

Knowledge claims and context: loose use

Wayne A. Davis

Received: 2 June 2006 / Accepted: 4 July 2006 / Published online: 3 October 2006
© Springer Science+Business Media B.V. 2006

Abstract There is abundant evidence of contextual variation in the use of “S knows p.” Contextualist theories explain this variation in terms of semantic hypotheses that refer to standards of justification determined by “practical” features of either the subject’s context (Hawthorne & Stanley) or the ascriber’s context (Lewis, Cohen, & DeRose). There is extensive linguistic counterevidence to both forms. I maintain that the contextual variation of knowledge claims is better explained by common pragmatic factors. I show here that one is variable strictness. “S knows p” is commonly used loosely to implicate “S is close enough to knowing p for contextually indicated purposes.” A pragmatic account may use a range of semantics, even contextualist. I use an invariant semantics on which knowledge requires complete justification. This combination meets the Moorean constraint as well as any linguistic theory should, and meets the intuition constraint much better than contextualism. There is no need for ad hoc error theories. The variation in conditions of assertability and practical rationality is better explained by variably strict constraints. It will follow that “S knows p” is used loosely to implicate that the condition for asserting “p” and using it in practical reasoning are satisfied.

Keywords Contextualism · Knowledge · Pragmatics · Semantics · Indexicals · Skepticism · Justification · Loose use · Assertion · Practical reasoning

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

W. A. Davis (✉)
Department Of Philosophy, Georgetown University, New North Bldg, Room 215, 37th And O
Streets, NW, Washington, DC 20057-1076, USA
e-mail: davisw@georgetown.edu

variation occurs when knowledge is ascribed to others. Here are two well-known and representative cases, each of which involves two contexts, A and B.¹

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

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 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.² *Contextualism* maintains that the level of justification necessary for

¹ DeRose (1992, p. 913) and Vogel (1999, p. 161). Cohen’s (2000a, p. 95) airport and Hawthorne’s (2004a, p. 69) capital cases have the same features.

² As Stanley (2005, Ch. 5) has stressed, knowledge claims make a temporal reference. Because “knows” is a present tense verb, “S knows p” refers to the time of utterance, which changes in our examples. “S knows p at t” allows reference to any time. However, because the mere difference in time does not explain the variation in our cases, will ignore the temporal reference.

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, pp. 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” 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.³ 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 and Stanley have recently championed subject contextualist views, which do not explain the bank and parking cases in terms of indexicality. Hawthorne (2004a, pp. 158–172) first suggested that the standard of justification necessary for knowledge is determined by the possibilities made salient to the subject by the subject’s context. Hawthorne (2004a, pp. 173–185) later decided that the standard is determined by the choices the subject faces and what is at stake.⁴ Stanley (2005) defends the practical interest view at greater length. The 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. Stanley and Hawthorne, however, reserve the term “contextualism” for ascriber relative views, on which knowledge claims are

³ See Lewis (1979, 1996), Cohen (1986, 1987, 1988, 1999, 2000a, b, 2001, 2004, 2005) and DeRose (1992, 1995, 1996a, b, 1999, 2002). See also Unger (1975, 1984, pp. 46–54, 1986), Stine (1976), Goldman (1976, p. 776), Dretske (1981, p. 367), Heller (1999), Hofweber (1999), Klein (2000), Fogelin (2000), Hawthorne (2004a, Ch. 2), Ludlow (2005) and Schaffer (2006).

⁴ See also Goldman (1976, pp. 772–773), Dretske (1981, p. 367) and Hawthorne (2004a, p. 66, n. 42).

indexical, and call their view “invariantism.” Whereas they wished to stress the differences between their view and indexicalism, I want to emphasize the similarities.

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 contextualism has many counterintuitive consequences. For example, subject contextualism rules that instances of (1) may be true assuming that the ordinary citizens are in contexts with low standards and the experts in contexts with sufficiently high standards.

- (1) Ordinary citizens know the tap water is safe but experts with even more and better evidence do not.⁵

It is absurd to suppose that someone with more and better evidence than one who knows does not. Stanley (2005, pp. 57–62) cites the unacceptability of statements like (1) as evidence against the view that “know” is an indexical like “large,” multiple occurrences of which can have different values in the same sentence. But it tells just as effectively against his own subject contextualist view. The Stanley–Hawthorne view similarly implies that (2) could well be true.

- (2) Hannah knew that the bank would open on Saturday until her husband reminded her that the money had to be deposited by Monday. After this made it more important that she know the bank would open, she no longer knew it.

How could making knowledge more important result in someone’s losing it? Hannah might stop knowing if her justification weakened, or if her belief wavered, but not if the only thing that changes is the importance of knowing (cf. Cohen, 2004, p. 488). Finally, the Stanley–Hawthorne view predicts that counterfactuals like (3) are true in context B. As Stanley (2005, p. 106) notes, such statements are “quite unintuitive.”

- (3) Hannah does not know that the bank will be open on Saturday because the schedule she remembers consulting may be out of date. But she would know if only she were in a context in which it did not matter whether the bank will be open.

We can gain knowledge by remedying deficiencies in our evidence, but not by moving to a context with lower standards of justification.

⁵ We must also assume, of course, that the water is safe, and that expert as well as ordinary citizens believe that it is. Hawthorne (2004a, pp. 160, 167) parries a similar objection from Cohen (1999, pp. 58–59) by observing that “Ordinary people know, but I (we) don’t” is unassertable because of the facticity of knowledge claims and the norm that proper assertion requires knowledge. But the unacceptability of third person statements like (1) cannot be due to similar factors. Hawthorne (2004a, pp. 166–168, 186) concedes that subject contextualism has such seemingly absurd consequences, while suggesting that they might not be as absurd as they seem. Stanley (2005, pp. 125–126) raises some legitimate doubts about Hawthorne’s epistemologist example, but they do not generalize to (1).

Finally, on the subject contextualist hypothesis, we should be able to describe the subject in the bank case as having *lost* knowledge, and as knowing *less* after the change than before. But we cannot. We cannot say that Hannah *used to know*, but *stopped*. Hannah not only would not say in B that she used to know that the bank would be open Saturday, she would say that she did not know it before any more than she does now.⁶ Subject contextualism fails to account for why we hear Hannah's second knowledge claims as *denying* her first. We hear Hannah as retracting her earlier affirmation of knowledge, and conceding that it was incorrect.⁷ Since the subject's context changed in the bank example, contextualism allows the earlier affirmation and subsequent denial to be equally true, even though they seem clearly *contradictory*.

One advantage of ascriber contextualism is that it accounts for the parking case as well as the bank case. The indexicalist theory, moreover, does not predict that (1)–(3) could be true. For example, the speaker who uses (1) is in either a low-standard or a high-standard context. Either way, part of (1) is false.⁸ Similarly, the fact that Hannah is in a low-standard context at one time and a high-standard context later is not relevant to the truth of “Hannah knows that the bank will be open” as asserted by someone else. The relevant standards will remain unchanged unless the ascriber's context changes. That this sentence is true when the ascriber is in one context rather than another does not imply that a counterfactual like (4) is true.

- (4) If the ascriber moved to a low-standard context, Hannah would know that the bank will be open.

For the standard of knowledge applicable to the consequent of (4) is fixed by the context in which (4) is uttered. Finally, the indexicalist correctly predicts that “They used to know but stopped” is false after the standards were raised in our two cases.

Ascriber contextualism has its own unexpected and 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 they 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). Furthermore, if ascriber contextualism were true, argument form (5) would be

⁶ Cf. Cohen (2004, p. 488, 2005, pp. 207–208). Hawthorne (2004a, p. 162ff) concedes this point, and tries to explain it away by postulating a hitherto unnoticed psychological tendency to overproject. That is, he adopts an error theory. See Cohen (2004, pp. 489–490) and Hawthorne (2004b, pp. 520–522) for further discussion.

⁷ Cf. DeRose (1992, p. 924), Schiffer (1996, p. 329), Feldman (1999, p. 103ff) and Stanley (2005, pp. 51–55).

⁸ That is, Cohen and DeRose avoid generating (1) because on their view “knows” does not behave like “large,” which often refers to standards defined by the subjects to which it is applied (small elephants are bigger than large cats). Stanley used the example of “large” to argue that if “know” were indexical, then it should be able to have different interpretations in the two places it appears in (1) (or at least, in a variant thereof in which the predicate after “do not” is made explicit). But Cohen and DeRose can plausibly take the relevant standards to be those prevailing in the speaker's context at the time at which the sentence containing “know” is uttered, rather than the time at which the word “know” is uttered.

invalid in contexts in which the speaker is in a high-standard context and S a low-standard context.⁹

(5) S speaks truly when saying I know p. Therefore, S knows p.

In fact, we understand (5) to be valid in all contexts. Similarly, ascriber contextualism predicts that in context B, Hannah should be willing to say the following:

(6) 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 (6) sounds completely contradictory.¹⁰ Cohen suggests that Hannah cannot assert the second conjunct of (6) in context B because knowledge claims are factive, and in B Hannah cannot say “I now know the bank will be open.” In my view, Hannah can assert the second conjunct of (6) because she still believes in B that “The bank will be open” is true, as facticity requires. However, even if facticity (or perhaps the assertion constraint discussed in the Section “The internalist assertion constraint”) blocks Hannah from asserting (6) in B, it will not block her from asserting it in high-standards contexts in which she can assert “I now know the bank did open.” And as Cohen observed, facticity would not block Hannah from asserting in B a variant of (6) in which “was true” is replaced by “might have been true.” Nor would it block a variant in which “was true” is replaced by “was true I believe.”

In general, we expect instances of the following version of the Tarski-schema to hold except to the extent that “p” is indexical, assuming that both occurrences of “p” have the same sense.

(7) p is true on any occasion of use iff p.

When we substitute sentences of the form “S knows that q” for “p” in (7), we find that failures arise for three reasons: the subject term “S” is a pronoun or other indexical; the verb “knows” is tensed; or the subordinate clause “q” contains an indexical. Except for these reasons, it does not seem intuitively to matter that the context in which such an instance of (7) is used might be different from the occasion on which “S knows that q” is used. Ascriber contextualism implies that it should (Cohen, 2005, p. 206).

Cohen has responded that context sensitive terms like “flat” behave the way “know” does in (5)–(7). Thus he submits that we also hear (8) as valid:

(8) S speaks truly when saying “The road is flat.” Therefore, the road is flat.

⁹ Cf. Hawthorne (2004a, p. 107, fn. 125). Contrast Cohen (2005, pp. 205–206).

¹⁰ Cf. Kompa (2002), Cappelen and Lepore (2003, pp. 29, 31), Davis (2005, p. 39), Bach (2005, p. 60) and Stanley (2005, pp. 54ff, 115, 119ff). Contrast DeRose (1992, pp. 925–926) and Cohen (2005, (205–206)).

Indeed, (8) *can* be interpreted in such a way that it is valid. That interpretation, moreover, is the most natural when (8) occurs out of context as it does here. But unlike (5), (8) does not seem valid in every context of use. Argument (8) can also be interpreted in such a way that it appears patently invalid. To hear such an interpretation, imagine that when S describes roads as flat or hilly, he is comparing them to the roads of the Pyrenees. For him, a road with a 2% grade is flat. Then imagine that we are interested in whether the road is flat enough for an easy bike ride. For us, a road with a 2% grade would cause serious exhaustion. We out-of-shape cyclists can grant that Lance Armstrong spoke truly when he described the mountain road as flat; but it would be foolhardy for us to infer that the road is flat (cf. the Ringle passage below). Similarly, we can grant that the professional basketball coach was correct when he described the 6½ foot center as short. But if we are comparing that 6½ foot center to other ten year olds, we cannot conclude that he is short. “Know” differs from “flat” in that there are no contexts in which (5) is intuitively invalid. Parallel differences would emerge with (6) and (7).

The strict assertion constraint

Hawthorne (2004a) argues that the ascriber contextualist theory is flawed because it does not respect what he calls “the assertion constraint.”

[I]n asserting that p, one represents oneself as knowing that p.... Plausibly, when someone asserts that p, she conveys that she knows that p though she does not assert that she knows that p.... Thus if someone asserts p, it is proper to criticize that person if she does not know that p.... (Of course, that is not to say that knowing p suffices for appropriate assertion. An assertion that p, even if one knows that p, might be out of place for any number of reasons: where it is irrelevant or because its implicatures mislead, or simply because one is in a conversation where p is up for debate.) (Hawthorne, 2004a, pp. 23–24)¹¹

We will take the *strict assertion constraint* to be the simple conjunction of two principles, one descriptive, the other normative:

(9) S’s asserting p conveys that, and is proper only if, S knows p.

This rule accounts for one way in which *asserting* that p is more specific than *saying* that p: the speaker intends to indicate that he knows p, and thus can be criticized if he does not.¹² Hence saying p as a result of a verbal slip, or as a joke, does not count as asserting p. Saying p for either reason does not convey or require knowing p. Williamson (2000, Section 11.3) used the assertion constraint to explain why we are 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 the assertion constraint accounts for “Moore’s

¹¹ Cf. Unger (1975, pp. 250–264), DeRose (1996b, 2002, p. 179ff), Williamson (2000, Ch. 11), Prades (2000, p. 126), Rysiew (2001, pp. 492–493), Stanley (2005, pp. 10–11) and Pritchard (2005). See also Grice’s (1989, p. 27) maxim of Quality, and its two submaxims.

¹² This is qualified in the Section “The variable assertion constraint”. I analyze speaker meaning and related notions in Davis (2003, Part I).

aradox.” 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, p. 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.

The strict assertion constraint entails that the assertability of “*p*” by *S* varies with the truth of “*S* knows *p*.” DeRose (2002) observed that assertability varies with the context as well as knowledge claims. To see this, consider a variant of the bank case in which what Hannah says in context *A* is “*The bank will be open*” rather than “*Yes*” or “*I know the bank will be open*.” Like the knowledge claim, the simple assertion is appropriate in the *A* context but not the *B* context.

Relying on (9), DeRose (2002, pp. 182, 187) uses the contextual variation of assertability as a further argument for contextualism.

Given invariantism about knowledge, the knowledge account of assertion is an untenable attempt to rest a madly swaying distinction upon a stubbornly fixed foundation. Less metaphorically, it is an attempt to identify what is obviously a context-variable standard (the standard for warranted assertion of “*P*”) with what one claims is a context-invariable standard (the relevant truth conditions of “*S* knows that *P*,” according to the invariantist). The knowledge account of assertion demands a contextualist account of knowledge and is simply incredible without it. (DeRose, 2002, p. 182)

For DeRose, however, “contextualism” means *ascriber* contextualism. Differences between ascriber and subject contextualism emerge when second and third person sentences are considered. So imagine that *we* are in a high-standard context. If ascriber contextualism is correct, we should judge that “*Hannah knows the bank will be open*” is false because we are in a high-standard context. But we should also judge that in the *A*-context, Hannah could properly and truly assert “*I know the bank will be open*.” If Hannah could correctly assert that, then by the facticity of knowledge, it had to be proper for her to make the weaker assertion “*The bank will be open*.” Thus we should judge (10) to be true given DeRose’s theory.

(10) In *A*, Hannah may properly assert that the bank will be open but does not know that it will.

This conjunction violates the assertion constraint. Hawthorne (2004a, p. 87) observes that (10) sounds very odd, and uses this as an argument for subject contextualism over ascriber contextualism. If subject contextualism is true, then we have to accept “*Hannah knows in A that the bank will be open*.” This is itself a problem, of course, but a different one (see the Section “The intuition constraint”).

As Cohen (2004, p. 486) realized, the ascriber contextualist must formulate the assertion constraint metalinguistically: *S may assert p only if “S knows p” is true in S’s context*.¹³ This seems to entail (9), but if ascriber contextualism is true, it does not. Given ascriber contextualism, what (9) claims is that *S* may assert *p* only if “*S* knows *p*” is true in the *ascriber’s* context. That is not *S*’s context unless *S* happens to be the speaker.

¹³ Cf. DeRose (2002, p. 182).

Hawthorne (2004a, p. 91) suggests that the ascriber contextualist will be forced to say that assertability claims are as ascriber dependent as knowledge claims. The contextualist in a high-standard context will thus feel pressure to deny that Hannah may properly assert “The bank will be open” in the A-context. But as indicated above, the ascriber contextualist cannot endorse the ascriber dependence of assertability without abandoning facticity. For the contextualist allows that Hannah can properly and truly assert “I know the bank will be open” in A.

Cohen (2004, p. 486) takes a different approach. He allows that (10) is true, and seeks to explain why it nonetheless sounds odd. He suggests that it sounds odd because we mistakenly take “Hannah may properly assert that the bank will be open” to imply that Hannah can represent herself as knowing by the high standards of our context. He might observe that when we do not make this mistake, as in the reasoning described above that leads the contextualist to endorse (10), then (10) does not sound odd at all. Cohen could also observe that “Smith can properly assert that Michael Jordan is tall even though he does not know that Jordan is tall” will sound odd when, as is most natural, both occurrences of “tall” are interpreted relative to the same standard. If we interpret the first occurrence of “tall” as meaning “tall relative to people in general” and the second as meaning “tall relative to basketball players,” the absurdity disappears.

I will agree with Cohen that (10) is literally true, although I will give a slightly different account of why it sounds odd. I will reject Hawthorne’s strict assertion constraint for independent reasons, and replace it with a variably strict constraint.

The internalist assertion constraint

Assertion plausibly indicates not only that the speaker *knows* what has been asserted, but also that the speaker *takes herself to know*. “The bank will be open, but I do not believe I know that” is just as odd as the first-order Moorean paradox. There is therefore an internalist form of the assertion constraint.

(11) S’s asserting p conveys, and is proper only if, S believes she knows p.

It is controversial whether believing one knows without actually knowing ever *suffices* for assertability.¹⁴ But because (11) merely states a necessary condition, it should be little more controversial than (9). But (11) too is problematic given ascriber contextualism. For on the simplest form of that theory, belief statement (12) is not something we should accept.

(12) Hannah believes that she knows the bank will be open.

If the use of “know” makes an indexical reference to the ascriber’s context, then *we* should accept (12) only if Hannah believes she is related to the proposition that the bank will be open by the relation that counts as knowledge in *our* context. For *we* are using “know” in (12), not Hannah. Given the simplest form of ascriber contextualism, (12) does not follow from (13) any more than (14) does.

¹⁴ Williamson (2000, Section 11.5) denies it; DeRose (2002, p. 180) affirms it “in a secondary sense.”

- (13) Hannah believes that “I know the bank will be open” is true in her context.
 (14) Hannah believes that I know the bank will be open.

Statement (14) does not follow from (13) because the “I” in (14) refers not to Hannah, but to me, the speaker using (14).

Whereas (12) would not follow from (13) if the simplest form of ascriber contextualism were true, (12) *does* follow from (13) given how we actually understand the word “know.”¹⁵ We do not understand “know” the way we understand “I.”

Contextualists generally respond to discrepancies between theoretical prediction and linguistic data by postulating that speakers are not aware of the context-sensitivity of “know.”¹⁶ They use this “error theory” to explain why in B Hannah regards her prior knowledge claim as erroneous rather than as true relative to its context. Ludlow (2005, p. 34) and Cohen (2005, pp. 202–206) take a different tack with the problem represented by (12), however. They adopt a mixed contextualist view according to which “know” always refers to the standards of the ascriber’s context when it is in the main clause of a sentence, while sometimes referring to the standards of the subject’s context when it is used in a subordinate clause headed by a speech act or propositional attitude verb.¹⁷ On this hypothesis, “know” behaves differently from the paradigm indexicals “I,” “here,” and “now.” The Ludlow–Cohen hypothesis does not eliminate the problem with (5), though, because “know” is in the main clause of the conclusion. Nor does it eliminate the problem with (6), in which “know” is mentioned rather than used. The Ludlow–Cohen epicycle also causes new problems. For it predicts that instances of semantic descent like (15) should be invalid.

- (15) S’s belief that he knows p is true. Therefore, S knows p.

But as competent speakers understand (15), it is valid. Contextualism provides no explanation as to why.

Cohen uses “tall” to illustrate how “know” is supposed to work. Thus (16) can be interpreted as saying that Suzie believes Michael Jordan is tall by *her* standards.

- (16) Suzie believes that Michael Jordan is tall.

One disanalogy is that “tall” can also be interpreted there as meaning tall by *our* standards. “Know,” however, does not have this kind of ambiguity. There is no interpretation of (15) on which it is invalid. There appear to be no established indexicals that behave the way Cohen and Ludlow postulate.

¹⁵ Cf. Hawthorne (2004a, pp. 98–104) and Stanley (2005, p. 54ff).

¹⁶ Cf. Cohen (1999, pp. 76–79, 2005, pp. 208–209), Hawthorne (2004a, pp. 107–111, 1145) and Stanley (2005, pp. 29–30). Contrast Schiffer (1996), Hofweber (1999), Rysiew (2001, p. 585); Davis (2004, p. 263ff), Bach (2005, p. 67), Stanley (2005, p. 116). I explore the error theory in greater depth in “Knowledge Claims and Context: Belief Difference.”

¹⁷ Following a suggestion of Williamson, Cohen (2005, p. 203) proposes that when we report beliefs about knowledge, we change to a context with the same standards as the subject’s if we know that the subject’s standards differ from our own. But then “Alan’s belief that he knows where his car is false because he does not know whether it was stolen” should be unequivocally false, whereas it seems unequivocally true.

Assertion sufficiency

The strict assertion constraint says that knowledge is *necessary* for proper assertion. In the passage quoted above, Hawthorne (2004a, pp. 23–24) noted that knowledge is not in general *sufficient* for proper assertion. It may not be proper for someone who knows *p* to assert *p* because doing so would be irrelevant, misleading, insulting, or question-begging. All these problems, however, are extrinsic to *S*'s justification for believing *p*. If *S* knows *p*, it is truisitic to say that *S* has enough justification for believing *p* to assert *p*. *S*'s assertion cannot be criticized on the grounds that *S* does not have enough evidence to support it. If *S* knows *p*, it seems plausible that *S* has enough evidence to assert *p* not only in her current context, but in any context. A statement like the following seems absurd. The two conjuncts appear incompatible.

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

“*S* knows *p*” conflicts with “*S* does not have enough evidence for *p*.” We therefore have a further connection between assertion and knowledge, which I shall call *assertion sufficiency*.

- (18) If *S* knows *p*, then *S* has enough evidence or justification for *p* to assert *p* in any context.

Neither subject nor ascriber contextualism validates assertion sufficiency. Subject contextualism, for example, allows that (17) will be true as long as Mary is actually in a low-standards context, where she does not need as much evidence to assert things as she does in courts of law or scientific journals. Ascriber contextualism similarly allows (17) to be true when the speaker is in a low-standards context.¹⁸ The absurdity of (17) and the plausibility of (18) thus support an invariantist semantics.

Pragmatic factors: variable 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, we should consider whether the observed variability can be explained without adopting a contextualist semantics. Variation in use is often due to *pragmatic* 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*.”

¹⁸ Contextualism may predict that (17) will seem false when Mary or the speaker is thinking about courts of law. For that may put Mary or the speaker in a high-standard context, making the knowledge claim false (cf. Cohen, 2004, pp. 487–488). But we do not need to be thinking about particular high-standards contexts to recognize that (24) is false. And (24) sounds not just false, but absurd: the conjuncts conflict.

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*.¹⁹ Implicature is the phenomenon of meaning or implying one thing by saying something else. Conversational implicatures are those that depend on the 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.²⁰

A few have suggested pragmatic accounts of the contextual variation of knowledge claims.²¹ But these have been quickly dismissed or ignored altogether. I reject the prevailing view in linguistics and philosophy of language that pragmatic accounts of linguistic phenomena are *prima facie* preferable to semantic accounts.²² But the breadth and depth of the problems with both subject and ascriber contextualism constitutes powerful evidence that the contextual variation of knowledge claims must have a non-semantic explanation. I will show that the variation in usage can be much more successfully explained in terms of pragmatic factors that operate very widely. One involves implicature.

My multivariate pragmatic theory (Davis, 2004, 2005) attributes 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.”

¹⁹ The theory of conversational implicature was developed by Grice (1975, 1978, 1989). See Davis (2005) for an introduction and references.

²⁰ Be careful not to give “some” in the negation the fall-rise intonation indicative of a “metalinguistic” negation (Horn, 1989, pp. 368; 374–375; 392–413).

²¹ See Unger (1975, Ch. 2, 1984, 1986, p. 141), Prades (2000), Rysiew (2001), DeRose (1999, 2002) and Pritchard (2005).

²² “*Grice’s Razor*” is cited in this connection. See for example Grice (1978, pp. 44–47), Levinson (1983, pp. 97–100, 132), Neale 1992, p. 535) and Bach (2000, pp. 276–277). For further references, and my argument against Grice’s Razor, see Davis (1998, pp. 18–27). Rysiew (2001, p. 505) correctly notes that Grice’s Razor would not apply anyway since the contextualist is not postulating any extra *senses*. He nonetheless argues that a pragmatic account is more economical on the grounds that it uses only “theoretical tools which we’ve clearly got on hand” (2001, p. 505). But the notion of truth conditions varying with the speaker or the subject context is a very standard tool. Cf. Unger (1984, pp. 10, 12).

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.²³ 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 (contrast the Sections “Semantic treatments of loose use” and “Relative knowledge?”).

Some types of sentences are almost always used with some looseness, such as those giving measurements on a continuous scale. Consider:

The Time Measurement Case

A. Wondering how hard the final exam was, I ask Mike how long he took to finish. He answers “Two hours.” B. When Nora says that she took two hours and four minutes to finish the exam, Mike responds “You took even longer than me. It took me two hours and two minutes.”

What Mike said in A is false if what he said in B is true. Yet what Mike undoubtedly meant in A was that it took him *about* two hours to finish—close enough to two hours for the purpose of the conversation at that point, which was just to satisfy my curiosity about how hard the exam was. In context B, the purpose of the conversation required greater precision. For the question there was whether he took longer than Nora. Since Nora’s time was given to the minute, Mike now has to be more precise and give his time to the minute too. We can easily imagine contexts in which even greater precision is required and supplied. But given the imprecision of all measuring procedures, we cannot hope to know the exact duration of any event.

Live examples of loose usage were provided by the Associated Press recently:

The Olympics Case

China has no chance of knocking the United States off the top of the medals table at the 2008 Beijing Games or beyond, a top Chinese official said Thursday.

“It is impossible,” Cui Dalin, assistant sports minister and vice president of China’s Olympic Committee, said. “Elite sports in the U.S. are very strong.” (Stephen Wilson AP Sports, 11/10/05)

²³ Bach (2000, p. 262, 2001, p. 250) uses “speaking loosely” to denote a much more general phenomenon: “omitting words that could have made what we meant more explicit.” One of Bach’s paradigm examples is the use of “Rick and Ann are engaged” to mean “Rick and Ann are engaged to each other.” This is very different from approximation. “Loose use” does not seem apt here.

Dalin surely recognized that there are a number of possible ways in which the Chinese could win more medals, some more likely than others. We would have missed the point, however, if we pedantically insisted that there is a non-zero probability that all the U.S. contestants would have heart attacks, or be disqualified on drug charges. The point was that the chances of such an outcome are negligible: for all practical purposes zero. The possibility of China's winning more medals is too remote to be worth considering. What Dalin said was strictly speaking false, but what he meant may well be true.

Other representative examples of loose use are illustrated by:

- | | |
|--|---|
| (19) I live in Washington, DC.
The water is pure.
My car has been trouble-free.
I played the trumpet. | Colorado is square.
The Subaru costs \$ 20,000.
The gas tank is full. |
|--|---|

I generally tell out-of-towners that I live in Washington even though I actually live in a suburb of Washington just outside the beltway. Most out-of-towners do not know where Springfield is, and are not interested in local jurisdictional boundaries. Loose use thus enables me to convey information that is more useful in a more effective way. If I discover that my audience knows the Washington area, and wants more precise information, I will say “Actually, I live in Springfield, Virginia.” Everyone should be able to illustrate loose uses of the other examples (with the likely exception of the trumpet example discussed in the Section “Semantic treatments of loose use”).

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.²⁴

Terms are used loosely in negative sentences too. Thus after I say “The coffee is all gone” in A, my wife might take a closer look, saying “The coffee isn’t all gone. There’s just enough to make a cup.” My wife is not just saying that the quantity of

²⁴ Compare and contrast Unger (1975, pp. 68–69, 83–87; 1984, p. 6ff) and Rysiew (2001, pp. 487–488).

coffee remaining is not absolutely zero. She is implying that the coffee is *not* close enough to being all gone to count as all gone for the purpose of making coffee.

Others who urge a pragmatic alternative to contextualism focus on related implicatures. Thus Prades (2000, p. 122) suggests that “I know p” is used in low-standard contexts to implicate something like “*There is no need for further investigation.*” Rysiew (2001, pp. 488, 490, 492) suggests “*My epistemic position is good enough*” and “*I can rule out salient alternatives.*” Speakers who use “I know p” loosely may implicate any or all of these propositions too. But these implicatures are not found in all cases, as we shall see. Whether they are depends on the conversational context and the speaker’s more specific intentions.

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. Dick and Jane are trying to answer certain questions. At first they wanted to know whether Alan remembers where his car is; later they wanted to know in addition whether his car may have been stolen. An affirmative answer to their second question requires Alan to be closer to knowing where his car is than an affirmative answer to the first. Because these purposes require knowledge or something close to it, the variable strictness theory accounts for the covariation of knowledge claims and practical interests that Hawthorne and Stanley detail.

The purposes referred to when a term is used loosely are determined by the *ascriber’s* context. This is obvious in the previous cases, in which there is no subject. But consider a variant of the coffee case.

Variant Coffee Case

- A. When the scoop comes up empty, I yell to my wife, “The coffee is all gone.”
- B. After my son announces that he needs a few coffee grounds for his science project, my wife says “Your father believes that it is all gone.”

There are at least two ways of taking my wife’s comment in B (cf. (16)). She may have meant that I believe the coffee is close enough to all gone either for my purposes (making coffee) or for her son’s purposes (the science project). My wife’s comment is thus either irrelevant or inaccurate. What she meant was determined by her intentions. In the same way, when “know” is used loosely, what purposes are referred to are determined by the speaker’s context, not by that of the subject she is describing. Thus in the parking case, what Dick and Jane say about Alan varies with their purposes, not Alan’s. They may have no clue what Alan is up to. In general, we may say that when “S knows p” is used loosely, it is used to implicate that S is close enough to knowing for *contextually indicated purposes*.

When a term used loosely occurs twice in a sentence, the purposes referred to are the same. Thus in (20) we can not hear the first “all gone” as meaning “for my purposes” and the second as meaning “for my son’s.”

- (20) When I was trying to make coffee, the coffee was all gone, but when my son told me what he needed for his project, it was not all gone.

This explains why (1) has no intelligible interpretation even though its two conjuncts could be interpreted as true in contexts where the speaker is referring to different purposes. For similar reasons, we cannot hear even a loose interpretation of (5) on which it is invalid.

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. The Chinese statement quoted above may have been hyperbole, but it could also have been a sober assessment of realistic expectations. 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."²⁵ Hyperbole is a figure of speech. Loose use is approximation.

Loose use must also 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. *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. 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

²⁵ See Partee (2004, pp. 155–157). Compare and contrast Hawthorne (2004a, p. 120).

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.

Another figure of speech resembling loose use is *metaphor*. If my wife asks whether I made caffeinated or decaffeinated coffee, I might answer “It’s high-test.” What I literally said is false, since coffee has no octane rating. What I meant, and therefore implicated, is that it is caffeinated. I also implied that what I made is similar in a certain way to high-test gasoline. There is no suggestion, though, that the degree of similarity is very great. So unlike loose use, I did not imply that the coffee is close enough to being high-test gasoline for our purposes. Everyday metaphors like this one further resemble loose uses in that neither the speaker nor the hearer focuses on the fact that what the speaker said is false. Indeed, my wife is likely to take me to have said that the coffee is caffeinated. Both thus differ from the more poetic metaphors Grice (1975, p. 34) examined, where the hearer, as intended, is initially struck by the obvious falsity of what the speaker says, and has to look for something similar the speaker may have meant that is true. His example was “You are the cream in my coffee.” The line between loose use and metaphor is very fine in some cases, as when we say “Big Blue is trying to capture your queen” or even “It knows where your queen is.” In these cases, the speaker may regard it as debatable whether what he said is literally true or false.

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, p. 39) considered it impossible, but implicatures are often entailed (Davis, 1998, p. 6). For example, if someone asks “Did anyone die?” we could answer by saying “Everyone did”; we would then implicate “Someone died.” When asked “Did you visit Greece or Italy or both this summer?” I could answer affirmatively by saying “I went to Italy.” Pace Cohen (1999, p. 60; 2000b, p. 138), we should not be surprised that the loose-use implicature of “know” is uncancelable. “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.²⁶ So the implicature should be cancelable. Witness: “S doesn’t know that the bank is open on Saturday, but he is close enough to knowing for current purposes.”

Bach (2005) attributes the view that terms like “know,” “flat,” and “empty” are ordinarily used loosely to Unger (1975), but rejects the idea that this involves implicature.

[Unger] was not suggesting that attributors, in saying something literally false, implicate or otherwise convey something true. After all, ordinary folk are not privy to skeptical arguments and do not make their knowledge attributions in

²⁶ Thanks to Cohen for this observation.

defiance of such arguments. Their loose use of ‘know’ is hardly self-conscious or sophisticated enough for them to intend their simple, unqualified knowledge attributions to be taken loosely or for them to recognize knowledge attributions made by others as intended to be taken loosely. Unger’s explanation of why people make simple unqualified knowledge attributions is nothing like a Gricean or pragmatic account of how people can say one thing and mean something else instead. (Bach, 2005, p. 73)

But in the A-context of the coffee case, what I said is something that is true only if there is no coffee at all left in the container. That is not what I meant. What I meant was something less precise. It is because I meant something less precise that what I said is to be construed loosely. I agree that there is no reason to think that ordinary folk mean anything other than what they are saying when they claim to know that they have fingers (see the Section “Other pragmatic factors”). But whereas ordinary folk may not have heard of brains in vats or Cartesian demons, they certainly know about schedule changes and car thefts. The loose uses we have illustrated in the coffee, exam, and Olympics cases involve a level of sophistication nearly universal in competent speakers of the language. Loose use may not involve the “working out” that Grice (1975) thought was characteristic of implicature. But few if any implicatures involve such reasoning (Davis, 1998), including the paradigm case of quantity implicature. The use of “some” to mean “not all” is just as unselfconscious as the use of “all gone” to mean “close enough to all gone.” Recognizing loose use requires no more cognitive ability than recognizing hyperbole, and less than many metaphors.

DeRose (1999, p. 198; 2002, p. 172) and Stanley (2005, p. 14) rightly warn that pragmatic explanations of apparently semantic intuitions can be carried too far. Using Jackson’s (1979) attempt to defend the material conditional analysis of indicative conditionals as a model, DeRose (2002, pp. 172–176) imagines a theory postulating that “know” means “believe” which accounts for the implication that S’s belief is true as a conversational implicature rather than a truth condition. Such an explanation would be completely *ad hoc*. Contrary to what is typical, the hypothesized conversational implicature would be detachable (“S believes p” has no such implicature) and not cancelable (or entailed). DeRose suggests that a pragmatic theory like mine would be similarly *ad hoc*.

To the extent that defenders of invariantism go beyond such bare maneuvers, their hints point in the direction of *special* rules for the assertability of ‘knows,’ such as “If someone is close enough, for present intents and purposes, to being a knower, don’t say that she doesn’t know, but rather say that she knows.” (DeRose, 2002, p. 176)

On the contrary, I am accounting for the observed behavior of “know” in terms of the *general* phenomenon of loose use. I use the close similarity between cases of contextual variation in the use of “know” and clear cases in which diverse terms are used with variable strictness as evidence that the variation is due in part to its being used loosely or strictly. The implicature theories of Prades and Rysiew may seem *ad hoc* in the way DeRose envisions. Nothing similar applies to the coffee or time measurement cases, for example. But the implicature I cite is

typical of loose use. DeRose (1999, p. 200) knew of no case in which what a speaker says appears true just because of a true implicature. Loose use provides countless examples, as do everyday metaphors.²⁷

DeRose goes on to say that

No defender of invariantism has proposed an account on which general rules of conversation together with their proposed invariantist account of the truth conditions of knowledge attributions predict the pattern of varying conditions of warranted assertability of knowledge attributions that we encounter in natural language. (DeRose, 2002, p. 176)

Stanley (2005, p. 15) similarly insists that we should not accept an explanation of apparently semantic intuitions in terms of implicatures “in the absence of a clear explanation of these intuitions from general conversational principles” (2005, p. 15). I presume DeRose and Stanley are here relying on Grice’s well-known “Calculability Principle,” according to which “the presence of a conversational implicature must be capable of being worked out” (Grice, 1975, p. 31). “Working out” an implicature involves deriving it from the meaning of an expression together with conversational principles, which for Grice were the Cooperative Principle and its maxims (Quantity, Quality, Relevance, and Manner). I have argued at length that *no* conversational implicature is capable of being derived from general conversational principles (Davis, 1998). 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.²⁸ 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 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.²⁹

²⁷ I reject the rule DeRose foists on the invariantist. If I *meant* in A that the coffee is literally all gone, then I have simply made a mistake, even if in fact it is close enough to being all gone. Conversely, if I wanted to point out that the coffee is not literally all gone, I could correctly have said “The coffee is not all gone” if that is what I meant.

²⁸ Rysiew’s (2001, p. 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, pp. 197–200) attempt to derive the “don’t know” implicature of “possible” from a variant of the Maxim of Quantity.

²⁹ Contrast DeRose (2002, p. 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.

An invariantist semantics

DeRose's charge against the invariantist proposing a pragmatic account of the contextual variation of knowledge claims suggests that such an account would have to provide an analysis of knowledge. There is no requirement that meaningful terms be definable, however, 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. However, a pragmatic account of the bank and parking cases in terms of variable strictness would work just as well if something weaker replaced "completely," such as "nearly completely" or "highly." 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, p. 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. In my view, though, 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 in this paper, 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.

I assume that what counts as complete justification does not vary from context to context.³⁰ But even if it does, one could still account for the variation in the bank and parking cases in terms of variably strict usage. As long as the range of variation

³⁰ Contrast Dretske (1981, pp. 363, 366–367, 370).

displayed in the bank and parking cases is outside that allowed by complete justification, Hannah could be saying in A that she is close enough to having a fully justified belief for current purposes, while denying this in B. So while the pragmatic account of the bank and parking cases I have proposed goes naturally with an invariantist semantics, it is not strictly necessary.³¹ We can expect, though, that even limited context-sensitivity would lead to the sorts of problems we have noted.

Semantic treatments of loose use

Cohen has suggested treating all my examples of loose usage as cases of indexical quantifier restriction or parameter specification. In the coffee case, for example, he suggests that the domain of quantification in A includes only quantities of coffee sufficient to make a cup, while in B the domain is widened to include smaller quantities. He infers that what I said in both contexts is true; the usage in the A-context is indexically restricted, not loose. Since Cohen takes quantifier restriction to be an instance of indexicality, he takes the similarity between the coffee and time-measurement cases and the bank and parking cases to support ascriber contextualism.

Some would argue that quantifier restriction is itself a pragmatic phenomenon, much the way Kripke (1977) argued that the distinction between the referential and attributive use of definite descriptions is an instance of the distinction between speaker meaning (pragmatic) and sentence meaning (semantic).³² But I am inclined to believe Cohen is right about quantifier restriction being semantic. When I proudly told my logic class “Everyone got an A,” I was talking about students in that class. I did not say that everyone in existence got an A. What I said, and meant, is that everyone in that class got an A. When I walk into my epistemology class and say “Not everyone got an A,” talking about this class now, I am not contradicting or retracting what I said to my logic class. In neither context do I speak loosely. If the domain of quantification included every person in existence, my first statement would be grossly false, not approximately true. And since quantifier restriction concerns reference, it determines what is said rather than what is implicated.

While *prima facie* plausible, Cohen’s suggestion that the coffee case involves quantifier restriction must ultimately be rejected for a variety of reasons. First, this indexical treatment produces many of the same counter-intuitive results we found with indexical treatments of “know.” One is the consequence that what I said in the B context does not contradict what I said in the A context. The coffee case differs markedly from the grade case in this respect. We do not hear what I said to my epistemology class as a denial, retraction, or correction of what I said to my logic class, whereas we do hear what I said after my son made his request in the coffee case as a denial, retraction, or correction of what I said before. Second, if context A

³¹ Stanley (2004, Section 5; 2005, p. 82) correctly observes, though, that the *contextualist* cannot account for the contextual variability of knowledge claims in terms of loose use, since loose use is a pragmatic phenomenon and contextualism is a semantic theory. See also Kompa (2005, pp. 23–24).

³² Cf. Atlas (1977, 1979, 1989), Sperber and Wilson 1986, pp. 182–183), Carston (1988) and especially Bach (1994, 2000, 2001). Contrast Stanley and Szabó (2000).

did restrict the domain of quantification as Cohen suggests, it should be natural or at least possible to interpret other sentences with the same restriction. It should be possible to interpret “No quantity of coffee is less than scoop-sized” as true in A. But this seems impossible given that quantities of coffee are divisible. (Such an interpretation would be possible in a context in which all the coffee is divided into piles that are scoop-sized or larger.) Similarly, it should be possible for my wife to contradict me by saying “At least one quantity of coffee is left.” Yet this would be uninterpretable in A (but okay in the divided pile context). In contrast, when I tell my class “Everyone got an A,” talking exclusively about students in that class, other quantified statements like “Everyone did the assignment,” and “No one is absent” are naturally interpreted with the same restriction. “Everyone is in this class” would not naturally be interpreted with such a restriction because we tend to avoid vacuity, but such an interpretation is at least possible. Finally, the notion of domain restriction does not apply naturally to “The coffee is all gone” because “coffee” is functioning as a mass term rather than a count noun, “the” is a determiner, and “all” means “100%” rather than “every.” “Every quantity of coffee is gone” is not an exact translation. If we switch to an example involving count nouns, all plausibility vanishes. Britain was overrun by millions of American soldiers in the build-up to D-Day. When Margaret’s town emptied out after the invasion, she said “The Yanks are all gone.” She may have meant this loosely, not caring whether one or two might be left. Her point was that there has been a big change. If even one Yank was left, however, then what she said was clearly false even though what she meant was true.

Loose use and quantifier restriction are independent phenomena. No matter what domain is specified in the grade and Yank cases, the speaker’s statement can be interpreted loosely or strictly. Domain restriction determines what the speaker said; whether the speaker is speaking loosely or strictly depends on what the speaker meant by saying that. Thus “Everyone got an A” is used loosely if what is meant is that close enough to everyone in the indicated set of people got an A, whether that be students in my intro class, or students I taught last semester. What is said is false if even one member of the indicated set got a B. With either domain, the sentence might naturally be used loosely by a dean complaining that I was contributing to grade inflation. Unless the dean happened to know who the exceptions were, there is no way for her to restrict her quantifier to the same effect.

In the time case, Cohen suggests that the measurement has an implicit “ $\pm\Delta t$ ” that is narrower in B than in A. That is, he takes “I took two hours” to be true iff the speaker took $2 \pm \Delta t$ hours, where Δt is contextually specified. If this were true, then “I took two hours” would not be incompatible with “I took Δt more than two hours,” as it is in every context for any $\Delta t > 0$. Moreover, Mike could have been speaking loosely even if his statement had an implicit “ ± 1 minute.” That is, he may have been approximating in A even if what he said were that he took two hours plus or minus one minute. What Mike would have meant in that case is that the amount of time he took on the exam was close enough to that range for current purposes. He may have been speaking more strictly in B even if what he said were “I took two hours and four minutes plus or minus a minute.”

It is also unlikely that speakers have any precise Δt in mind. Speakers more plausibly mean something vague like “I took *about* two hours,” where the range allowed by “about” varies from context to context. But the sentence “I took two hours” itself does not mean or have the truth conditions of “I took about two hours,” even on a particular occasion of use. The former entails the latter, but not

vice versa. The latter is compatible with “I took slightly more than two hours,” the former is not. The speaker may have meant “I took about two hours,” but that is not what he said.³³

Similar problems arise if the coffee case is treated as involving parameter rather than domain restriction. Cohen might suggest that what I meant is that the coffee is $100 \pm \Delta p\%$ gone, with Δp small in A and smaller in B; or that the coffee is *nearly* all gone. But then what I said in A would not be incompatible with “The coffee is slightly less than all gone.” It is unlikely that I had any precise Δp in mind. And with any Δp , the resulting statement could be used loosely or strictly.

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. the Section “The moorean constraint”). That conclusion would be absurd indeed, but it does not follow. I can know that the stick is *about* 3 ft. long, and that suffices to know how long it is. What I cannot know is *precisely* how long the stick is, with no margin of error. But knowing that is not necessary for knowing how long it is. Furthermore, even though “I know the stick is 3 ft. long” is strictly speaking false, I can nonetheless use it to imply that I know how long the stick is by using it loosely.

Finally, the bank and parking cases also resemble examples of loose use for which semantic treatments are not plausible at all. When I am talking to people who are not very knowledgeable about music, I generally say that I played the trumpet in school. When I am talking to people who know brass instruments, I will tell them that I played the cornet, not the trumpet. No quantifier restriction or parameter widening makes “I played the trumpet” true in any context. It is just false, although for most purposes it is close enough to being true.

Relative knowledge?

My loose use account presupposes that whether someone is close enough to knowing something depends on one’s purposes. Contextualism, in contrast, implies that whether someone knows is dependent on purposes (cf. Hawthorne 2004a, p. 66, fn. 42). The conventional usage of “know” in English does not confirm this prediction. Compare:

- (21) (a) Does Steve have enough information? That depends on what he is doing.
 (b) Did Steve ruin his device? That depends on what he uses it for.
 (c) Does Steve know p ? That depends on what he is trying to do.

Whereas the responses in (21)(a) and (b) are perfectly appropriate, the response in (21)(c) seems uninterpretable. I would expect “Huh?” from the questioner.

Cohen offers “flat” as a model of the sort of indexicality “know” would possess if his theory were true.

³³ Scientists have a convention whereby the number of digits after the decimal point indicates the level of precision intended when measurements are reported. By this convention, “The stick is 3.00 ft. long” implies that it is $3.000 \pm .005$ ft. long. Using this convention, “The yardstick is 3.00 ft long” is not equivalent to “The yardstick is 3.0 ft. long,” as it is in conventional English.

Many, if not most, predicates in natural language are such that the truth-value of sentences containing them depends on contextually determined standards, e.g. ‘flat’, ‘bald’, ‘rich’, ‘happy’, ‘sad’.... These are all predicates that can be satisfied to varying degrees and that can also be satisfied simpliciter. So, e.g., we can talk about one surface being flatter than another and we can talk about a surface being flat simpliciter. For predicates of this kind, context will determine the degree to which the predicate must be satisfied in order for the predicate to apply simpliciter. So the context will determine how flat a surface must be in order to be flat. (Cohen, 1999, p. 60)

One major disanalogy is that “know” does not admit of the range of comparative and superlative modifiers that we find with “flat” (Stanley, 2005, Ch. 2). There is nothing comparable to “flat, flatter, flattest,” “very flat” versus “perfectly flat,” or “flat compared to” We never say “*He knows it’s true, but not perfectly,*” and would not know how to interpret it. Adverbials of degree that can modify other propositional attitude verbs, such as “to some/a great/a greater degree (or extent),” cannot modify “know.” “Much,” “more,” and “most” can be used as adverbs to compare the amount of knowledge someone has, but are not used in “S knows ___ p” to compare the degree of knowledge. The comparative adverbs “well,” “perfectly well,” and “better than” can be used in this context, but denote directness of the evidence rather than level of justification (Dretske, 1981, n. 1). If you have run marathons, then you know better than I do that they are hard. I am nonetheless completely justified in believing that. When applied to the adjective “known,” these adverbs compare how widely something is known. “Knows” cannot even be modified by “more/most intensely,” which can modify many verbs that do not admit “more/most” in the same sense, such as “laugh” and “hit.” Cohen (1999, p. 60) takes “know” to make an indexical reference to a level of justification. Then we should expect “*He knows p with less justification than she does*” to be semantically as well as grammatically acceptable English. But we never say things of this form, and they strike me as not making sense—in marked contrast to what results when “believes” replaces “knows.”

To support the thesis that “S knows p” has an implicit argument function for standards of knowledge, Ludlow (2005, pp. 20, 30–31) observes that “know” occurs with adverbial qualifiers like “by objective/scientific/high/today’s standards.” But “S knows by subjective/unscientific/low/outdated standards” all sound self-contradictory. Moreover, we never make contrasts like “*He knows by low standards, but not by high standards.*” Stanley (2005, p. 70) observed that these qualifiers are not attached grammatically to the word “know.” They serve the same function in sentences without “know,” as in “By strict standards, France is not hexagonal” or “By my standards, he played the sonata perfectly.” Whenever there are differing standards by which people say or judge that p, we can say “By X standards, p.” We use this formula to explicitly state that X standards have been met. In some instances, there are different standards for saying or judging that p because “p” makes an indexical reference to standards; an example is “By the .05 standard, the results are statistically significant.” In other cases, however, there are different standards because “p” is used with variable strictness (see the Section “Pragmatic factors: variable strictness”), or because there are differences of opinion on the subject (see the Section “Other pragmatic factors”). Consider, respectively, “By wood working/road building standards, the surface is perfectly flat” and “By religious/scientific

standards, we are completely justified in believing that there are supernatural beings.” From the meaningfulness of such sentences, we cannot infer that “perfectly flat,” “completely justified,” or “entirely sober” are context-sensitive.³⁴

Further evidence against “know” having the relativity of terms like “flat,” “tall” or even “moving” is that we never mark contrasts by saying:

(22) He knows relative to this, but not that. It is knowledge for him, but not her.³⁵

Moreover, nothing like the following discourse is possible with “know.”

It’s a Mountain, Not a Molehill

“Tiny”?

Please tell Jonathan Yardley that Mount Marcy, at 5,344 feet above sea level, is the highest peak in New York [Style, Nov. 1]. We New Yorkers might take “tiny” from the Dalai Lama, who probably thinks in Mount Everestian terms, or from Sen. Ted Stevens (R-Alaska), who can trump us with Mount McKinley (20,320 feet).

But we can’t abide it from Yardley, a native of Pennsylvania, where they’re awestruck by their tallest: Mount Davis – a hillock, a pimple of a mountain, at 3,213 feet.

–William Ringle, *McLean* [Washington Post, 11/05/05]

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 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 woodworker 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.

³⁴ Ludlow also observed that “knows” is sometimes qualified by “with some/complete confidence.” However, “He knows without complete confidence that p” sounds contradictory; “S knows with complete confidence that p” sounds like a redundancy used for emphasis; and “S knows with some confidence that p” is most naturally interpreted either as an understatement or as an overstatement. In the latter case, we hear it as implicating that S does not really know p. Note that “Hannah knows with some confidence that the bank is open on Saturday” is no more assertable in the B-context than its unqualified root is. Even more often, “know” appears with “for sure.” I believe the latter has the same “parenthetical” function it has when it occurs without “know,” as in “I’ll be there, for sure.” We never say “*He knows that p, but does not know for sure that p.*”

³⁵ Cf. Schiffer (1996, pp. 326–327), Hofweber (1999, p. 95ff), Schaffer (2006, Section 3). Cappelen and Lepore (2003, pp. 39–42) and Davis (2004, pp. 267–268) observe in different ways that “S knows p in context c” does not have the requisite interpretation either. “In c” functions as an adverbial of place, not as a sentence operator (or “monster”).

Hawthorne argues that “know” is disanalogous to such absolute terms. Whereas we can say that a surface is perfectly flat iff nothing could possibly be flatter, and a jar completely empty iff nothing could be emptier, no similar definitions are possible for “S knows p.” In saying this, Hawthorne intends to go beyond the linguistic observation that there is no comparative form of “know” similar to “flatter” or “emptier.”

Speaking with maximal abstractness: mastery of the comparative use of a comparative adjective requires a conceptual grip on some kind of scale. Where the scale has an upper bound, it will provide the basis for a concept that applies to a thing iff it is at that upper bound.... The question at hand is whether our mastery of ‘know’ has a scale of the required kind as a conceptual backdrop; and with no comparative use of ‘know’ to rely on, it is very much an open question whether our mastery of ‘know’ is dependent on a scale, let alone one with an upper bound. (Hawthorne, 2004a, pp. 138–139)

My pragmatic account of the contextual variability of “know” does not presuppose that we can define “know” in terms of such a background scale. It is enough that we have a conception of when “know” strictly speaking applies to a situation, and a belief that other situations resemble that one more or less closely. Since knowledge involves true justified belief, there are at least three dimensions along which someone may be more or less close to knowing.

While the determinant of closeness is nearness to the truth in some cases (e.g., the variant time measurement case in the Section “The intuition constraint”), and degree of belief in others, the relevant scale governing the variably strict usage of “know” in the bank and parking cases is the subject’s level of *justification* for the belief p.³⁶ S knows p only if S is completely justified in believing p. The more complete S’s justification for believing some truth is, the closer S is to knowing it. In context A, Hannah’s knowledge claim was strictly speaking false, because she was not completely justified in believing that the bank would be open on Saturday. Her justification left much room for doubt. However, she had more than enough justification for their purposes in context A, so it was appropriate for her to use “know” loosely. When the conversational purpose changed, more justification was required; so she used the term more strictly. Note that “completely” and “fully justified” can themselves be used loosely or strictly. They could be affirmed loosely in the A-contexts of our cases, while denied in the B-contexts.

Lewis (1979, p. 355) observed that the standards for knowledge claims are more easily raised than lowered.³⁷ It is hard to find natural cases in which the B-context precedes the A-context. This “downward stickiness” is not what we would expect if contextualism were correct. The standards for claims of statistical significance, for instance, can go down as easily as up. The domain restriction on quantifiers moves either way (Stanley, 2005, p. 65). In the case of “flat,” we regularly and smoothly move from woodworking contexts to driving contexts, which takes us from high standards of flatness to low standards. We cannot identify similar transitions with “know” because we cannot in the same way describe contexts in which “know” is

³⁶ Cf. Cohen (1999, p. 60), Rysiew (2001, pp. 488, 491). Stanley (2005, pp. 75–77) takes Cohen to be arguing that the context-sensitivity of knowledge can be inferred from that of justification; that is not what I am arguing.

³⁷ See also DeRose (1995, p. 8, fn. 12).

used relative to low standards and others in which it is used relative to high standards. “Everyday contexts” does not identify a standard, because the B-contexts in our examples are just as everyday as the A-contexts.

Cohen accounts for downward stickiness in terms of his hypothesis that standards tend to rise when error possibilities become salient. This would account for the phenomenon given that once something has become salient in conversation, it tends to stay salient. What now needs to be explained is why standards tend to rise in this way. Making error possibilities salient does not tend to make the significance level rise in experimental studies. Why should it make standards of justification rise?

The downward stickiness of knowledge claims does conform to the general pattern for variably strict usage. When terms have been used loosely, it is easy to raise the standard to offer or demand greater precision. But once terms have been used strictly, it is hard to use them loosely without being misunderstood. Thus once we raise the standard for coffee to be all gone, we would naturally use high standards in describing the sugar or the milk as being all gone. To revert to a low standard again, we would have to say something like “just roughly” or “for all practical purposes” or “compared to the cookie jar.” Downward stickiness occurs even when terms are used strictly in negative sentences. If someone says “Hannah does not weigh 151 pounds,” a statement like “Hannah weighs 150 pounds” would be naturally taken more strictly, even though as a conversation starter it would just as naturally be interpreted loosely. Similarly, someone focusing on Michigan or Florida would naturally describe Colorado as “perfectly square,” at least until someone observes that it is longer from east to west than it is from north to south. Once that observation has been made, it is very hard to interpret “Wyoming is perfectly square” loosely, even though it is significantly closer to being perfectly square than Colorado is. So the downward stickiness of knowledge claims provides further support for the variable strictness account.

The fact that “know” is disanalogous in many ways to known relative terms, and behaves more like an absolute term, is not a conclusive argument against contextualism. It could be an exceptional relative term. But the disanalogies cited are evidence against the theory. The fact that the pragmatic theory can explain the contextual variation in the use of “know” in terms of a well-established, general phenomenon (loose use) without postulating that “know” is a *sui generis* relative term is, by standard scientific methodology, a strong argument in its favor.

The variable assertion constraint

We saw in the Section “The strict assertion constraint” that subject contextualist theories could reconcile the contextual variation of assertability with the strict assertion constraint (9). For (9) entails that the assertability of “p” by S varies with the truth of “S knows p.” If knowledge varies with the subject’s context as contextualist theories claim, then so does assertability. The variable strictness theory of knowledge claims does not similarly imply that assertability varies from context to context. For even though speakers can use “knows” loosely, the literal meaning of the term is the strict meaning. So (9) ties assertion and assertability to knowledge strictly speaking.

A variant of (9), however, is at least as plausible, and would imply that assertability varies with the same factors that influence knowledge claims on the variable strictness theory. I call it the *variable assertion constraint*.

- (23) S's asserting p conveys that, and is proper only if, S is close enough to knowing p for current purposes in S's context.³⁸

Independently of any theory about whether knowledge is context sensitive, the variable assertion constraint implies that assertability will vary with the purposes of the asserter's context. This is exactly what we found in the variant of the bank case. Given the variable strictness factor, the pragmatic theory also predicts that knowledge claims will vary with the purposes of the context, in exactly the same way. For "I know p" is used to implicate that the conditions for the speaker to properly assert "p" are satisfied. It is not that assertability is variable because knowledge is context sensitive, on my view. It is that knowledge claims and assertability vary with the same independent factors because of what "I know p" is used to implicate.

The variable assertion constraint differentiates asserting from other ways of saying just as effectively: Saying that p counts as asserting p only if the speaker intends to indicate that he is close enough to knowing p, and thus has misled his audience if he is not. Similarly, we can explain Moore's paradox just as well given the variable assertion constraint. For asserting "I don't know p" implicates that I am not close enough to knowing p for the purposes at hand, which contradicts what is conveyed by my asserting "p." The variable assertion constraint also accounts for the observation that "How do you know?" is always an appropriate response to an assertion. For even if its presupposition is strictly speaking false, the question "How do you know?" is appropriate in any context in which the addressee is close enough to knowing to count as knowing. Finally, given our ability to use "know" loosely, sentence (9) can naturally be used to express what (23) asserts. Thus the pragmatic theory of knowledge claims together with the variably strict assertion constraint can account for the oddity of (10) at least as well as subject contextualism. To force a consistent interpretation of (10), we have to deviate from the conventional practice governing purpose selection in loose use.

Support for the variable assertion constraint *over* the strict form is provided by the intuitive acceptability of (24).

- (24) In A, Hannah properly asserts that the bank will be open but does not strictly speaking *know* that it will.

The addition of "strictly speaking" plus the emphasis on "know" makes (24) perfectly acceptable. The strict assertion constraint predicts that (24) should seem even worse than (10). Similarly, people readily assert that their yardsticks are three feet long. And in everyday contexts, these assertions seem a model of propriety. Yet

³⁸ Hawthorne (2004a, p. 134) mentions this possibility, but instantly dismisses it without argument. Pritchard (2005) explores Grice's maxim of evidence ("Do not say that for which you lack adequate evidence"), which is the implication of (23) relevant to the bank and parking cases. Note that I am not tying the assertability of "p" to the assertability of "S knows p," the view DeRose (2002, p. 182) rejects. What (23) implies is that "S knows p" is assertable only when S is close enough to *knowing* that S knows p.

what they said is false if the yardstick is either more or less than three feet long—by any amount, no matter how small. So what they are asserting has a precision that exceeds our ability to measure. While they are close enough to knowing for everyday purposes, they strictly speaking do not and cannot know that their yardsticks are three feet long. In addition to lacking sufficient justification for a precise belief, it is likely to be false. So there is good reason to believe that the strict assertion constraint is too strict.³⁹ Note that the yardstick case is one in which subjects may be close enough to knowing something even though their belief is false, because it is close enough to being true.

Williamson's primary argument for the strict assertion constraint is that it accounts for our disinclination to flat-out assert that someone will lose a lottery if one's only evidence for that outcome is its high probability. This disinclination is not absolute, however. The lottery case described below (In the Section "The intuition constraint") involves a career dishwasher named Harry who is seriously considering buying a million dollar Ferrari, and is counting on winning the big lottery afterwards to pay for it. For the purposes of deciding whether Harry should buy the Ferrari, Joe is close enough to knowing that Harry will lose. Consequently Joe could appropriately assert either "You're not going to win any lottery" or "You know you'll lose."⁴⁰ This is true even though Joe does not strictly speaking know that Harry will lose. So this is another example in which the strict assertion constraint is too strict.

The practical rationality connection

Hawthorne (2004a, pp. 29–31, 87–89, 147) and Stanley (2005, pp. vi, 12–13) postulate that knowledge is normatively connected to practical rationality as well as assertibility. Neither develops the point in detail, but the following formulation is suggested by what they do say.

(25) S may use p as a premise in practical reasoning iff S knows p.

Hawthorne argues that it would be unacceptable for Sheila to conclude that she should sell her lottery ticket for a penny from the premise that the ticket is going to lose, if her only reason for believing that premise is the low odds of the ticket winning. In contrast, if Sheila knew the ticket was going to lose, perhaps because she learned how the drawing was going to be fixed, then her practical reasoning would be perfectly acceptable.

Stanley and Hawthorne use the practical rationality constraint to support subject contextualism. They observe, for example, that in context A it was rational for Hannah to put off going to the bank until Saturday. Since that decision was based on the belief that the bank would be open on Saturday, they infer that Hannah's belief must constitute knowledge in context A. But in context B, it would be irrational for Hannah to act on the same belief. So in B, that belief must not count as knowledge.

³⁹ Williamson (2000, 256–258) raises this problem, but does not, I believe, satisfactorily address it. Another counterexample arises when agents properly assert that they will do something because they fully intend to do it, but do not *know* they will.

⁴⁰ Unger (1986, p. 145) described a similar case.

Hawthorne and Stanley appear to claim that knowing “p” is both necessary and sufficient to use “p” in practical reasoning. The sufficiency claim is too strong in one respect, too weak in another. Knowing “p” does not ensure that “p” is relevant to the practical conclusion drawn, nor preclude begging the question. Knowledge at most ensures adequate evidence or justification. On the other hand, if S really knows p, then S would seem to have enough evidence to use “p” as a premise in *any* practical reasoning, not just the reasoning S happens to be engaged in. This gives us a parallel to assertion sufficiency:

- (26) If S knows p, then S has enough evidence to use p as a premise in any practical reasoning.

Hawthorne objects to ascriber contextualism because it denies the right-left conjunct of (25). “S knows p” might be true in the speaker’s context because its standards are low, while “S can use ‘p’ as a premise in practical reasoning” is false in any context because S does not have enough evidence for the truth of “p” given the importance or the nature of the reasoning S is engaged in. This is problematic because “S fails to have enough evidence that p” seems incompatible with “S knows p” in any context. Cohen (2004, 486–487) acknowledges this problem, offering several considerations to soften the blow. One is that subject contextualism has similar problems. It allows, for example, that “S knows p” might be true because S happens to be in a low-standards context, even though “S can use ‘p’ as a premise in any practical reasoning” is false because S does not have enough evidence for the truth of “p” to engage in the sort of practical reasoning that would move S into a high-standards context. Cohen focuses on the objectionable implication (noted earlier) that S would stop knowing in the high-standards context despite being in the same evidential position.⁴¹ The current objection is the implication that S does not have enough evidence despite knowing. We see again that contextualism in general is inconsistent with knowledge sufficiency.

The claim that knowledge is *necessary* for rational practical reasoning is questionable for different reasons. First, a popular alternative is provided by decision theory, which was designed to handle decision making under uncertainty. If we suppose that the degree to which S believes p is proportional to the degree to which the evidence supports the proposition p (i.e., proportional to the objective probability), and that the degree to which S desires an outcome O is proportional to the desirability or value of O, then it is at least plausible that S’s action is rational provided the “subjective expected utility” of the action is maximal. This implies that if the outcome of action A given p is O and the outcome given not-p is O’, then the greater the difference in value between O and O’, the more evidence S will need as to whether p or not-p before deciding whether or not to do A. This adequately explains why Hannah needs more evidence in the B-context of the bank case.

⁴¹ Cohen (2004, p. 487) also observes in mitigation that speakers will not be able to say “S knows p but cannot use ‘p’ in reasoning” when they are thinking of particular instances of illicit reasoning. For when thinking about such examples, speakers will be in a context in which “S knows p” is false. But as Cohen concedes, the statement may be true when the speaker is not thinking about such examples, and is nonetheless objectionable.

Second, the data Hawthorne and Stanley provide in support of their constraint can be handled equally well by an alternative that parallels the variable assertion constraint (see the Section “The variable assertion constraint”).

(27) S may use *p* as a premise in practical reasoning only if S is close enough to knowing *p* for S’s purposes.

Thus in context A, Hannah could rationally put off going to the bank until Saturday because she was close enough to knowing that the bank would be open given what was at stake (convenience). She had enough evidence. But in context B when more was at stake (bouncing checks), Hannah was not close enough to knowing. She no longer had enough evidence. In both contexts, she needs more than a mere belief. Similarly, it would be irrational for Sheila to sell her ticket for a penny not just because she does not know that it is going to lose, but because she is not even close enough to knowing for her purposes.

Third, suppose Sheila were offered not a penny, but ten thousand dollars for her ticket. If she happily sold the ticket because she believed it was going to lose, would she be acting irrationally? Not in my book. Or imagine that I have good reason to believe that a student is coming to see me at noon. Students being what they are, I cannot know that the student is going to show up. Nevertheless, it seems rational, indeed obligatory, for me to act on my belief and be in my office at noon. Sheila is close enough to knowing that her ticket will lose, and I that my student will show up, for the purposes of making these decisions.

Finally, there is a question about justification that is begged by (25) but left open by (27): the issue of externalism versus internalism. Calculators these days are incredibly reliable. Suppose S uses a brand that has never before malfunctioned, and acts on the results. Suppose further that despite being well justified, one of the beliefs produced by the calculator is false. Did S act improperly simply because the belief is false? The Stanley–Hawthorne rule (25) says S did act improperly, and thus rules for a very strict externalism. Rule (27) implies that whether S acted improperly depends on whether the truth of S’s belief was required for it to be close enough to knowledge for the purposes at hand. If the error was in the third decimal place, and S’s purposes only required precision to the second place, it is especially hard to argue that S acted irrationally.

Even though “S knows *p*” does not strictly speaking express a necessary condition on rational practical reasoning, it may be used loosely to say that the variably strict practical reasoning and assertion constraints are satisfied. These are reasons why the loose use implicature is important for us.⁴²

⁴² Cf. Cohen (2004, p. 487): “[T]o say that S knows P is to say that the probability of P on his evidence is close enough to 1 for S to appeal to it in reasoning.” On ascriber contextualism, however, this claim is false. The truth of “S knows *p*” depends on the standards of the speaker’s context, not the subject’s. “S knows *p*” may be true in the speaker’s low-standard context even though “The probability of ‘*p*’ on S’s evidence is close enough to 1 for S to appeal to it in reasoning” is false because S is in a high-standard context.

The Moorean constraint

Some reject a pragmatic account of the contextual variability of knowledge claims on the grounds that it endorses skepticism.⁴³ Hawthorne (2004a, pp. 116–118), 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. Hawthorne then casts doubt on whether the theory is coherent. If “I know p” is always false, he asks, “can I even be reckoned to *assert* the falsehood that is posited as the semantic value?” Hawthorne may be relying on his assertion constraint here. Indeed, if “I know p” is always false, then arguably saying “p” would not *convey* knowledge that p, and so on Hawthorne’s view would not count as asserting that p. This argument can be parried a bit by replacing the strict knowledge constraint with the variable constraint. But still, it seems clear that if saying that p is enough of an exaggeration, then we have not asserted that p. It is as incorrect to say that Jane *asserted* that she knows she got everything wrong as it would be if she were just joking.

My account of the variation in usage in the bank and parking cases has a pragmatic component and a semantic component. The former is the thesis that terms can be used loosely or strictly, with loose use explained in terms of implicature. This component accounts for why the speaker changes from affirmation to denial, while both seem appropriate. The fact that someone uses a sentence loosely does not imply that it is false, or even that the speaker believes that it is false. Even hyperbole does not imply that what the speaker said is false, only that he believes it is false. Analogy: the hypothesis that the A and B contexts of the coffee case differ because “all gone” is used with variable strictness does not imply that the coffee is never strictly speaking all gone, nor even that it is not in fact all gone in this case. As these cases illustrate, loose use is not exaggeration. When “S knows p” is an exaggeration, the speaker thinks S is *far* from knowing, and does not imply that S is close enough to knowing for current purposes. In the bank and parking cases, the speaker does convey that the subject is close enough to knowing, and thus that the variably strict assertion constraint is satisfied.

The second component of my theory is the invariantist semantics on which knowledge requires complete justification. It is this component that accounts for our intuition that the affirmations and denials in the bank and parking cases are contradictory, and our intuition after careful reflection that what the speaker says in the A contexts is not strictly speaking true even though what he means is. Clearly, it is the invariantist semantics, not the variably strict pragmatics, that might lead to skepticism.

I have already observed that the variably strict pragmatics could be combined with a weaker invariantist semantics, or even a contextualist semantics that allows only limited variation, requiring less than complete justification. But even if we accept that knowledge requires complete justification, we need not affirm radical

⁴³ Mylan Engel raised this objection at the Mainz conference on contextualism (September 2003). Cf. DeRose (1999, p. 202), Cohen (1999, p. 83), Stanley (2004, p. 141ff; 2005, p. 84).

skepticism, the thesis that we do not or cannot know anything, nor even the less extreme view that “ordinary knowledge claims are generally false” (Stanley, 2005, p. 82). For these conclusions require the independent and in my view unwarranted thesis that subjects are never (or not ordinarily) completely justified in believing anything. When she arrives at the bank Saturday morning after nine o’clock and finds it bustling with activity, I think it is literally and strictly speaking true that Hannah now knows the bank is open. What she now sees and hears does suffice to establish that the bank is open, so she is completely justified in believing this. In both the A and the B contexts, moreover, “I know that I am speaking,” “I know you exist,” “I know we are driving,” “I know it is Friday,” and countless other knowledge claims, are all true and satisfy the complete justification requirement. My theory is thus fully compatible with what Hawthorne called the “Moorean Constraint.”

Very many ordinary positive knowledge ascriptions are true. Lewis thus writes: ‘It is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary’ (1996: 549). (Hawthorne, 2004a, p. 111)

Cohen believes that the invariantist semantics I have adopted would nonetheless lead to a “fairly radical” skepticism. He submits, for example, that I am forced to deny a truism about knowledge, namely: “much of what we know comes from the testimony of others.” An adequate reply is beyond the scope of this paper, and perhaps even my ability. We would first have to determine whether “know” is being used loosely or strictly here. It is surely a truism that testimony gives us enough justification for most practical purposes. It is less clear whether we can ever strictly speaking *know* that something is the case just because someone has said so. Newspapers, for example, routinely require independent confirmation before printing something based on a report. Scientists insist on replicating reported results. The additional premise that *testimony is generally reliable* does not enable us to know anything on the basis of testimony, just as the premise that nearly all entrants of a lottery lose does not enable us to know that a particular entrant lost. But suppose my wife screams in pain and calls urgently for my help. I say “Where are you?” and she replies “In the kitchen.” I would judge in this case that *my wife would not have said such a thing unless it were true*.⁴⁴ I would presume that she is saying what she believes, and that her belief is completely justified. Moreover, I have confirmation in the fact that her voice sounded like it came from the kitchen, and I recall hearing her walk into it a few minutes ago. I know she does not play practical jokes. Am I not completely justified in believing that my wife is in the kitchen? The answer is debatable, to be sure. Anyone who answers “No,” however, will have a hard time maintaining that I do know she is in the kitchen, and *a fortiori* that it is a truism that we can know things on the basis of testimony. Following Lewis (1996), contextualists maintain that when far-fetched skeptical possibilities become salient, such as that I am a brain in a vat, then “I know she is in the kitchen” is not true. I do not believe such possibilities are grounds for doubt, so I do not concede this point to the skeptic. Contextualists are not Moorean enough.⁴⁵

⁴⁴ Cf. Dretske (1971).

⁴⁵ I discuss this point briefly in Davis (2004), and develop it further in “Knowledge Claims and Context: Belief Differences.”

The quotation from Hawthorne introducing the Moorean Constraint contains two formulations. Lewis says *we know a lot*. Hawthorne says *many knowledge ascriptions are true*. These formulations are not equivalent. Hawthorne's formulation could be true while Lewis's is false if one person had all the knowledge. Even if we take "of many different people" to be implicit in Hawthorne's formulation, it does not entail Lewis's if contextualism is true. According to subject-contextualists, what we know depends on what context we are in. In some contexts we may know a lot. But in other contexts we know little or nothing. According to ascriber contextualism, whether "We know a lot" is true depends on the context of the ascriber. In some contexts it is true, in others it is false. The thesis that we know a lot, if contextualism is true, is like the thesis that many roads are flat. Relative to the standards of the Swiss or Nepalese, many roads are indeed flat. But relative to geometric standards, none are. Contextualists cannot even say without qualification that skepticism is false. For given what they say about epistemological contexts, "We know nothing" is true therein.⁴⁶ The fact that contextualists cannot maintain without qualification that we know a lot is a second respect in which the thesis is not Moorean enough.

DeRose adopts a metalinguistic and methodological version of the Moorean constraint. He says that the contextualist theory

is designed largely with the goal in mind of crediting most of our attributions of knowledge with truth (1995: 46).

The idea that "most" knowledge attributions—literally, over fifty percent of all the cases in which anyone anywhere has affirmed a statement of the form "S knows p"—is at best unverifiable. To justify this claim, DeRose cites a related constraint that is weaker in some respects and stronger in others.

We in general take it as a strike against a theory of a common term of a natural language that it involves the speakers of that language in systematic and widespread falsehood in their use of that term.⁴⁷ (DeRose, 1995, p. 46)

One respect in which DeRose's constraint is stronger than Hawthorne's is that it applies to denials as well as affirmations. This is important because contextualism is based on the assumption that the denials in the B-contexts are just as true as the affirmations in the A-contexts. Given the meaning of "not," this is what forces the conclusion that "S knows p" can be true in one context and false in another, and the further conclusion that our linguistic intuition that the denials contradict the affirmations is illusory. I am not sure what "systematic" means here. But by any reasonable standard, there would appear to be *widespread* falsehood in the use of "know." Think of all the instances in which "p" is about the uncertain future, the distant past, chance events, things misperceived, religious matters, outdated scientific results, precise quantities, stereotypes, urban myths, rumors, propaganda, and disinformation. Think too of all the instances in which S is simply an ignoramus.

⁴⁶ Contextualists could avoid this result by denying that there are any contexts in which standards of justification are so "high" that "I know I have a hand" and the like come out false. But then a large part of the motivation for the theory must be given up too. Cf. Lewis (1979, p. 355), Cohen (1999, pp. 64–67) and DeRose (1992, p. 917, 1999, Section 6; 2002, p. 168). Contrast Feldman (1999, pp. 104–107, 111) and Davis (2004).

⁴⁷ Cf. DeRose (1992, p. 924) and Cohen (1999, pp. 65, 80, 83). Compare and contrast Unger (1984, pp. 37–38, 70–77; 1986, p. 173ff).

It is not clear why making knowledge claims generally true should be cited as an advantage on *linguistic* grounds. “Witch” is still a common term of English, and used to be used regularly in daily life. Nevertheless, its meaning is such that all instances of “S is a witch” are and were false. “Alien” in its extraterrestrial sense is a perfectly good English noun; yet it is doubtful that any predication has been true. No such statement involving “Santa Claus” will ever be true. And consider the term “god,” one of the most common and widely used terms in English. Given that there are so many different religions making incompatible claims about gods, there is widespread and systematic falsehood in the use of “god.” Only those speakers of the true religion, if there is one, are speaking truly. Finally, nothing in linguistics forces us to hold that the uses of “all gone,” “two hours,” and “no chance” in our paradigm examples of loose use are strictly and literally true.

A range of knowledge attributions would have to be true if “knowledge” were an observation term like “water,” “cat,” or “thought,” expressing a concept acquired as a result of perceptual or introspective contact with its instances. But “knowledge” is not observational. Nothing like the causal theory of reference is plausible. Even if it were, the most we could infer is that *some* attributions must be true.

Finally, let us grant for the sake of argument that speakers generally use “S knows p” to communicate truths. A separate argument is needed to claim that the truths were communicated as part of what is said rather than as something implicated. When speakers use “know” to say something false while implicating something that is true, they are communicating truths. If such usage is commonplace, as the evidence I have presented indicates, it is hard to say whether speakers are engaged in “systematic and widespread falsehood in their use of the term.” It is also unclear how DeRose’s principle could be strengthened so as to favor truth in what is said over truth in what is implicated. No principle will be plausible, for example, if it implies that most things must have the exact measurements we ascribe to them, or that speakers cannot use figures of speech. This is especially true if it implies that speakers are generally ignorant of the indexicality of common terms, the contextualist’s error theory.

DeRose argues for his methodological principle as follows.

[Suppose] that a crazed philosopher claimed that there are no physicians, because, in addition to holding a medical degree, a necessary condition for being a physician is that one be able to cure any conceivable illness. On what grounds should we reject this bizarre conjecture in favor of a more traditional and less demanding account of what it is to be a physician?... it’s eminently reasonable to suppose that such facts as these, regarding our use, in thought and in speech, of the term ‘physician’ are involved: that we take to be physicians many licensed practitioners of medicine who don’t satisfy the demanding requirement alleged; that we seriously describe these people as being physicians;... etc. (DeRose, 1995, p. 47)

That competent speakers routinely apply “physician” to people they *take* to be far from satisfying that requirement is indeed evidence that it is not part of the meaning of the term. What “physician” means is determined by what *beliefs* or *thoughts* speakers use it to express; whether they are true or false is irrelevant.⁴⁸ If speakers

⁴⁸ For a much more complete answer to this sort of question, see Davis (2003).

did use “physician” to convey the belief that the subject can cure every possible illness, then the objective fact that they cannot would not undermine the conclusion that the requirement is part of the meaning of the term.

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 people who use “S knows p” commonly do not have such a belief, are not trying to indicate that they do, and may even agree that S cannot rule out all conceivable alternatives when the question is posed. If we found on the contrary 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 of “S knows p” even if that meant that its instances are generally false. Whether a particular belief is commonly expressed by users of “S knows p” is relevant to what “know” conventionally means; the truth of the belief is irrelevant.

The intuition constraint

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.

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. On subject contextualism, however, this change in our context is not relevant to the truth conditions of “Hannah knows.” So unless it is supplemented by an *ad hoc* error theory,⁴⁹ subject contextualism predicts that there should be no change in intuition on reflection, and is thus undermined by the additional intuitive data.

The pragmatic theory accommodates the intuitive data 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.⁵⁰ The raising of the stakes that occurs in the bank case makes even small differences between knowledge and something close to knowledge important, and thus makes us reflect especially carefully.

In support of my account of the intuitive data, I note that a similar process occurs with other expressions that can be used more or less strictly, with no affect on truth conditions. Thus at first we have the intuition that my claim “The coffee is all gone” is true. But when someone raises the question “Is there absolutely none left?” we reflect and judge that what I said was false. Similarly, when someone says “The yardstick is 3 feet long,” at first we would intuitively judge that the statement is true. But if someone raises the question, “Is the yardstick exactly 3 feet long?” we should reflect and judge that it might not be (unless the question is itself interpreted loosely).

Stanley (2005, p. 97) describes a similar case in which subject contextualists fail to account for our linguistic intuitions: the case in which the speaker is in a high-stakes context and the subject a low-stakes context. If it is very important to us whether the bank is open on Saturday, but not to Hannah, we will judge that Hannah does not know the bank will be open if her only evidence is that she has been there a few times on Saturdays. Subject contextualism also fails to account for our intuitions in Vogel’s parking case. For the change from A to B is a change in the ascriber Dick’s context that has no effect on the subject Alan’s context. Stanley’s response is that he is looking at intuitive judgements “to make vivid our commitment to the conceptual connection between knowledge and practical reasoning” (2005, pp. 98, 114–115). But the intuition casts doubt on whether we have that commitment. It is not the intuition

⁴⁹ Such auxiliary hypotheses are proposed by Hawthorne (2004a, pp. 162–166) and Stanley (2005, pp. 100, 102, 116). That they make faulty predictions in other cases is shown by Stanley (2005, p. 101) and Schaffer (2006, pp. 92–93). A further problem with Stanley’s auxiliary is that we who are evaluating S’s claim to know p may have no interest in whether “p” is true or not; so we ourselves are in a low-stakes context even when we judge after reflection that S does not know p.

⁵⁰ This is not to say that what the contextualist takes to be variation in truth conditions is really variation in propriety conditions. If that were true, then we should not find the knowledge denials in the B-contexts to be appropriate, as DeRose (2002, pp. 189–190) observed. I avoid DeRose’s (2002) talk of “warranted assertability” in this context. For the “know” in “S is warranted in asserting that A knows p” can be taken loosely or strictly.

we would expect if knowledge were connected with rational action in the way Stanley imagines. Further doubts were presented in the Section “The practical rationality connection.”

More importantly, our intuition about the high speaker-stakes, low subject-stakes case is evidence against subject contextualism because it predicts that we should not have that intuition. Stanley attempts to rebut this objection by claiming that “IRI is not a semantic thesis at all; it is rather a metaphysical thesis about the nature of the knowledge relation” (2005, p. 120). As Stanley defines it, “Bare Interest-Relative Invariantism (henceforth IRI) is simply the claim that *whether or not someone knows p may be determined in part by practical facts about the subject’s environment*” (2005, p. 85, my emphasis). This claim is about knowledge, not about the word “knowledge.” So Stanley is correct that linguistic intuitions about uses of “know” are not directly relevant to IRI. But for the same reason, IRI does not account for our intuitions about the bank and parking cases either—despite Stanley’s claim that it “gives a charitable explanation for all intuitions except for High Attributor-Low Subject” (2005, p. 98).

While IRI is not itself a semantic thesis, Stanley presupposed a semantic thesis when he stated IRI. In using the English word “know” to formulate his thesis, he presupposed that “*S knows p*” has a meaning on which its truth conditions are determined in part by practical facts about *S’s environment*. This is a semantic thesis, and it is a version of what I am calling “subject contextualism.” Call it *semantic-IRI*. Given that Stanley stated IRI using English, IRI is true if and only if semantic IRI is true. In virtue of this contingent equivalence, linguistic intuition is *indirect* evidence for IRI, which *implies* an explanation of that evidence. Our linguistic intuitions in the bank case support the semantic thesis Stanley presupposes. However, our intuitions in the high speaker-stakes, low subject-stakes case is evidence against semantic-IRI.

Ascriber contextualism also has a non-linguistic counterpart: *whether S knows p relative to c is determined by the standards of justification prevailing in c*, where *c* is a context of utterance. Since *S* does not necessarily exist in *c*, knowledge is a triadic relation on this view, rather than the dyadic relation assumed by subject contextualism. The triadic relational predicate I used here, however, is not a formula of standard English. One of the criticisms of ascriber contextualism as a semantic thesis is that if it were true, we would expect “know” to occur in such relational formulas in standard uses of natural language. But it does not (see the section “Relative knowledge”).

Stanley (2005, p. 24) argues that ascriber contextualism cannot account for our intuitions about cases in which the speaker is in a low-stakes context and the subject is in a high-stakes context. Consider:

The Copier Case

It makes no difference to Larry whether the drug store has a copier machine. Kelly, however, is late, and has to make copies for a presentation. If she stops at the drug store and it has a copier machine, everything will be fine. But if she stops in vain, she will have no time to stop anywhere else. Knowing that Kelly has often used the copier before, but not how important its having a copier today is, Larry says “Kelly knows the drug store has a copier.”

I believe that speaker contextualism does rule that Larry's statement is true, and that our intuitive judgement is that his statement is false. So I believe Stanley has identified a further problem for ascriber contextualism. Subject contextualism has a similar problem. It rules that "*Larry* knows the drug store has a copier" is true because Larry is in a low-standards context. But intuitively, we judge that it is false too. Since the fact that the drug store has often had a copier in the past does not establish that it has one today, we realize that "*Kelly* (or *Larry*) knows the drug store has a copier" is not strictly speaking true. Furthermore, since we know how important it is to Kelly that the drug store have a copier, we do not regard either statement close enough to being true for the contextually indicated purpose. We should also recognize, of course, that Larry had different purposes in mind, and that for his purposes, Kelly was close enough to knowing. So we should count Larry as having implicated something that is true. Our intuitive judgements, however, are about what Larry said and what that implicates to us. We regard's Larry's claim as misguided because he is not aware of Kelly's situation.⁵¹

Further difficulties for contextualism are illustrated by a variant of the time-measurement case.

Variant Time Measurement Case

A. Elaine asks Mike whether he knows how long he took to complete the final exam. Recalling that he looked at his watch when he handed the exam in, he answers "Yes, two hours." B. When Elaine learns that Nora said she took two hours and four minutes to complete the exam, she asks Mike whether he really knows that he took two hours and not perhaps a bit longer. He answers "No, my watch indicated two hours and two minutes."

First, Mike's answers in this variant seem intuitively as correct as the knowledge claims in the bank case. Yet if what Mike says in the B-context is true, and it actually took him two hours and two minutes to finish, then we have to conclude that Mike's knowledge claim in the A-context is false.⁵² For "*S* knows *p*" is true only if "*p*" is true. We cannot take the intuition that Mike's knowledge ascription is true at face value because that would violate the facticity of knowledge claims. This provides indirect evidence that we should not take our intuitions in the bank and parking cases at face value either. Second, the change in usage in this case cannot be explained in terms of a change in the contextually indicated standards of justification. For there was no change in those standards: consulting a watch is adequate evidence in both contexts. Nor was there any change in stakes or relevant alternatives.

My pragmatic theory accounts for the second time-measurement case just as well as the bank case, without rejecting facticity. Mike's claim in the A-context seems correct because he is close enough to knowing that he took two hours for that

⁵¹ The same explanation applies to our intuitions about Stanley's (2005, p. 5) "Ignorant High Stakes" case. Stanley's own low speaker-stakes, high subject-stakes example is more complex than my copier case, and as a result is less clear. To make a parallel change in my copier case, add to the description that Larry has decided to go to the drug store himself, just for fun, to see if he meets Kelly. This addition makes me focus on Larry's situation rather than Kelly's. Consequently I am inclined to judge that Larry's statement was close enough to being true.

⁵² As discussed in the Section "Semantic treatments of loose use," Cohen tries to avoid this conclusion with an indexical treatment of loose use.

context. Since context B demands greater precision, he is not close enough to knowing that he took two hours. The standard of precision relates to how loosely or strictly the knowledge claims are to be interpreted. In this case, the relevant respect of similarity to knowing that he took two hours is not the level of justification Mike possesses for what he believes, but the closeness of what he said he knows to both what he believes and the truth. The contextualist could allow that contexts vary in dimensions other than standards of justification. Yet the truth of what is known is not one of them.⁵³

A final problem for both ascriber and subject contextualist theories 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.⁵⁴ 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 isn’t “S knows that he will lose” true in low-standard contexts? Why wouldn’t the lottery case be treated just like the parking case?

The 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.

Cohen (1988, p. 107, 2000, p. 135) correctly 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

⁵³ Cf. DeRose (1992, p. 922, fn. 18), Partee (2004, p. 154) and Ludlow (2005, p. 25). Bach (2005, p. 62) notes that cognitive scientists and social scientists, among others, often use the term “knowledge” without consideration of whether the belief is true or justified. This can be viewed as a very loose usage appropriate in scientific contexts in which the only important component of knowledge is belief.

⁵⁴ Contrast Hawthorne (2004a, p. 161) and Cohen (2000, pp. 135–136, 2005, p. 200). Hawthorne (2004a, p. 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, p. 117ff) when criticizing Cohen.

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 hyperbole, but not Hannah or Dick. As far as I can see, however, Cohen's observation does not help the contextualist's cause. Cohen thinks that the salience of the chances of error will create a high-standard context. But the A context is a low-standard context. Since the fact that the chances of error are extremely remote is also salient, it should be evident that the standard of justification necessary to assert "Harry knows he will lose" in the A context is easily met. So if contextualism were true, it should seem evident that Harry's statement is true. But that is not our intuitive judgement.

Following Hawthorne (2000, p. 118), Cohen observes that in contexts like A, the knowledge claim is typically made with a special intonation. In addition to the inflection of "practical advice" they mention, it might be made with the intonation of hyperbole or scolding. Cohen suggests that the special inflection keeps the standards from rising, and thus makes the utterance seem true. But the standard remained low, while what Joe said seems false. I believe that the intonation serves as a signal of a non-literal utterance, and thus helps the hearer get beyond the obvious falsity of what was said and recognize what is meant.

Citing arguments of DeRose (1999), Stanley worries that "the tendency philosophers have to give pragmatic rather than semantic explanations of apparently semantic intuitions threatens to undermine the whole enterprise of giving semantic explanations" (2005: 14). Unlike those who invoke "Grice's Razor,"⁵⁵ my argument is not based on any general rule favoring pragmatic over semantic explanations of intuitions. Instead, I have argued that a theory combining an invariantist semantics with the pragmatic factor of variably strict usage accounts for a large range of intuitive data much better than contextualist theories. Both subject and speaker contextualism are supported by some semantic intuitions, to be sure. But they are undermined by many many others, as we have seen throughout this paper. A pragmatic theory avoids the faulty semantic predictions of contextualism, while providing a plausible account of the intuitions that support contextualist theories. It is supported by the strong analogy between the contextual variability of knowledge claims and that of claims of emptiness, measurement, impossibility, perfect flatness, and so on. So its hypothesis are not at all *ad hoc*, unlike the thesis that "know" has a *sui generis* kind of indexicality and the error theories posited by contextualists to explain away apparent discrepancies with the data.

Other pragmatic factors

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. Variable strictness does not explain the main stimulus for contextualist theory: the epistemology case.

Before students take epistemology, they typically affirm "I know that I have a hand" with no doubt or hesitation. After the professor explains the brain in the vat or evil demon hypothesis, they begin to wonder. Many come to deny that they really

⁵⁵ See fn. 22.

know they have a hand. But inevitably, students eventually return to affirming that they know. This case does not involve loose use or hyperbole. In everyday contexts, people cannot imagine a clearer case of knowledge. Careful reflection does not ordinarily prompt a retraction or correction of the claim to know one has a hand, they way it prompts Hannah to retract her claim to know the bank will be open on Saturday. Students in epistemology resist the suggestion that they do not really know, and experience mental turmoil, at least initially. In the bank and parking cases, there is no such conflict. It is natural to describe students as changing their minds as a result of studying epistemology. We would not describe Hannah as having changed her mind about whether she knows the bank will be open. The difference in the use of “know” in the epistemology case appears to be due to a *difference in belief* about what the subject knows (Davis, 2004, 2005). Courses in epistemology have the same effect on students’ beliefs about knowledge as courses in philosophy of religion may have on their beliefs about the existence of God.

To adequately explain *all* the variation in the use of “know,” then, we need a *multivariate* pragmatic theory. I will explore how belief differences account for variation in knowledge claims in greater depth in a sequel.⁵⁶

Acknowledgements I would like to thank the members of my Fall 2005 seminar, Emily Evans, Paul Naquin, Nate Olson, David Pierce, Mark Pitlyk, Diana Puglisi, and Dan Quattrone, for many hours of fruitful discussion in which many of the ideas of this paper were tried out or tested. Jason Stanley graciously providing the proofs of *Knowledge and Practical Interest* for us to work through. I am also grateful to Paul Portner and Elena Herburger for the opportunity to present part of this paper in their Spring 2006 seminar on implicature. Steven Gross also provided helpful comments. My greatest debt is to Stewart Cohen, who provided pages of extremely helpful and thought-provoking objections and replies. Attributions to Cohen without references were personal communications.

References

- Atlas, J. D. (1977). Negation, ambiguity, and presupposition. *Linguistics and Philosophy*, 1, 321–336.
- Atlas, J. D. (1979). How linguistics matters to philosophy: Presupposition, truth, and meaning. In C.-K. Oh, & D. A. Dinneen (Eds.), *Syntax and semantics, 11: Presupposition*. New York: Academic Press, pp. 265–281.
- Atlas, J. D. (1989). *Philosophy without ambiguity*. Oxford: Oxford University Press.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Bach, K. (2000). Quantification, qualification and context: A reply to Stanley and Szabó. *Mind and Language*, 15, 262–283.
- Bach, K. (2001). Speaking loosely: Sentence nonliterality. In P. French, & H. Wettstein (Eds.), *Midwest studies in philosophy*, 25, *figurative language*. Oxford: Blackwell, pp. 249–263.
- Bach, K. (2005). The emperor’s new ‘knows’. In G. Preyer, & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth*. Oxford: Clarendon Press, pp. 51–90.
- Cappelen, H., & Lepore, E. (2003). Context shifting arguments. *Philosophical Perspectives*, 17, 25–50.
- Carston, R. (1988). Implicature, explicature, and truth-theoretic semantics. In M. Kempson (Ed.), *Mental representations: The interface between language and reality* (pp. 155–81). Cambridge: Cambridge University Press; Reprinted in S. Davis (Ed.), *Pragmatics: A reader* (pp. 33–51). Oxford: Oxford University Press (1991).
- Cohen, S. (1986). Knowledge and context. *Journal of Philosophy*, 83, 574–583.
- 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.

⁵⁶ “Knowledge Claims and Context: Belief Differences.” See also Davis (2004, 2005).

- 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), 132–139.
- 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), 482–491.
- 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, 267–268.
- Davis, W. A. (2005). Contextualist theories of knowledge. *Acta Analytica*, 20, 29–42.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–929.
- 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, 193–197.
- DeRose, K. (1996b). Knowledge, assertion, and lotteries. *Australasian Journal of Philosophy*, 74, 568–580.
- DeRose, K. (1999). Contextualism: An explanation and defense. In J. Greco, & E. Sosa (Eds.), *The Blackwell guide to epistemology* (pp. 187–205). Oxford: Blackwell.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111, 126–203.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
- Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40, 363–378.
- Feldman, R. (1999). Contextualism and skepticism. *Philosophical Perspectives*, 13(Epistemology), 91–114.
- Fogelin, R. J. (2000). Contextualism and externalism: Trading in one form of skepticism for another. *Philosophical Issues*, 10(Skepticism), 43–57.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Grice, H. P. (1975). Logic and conversation. In P. Cole, & J. Morgan (Ed.), *Syntax and semantics*, 3: *Speech acts*. New York: Academic Press; Reprinted In H. P. Grice (Ed.), *Studies in the Way of Words* (pp. 22–40). Cambridge, MA: Harvard University Press (1989).
- Grice, H. P. (1978). Further notes on logic and conversation. In P. Cole (Ed.), *Syntax and semantics*, 9: *Pragmatics* (pp. 113–128). 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), 117–120.
- Hawthorne, J. (2004a). *Knowledge and lotteries*. Oxford: Clarendon Press.
- Hawthorne, J. (2004b). Replies. *Philosophical Issues*, 14(Epistemology), 510–523.
- Heller, M. (1999). Contextualism and anti-luck epistemology. *Philosophical Perspectives*, 13(Epistemology), 115–129.
- Hofweber, T. (1999). Contextualism and the meaning intention problem. In K. Korta, E. Sosa, & J. Arrazola (Eds.), *Cognition, agency, and rationality* (pp. 93–104). Dordrecht: Kluwer Academic Publishers.
- Horn, L. R. (1989). *A natural history of negation*. Chicago: University of Chicago Press.
- Jackson, F. (1979). On assertion and indicative conditionals. *Philosophical Review*, 88, 565–589.
- Klein, P. (2000). Contextualism and academic skepticism. *Philosophical Issues*, 10(Skepticism), 108–116.
- Kompa, N. (2002). The context sensitivity of knowledge ascriptions. *Grazer-Philosophische Studien*, 64, 1–18.
- Kompa, N. (2005). The semantics of knowledge ascriptions. *Acta Analytica*, 20, 16–28.
- Kripke, S. (1977). Speaker reference and semantic reference. *Midwest Studies in Philosophy*, 2, 255–278.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Lewis, D. (1979). Scorekeeping in a language game. In R. Bäuerle et al. (Eds.), *Semantics from different points of view* (pp. 172–187). Berlin: Springer.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.

- Ludlow, P. (2005). Contextualism and the new linguistic turn in epistemology. In G. Preyer, & G. Peters (Eds.), *Contextualism in philosophy* (pp. 11–50). Oxford: Oxford University Press.
- Malcolm, N. (1952). Knowledge and belief. *Mind*, 51, 178–189.
- Moore, G. E. (1962). *Commonplace book: 1919–1953*. London: Allen & Unwin.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15, 509–559.
- Partee, B. (2004). Comments on Jason Stanley's "On the linguistic basis for contextualism". *Philosophical Studies*, 119, 147–159.
- Prades, J. L. (2000). Scepticism, contextualism, and closure. *Philosophical Issues*, 10(Skepticism), 121–131.
- Pritchard, D. (2005). Contextualism, skepticism, and warranted assertibility manoeuvres. In J. Keim-Campbell, M. O'Rourke, & H. Silverstein (Ed.), *Knowledge and skepticism*. Cambridge, MA: MIT Press.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Nous*, 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, 317–333.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Stanley, J. (2004). On the linguistic basis of contextualism. *Philosophical Studies*, 119, 119–146.
- Stanley, J. (2005). *Knowledge and practical inference*. Oxford: Oxford University Press.
- Stanley, J., & Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15, 219–261.
- Stine, G. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–261.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Clarendon Press.
- Unger, P. (1984). *Philosophical relativity*. Minneapolis: University of Minnesota Press.
- Unger, P. (1986). The cone model of knowledge. *Philosophical Topics*, 14, 125–178.
- Vogel, J. (1999). The new relevant alternatives theory. *Philosophical Perspectives*, 13(Epistemology), 155–180.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.