **74**, 568-80
- DeRose, K. (1999) Contextualism: An explanation and defense. In *The Blackwell Guide to Epistemology*, ed. J. Greco & E. Sosa, pp. 187-205. Oxford: Blackwell
- DeRose, K. (2002) Assertion, knowledge, and context. *Philosophical Review*, **111**, 126-203
- Dretske, F. (1981b) *Knowledge and the Flow of Information*. Oxford: Basil Blackwell
- Feldman, R. (1999) Contextualism and skepticism. *Philosophical Perspectives*, **13**(Epistemology), 91-114
- Feldman, R. (2001) Skeptical problems, contextualist solutions. *Philosophical Studies*, **103**, 61-85
- Fogelin, R. J. (2000a) Contextualism and externalism: Trading in one form of skepticism for another. *Philosophical Issues*, **10**(Skepticism), 43-57

- Goldman, A. (1976) Discrimination and perceptual knowledge. *Journal of Philosophy*, **73**, 771-91
- Grice, H. P. (1975) Logic and conversation. In *Syntax and Semantics, 3: Speech Acts*, ed. P. Cole & J. Morgan. New York: Academic Press. Reprinted in *Studies in the Way of Words*, ed. H. P. Grice, pp. 22-40. Cambridge, MA: Harvard University Press (1989)
- Grice, H. P. (1978) Further notes on logic and conversation. In *Syntax and Semantics, 9: Pragmatics*, ed. P. Cole, pp. 113-28. New York: Academic Press. Reprinted in *Studies in the Way of Words*, ed. H. P. Grice, pp. 41-57. Cambridge, MA: Harvard University Press (1989)
- Grice, H. P. (1989) *Studies in the Way of Words*. Cambridge, MA: Harvard University Press
- Hawthorne, J. (2000) Reply to Cohen. *Philosophical Issues*, **10**(Skepticism), 117-20
- Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: Clarendon Press
- Heller, M. (1999) Contextualism and anti-luck epistemology. *Philosophical Perspectives*, **13**(Epistemology), 115-29
- Hofweber, T. (1999) Contextualism and the meaning intention problem. In *Cognition, Agency, and Rationality*, ed. K. Korta, E. Sosa & J. Arrazola, pp. 93-104. Dordrecht: Kluwer Academic Publishers
- Horn, L. R. (1989) *A Natural History of Negation*. Chicago: University of Chicago Press
- Klein, P. (2000) Contextualism and academic skepticism. *Philosophical Issues*, **10**(Skepticism), 108-16
- Kompa, N. (2002) The context sensitivity of knowledge ascriptions. *Grazer-Philosophische Studien*, **64**, 1-18
- Lewis, D. (1979b) Scorekeeping in a language game. In *Semantics from Different Points of View*, ed. R. Bäuerle & et. al., pp. 172-87. Berlin: Springer
- Lewis, D. (1996) Elusive knowledge. *Australasian Journal of Philosophy*, **74**, 549-67
- Ludlow, P. (2005) Contextualism and the new linguistic turn in epistemology. In *Contextualism in Philosophy*, ed. G. Preyer & G. Peters, pp. 11-50. Oxford: Oxford University Press
- Malcolm, N. (1952) Knowledge and belief. *Mind*, **51**, 178-89. Revised version reprinted in *Knowledge and Belief*, ed. A. P. Griffiths, pp. 69-81. Oxford: Oxford University Press, 1967
- Moore, G. E. (1962) *Commonplace Book: 1919-1953*. London: Allen and Unwin
- Partee, B. (2004) Comments on Jason Stanley's "On the linguistic basis for contextualism". *Philosophical Studies*, **119**, 147-59
- Rysiew, P. (2001) The context-sensitivity of knowledge attributions. *Noûs*, **35**, 477-514
- Schaffer, J. (2006) The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies*, **127**, 87-107
- Schiffer, S. (1996) Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, **96**, 317-33
- Stanley, J. (2004) On the linguistic basis for contextualism. *Philosophical Studies*, **119**, 119-46
- Stanley, J. (2005) Fallibilism and concessive knowledge attributions. *Analysis*, **65**, 126-31
- Stanley, J. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press
- Stine, G. (1976) Skepticism, relevant alternatives, and deductive closure. *Philosophical*

- Studies*, **29**, 249-61
- Unger, P. (1975) *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press
- Unger, P. (1984) *Philosophical Relativity*. University of Minnesota Press: 1984
- Unger, P. (1986) The cone model of knowledge. *Philosophical Topics*, **14**, 125-78
- Vogel, J. (1999) The new relevant alternatives theory. *Philosophical Perspectives*, **13**(Epistemology), 155-80
- Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press