

The Assessment Sensitivity of Knowledge Attributions*

John MacFarlane[†]

June 28, 2004

Recent years have seen an explosion of interest in the semantics of knowledge-attributing sentences, not just among epistemologists but among philosophers of language seeking a general understanding of linguistic context sensitivity. Despite all this critical attention, however, we are as far from consensus as ever. If we have learned anything, it is that each of the standard views—invariantism, contextualism, and sensitive invariantism—has its Achilles’ heel: a residuum of facts about our use of knowledge attributions that it can explain only with special pleading. This is not surprising if, as I will argue, there is a grain of truth in each of these views.

In this paper, I propose a semantics for “know” that combines the explanatory virtues of contextualism and invariantism. Like the contextualist, I take the extension of “know” to be sensitive to contextually determined epistemic standards. But where the contextualist takes the relevant standards to be those in play at the context of *use*, I take them to be those in play at the context of *assessment*: the context in which one is assessing a particular use

*I presented versions of this paper to the Stanford Philosophy Department on October 17, 2003, to the Themes in Philosophy of Language conference at Yale on November 8, 2003, and to the Department of Logic and Philosophy of Science at UC Irvine on December 5, 2003. I am grateful to audiences at all three talks for stimulating discussions, and especially to Keith DeRose, who commented on my paper at Yale. I would also like to thank Kent Bach, Gilbert Harman, Ram Neta, Jonathan Schaffer, Lionel Shapiro, and Matt Weiner for useful correspondence and discussion. This work was made possible in part by an ACLS/Andrew W. Mellon Fellowship for Junior Faculty and a UC Berkeley Humanities Research Fellowship.

[†]Department of Philosophy, University of California, Berkeley. 314 Moses Hall, Berkeley, CA 94720–2390. E-mail: jgm@berkeley.edu.

of a sentence for truth or falsity. Thus, I can agree with the invariantist that “know” is not sensitive to the epistemic standards in play at the context of use, while still acknowledging a kind of contextual sensitivity to epistemic standards. The proposed semantics for “know” is contextualist along one dimension (contexts of assessment) and invariantist along another (contexts of use).¹

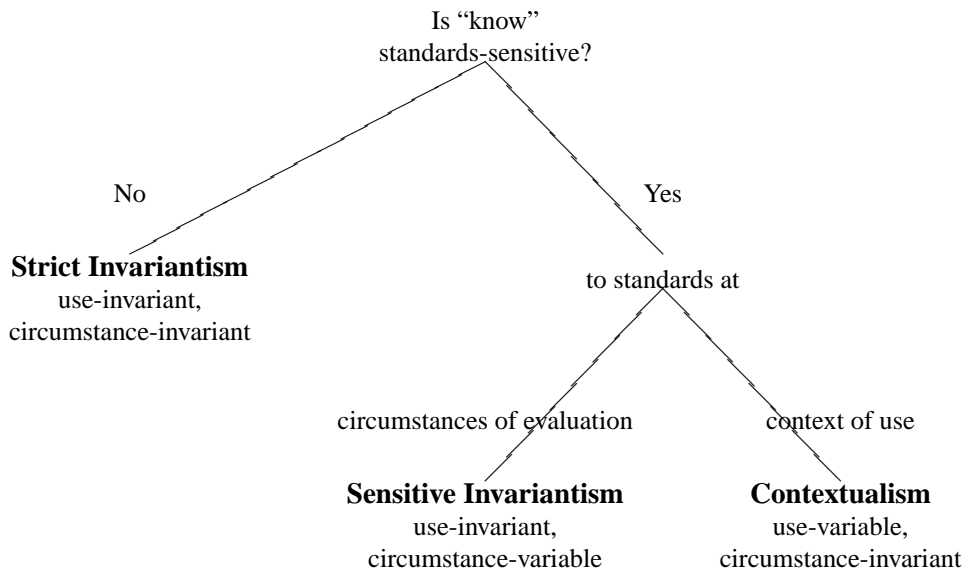
In the first part of the paper, I motivate my proposal by considering three facts about our use of “know” (§2) that collectively cause trouble for *all* of the standard views about the semantics of “know” (taxonomized in §1). I argue that the usual attempts to explain away the anomalies by appeal to pragmatics or to speaker error are unpersuasive (§3). In §4, I show how standard semantic frameworks must be modified to make room for my “relativist” semantics, and I show how the proposed semantics makes sense of the features of our use of “know” that proved puzzling on the standard views. Finally, in §5, I respond to worries about the coherence of relativist semantics by describing the role assessment-relative truth plays in a normative account of assertion.

1 A taxonomy

For our purposes, the standard views about the semantics of “know” can be divided into three main classes. *Strict invariantists* hold that “know” is associated with a fixed epistemic standard, in much the same way as “six feet apart” is associated with a fixed standard of distance. A person and a fact satisfy “ x knows y ” just in case the person’s epistemic position with respect to the fact is strong enough to meet this fixed epistemic standard. *Sensitive invariantists* allow the epistemic standard to vary with the subject and the *circumstances of*

¹For a kindred view, developed rather differently, see Richard 2004. I learned of Richard’s work too late to take account of it in this paper. There are also some affinities between the present proposal and the “perspectival” view of knowledge attributions defended in Rosenberg 2002: ch. 5 (see esp. 148–9, 163–4), though Rosenberg does not develop his proposal in a truth-conditional framework.

Figure 1: Standard taxonomy of positions on the semantics of “know”



evaluation (in the sense of Kaplan 1989), in much the same way as the standard of distance expressed by “as far apart as Mars and Jupiter” varies with the circumstances (for instance, the time) of evaluation. And *contextualists* allow the epistemic standard to vary with the *context of use*, like the standard of distance expressed by “as far apart as my two hands are right now.” The differences are summed up in Figure 1.

This is of course only one way of carving up the range of positions that have been taken, and it lumps together positions that may seem very different, even from a semantic point of view.² The advantage of this taxonomy is that it will allow us to see in a perspicuous way what is wrong with *all* of the views it encompasses. Because a “formal” taxonomy will be enough for our purposes, I leave it completely open here what an epistemic position is, how an epistemic standard might be specified, and what features determine which epistemic standard is relevant in a given context or circumstance. In particular, although I will sometimes talk of “high” and “low” standards, I wish to leave it open whether standards

²For example, it classes the view advocated in Kompa 2002 as a form of contextualism, even though on Kompa’s view “know” always expresses the same, “unspecific” relation, and so would be counted a form of invariantism on some criteria. Kompa’s view will be discussed later, in §4.2.

vary on a linear scale—from “low” to “high”—or in a more complex and qualitative way, as on “relevant alternatives” theories. Different views in each of our formal categories will cash out these notions in different ways. The arguments that follow abstract from these details.

The differences between contextualism and sensitive invariantism tend to be obscured when we consider first-person, present-tense knowledge attributions. For in these cases the epistemic standards in play at the context of use will coincide with those in play for the subject at the circumstances of evaluation. To see the differences, we need to vary the context of use while keeping the circumstances of evaluation constant—say, by considering “On Tuesday, Joe knew that whales are mammals” as uttered by Sally on Wednesday and by Fred on Thursday—and vary the circumstances of evaluation while keeping the context of use constant—say, by considering both “On Tuesday, Joe knew that whales are mammals” and “On Wednesday, Joe knew that whales are mammals” as uttered by Fred on Thursday. Contextualism predicts that the epistemic standard one must meet in order to count as “knowing” should shift as we shift the context of use (even if the circumstances of evaluation are kept fixed), while sensitive invariantism predicts that it should shift as we shift the circumstances of evaluation (even if the context of use is kept fixed). Thus we may aptly describe a contextualist semantics for “know” as *use-variable* and a sensitive invariantist semantics as *circumstance-variable*. Standard versions of contextualism are *circumstance-invariant*, and standard versions of sensitive invariantism are *use-invariant*—though of course one might also have a hybrid view that was both use-variable and circumstance-variable. Strict invariantism is both use-invariant and circumstance-invariant.

2 Some facts about our use of “know”

I now want to look at three facts about our use of knowledge attributions that have figured prominently in discussions of the semantics of “know.”

2.1 Variability of standards

Normally, I am perfectly happy to say that I know that my car is parked in my driveway. I will say this even when I’m at work, several miles away. But if someone asks me how I know that my car has not been stolen (and driven away), I will admit that I do not know this. And then I will have to concede that I do not know that my car is in my driveway: after all, if I knew this, then I would be able to deduce, and so come to know, that it has not been stolen.

How should we think of my shift from claiming to know to denying that I know? It doesn’t seem right to describe me as having learned something, or as correcting a mistake. If I have learned something, what exactly have I learned? It’s not as if I was unaware of the existence of car thieves when I made my original knowledge claim. Besides, the next day I will go right back to saying that I know that my car is in my driveway. Am I so dense as never to learn from my mistakes?

Nor does it seem right to say that when I claimed to know, I didn’t mean it literally. I would have said the same thing in a forum where non-literal speech is discouraged, like a courtroom. And I would have said the same thing if I had been instructed to say just what I meant, without exaggeration, artifice, or innuendo. Indeed, I would have said the same thing in a crowd of epistemologists, so it was not just a matter of “speaking with the vulgar.”

Perhaps my mistake lies in conceding that I don’t know that the car is in the driveway. Perhaps the fact that I cannot rule out auto theft is actually irrelevant to whether I know.

But what, then, *should* I say when the possibility is floated? Should I ignore it or dismiss it as irrelevant? That might be the right response to certain far-fetched sceptical worries—say, “how do you know that the matter in your car has not spontaneously reorganized to form a giant lizard?”—but it hardly seems appropriate in response to a perfectly mundane worry about thieves. Should I say “Although I know that the car is in my driveway, there’s always a chance that it has been stolen and is not in my driveway”? This sounds close to contradictory. Or should I say “Since I know that the car is in my driveway, I know that it hasn’t been stolen”? That too seems wrong. I am not in a position to know that the car has not been stolen. If I am making a mistake, it is not one that ordinary speakers recognize as a mistake.

If I was speaking literally both times and didn’t make a mistake, then presumably the standards I must meet in order to count as “knowing” must have changed. I met the laxer standards that were in play at the time of my first knowledge claim, but not the stricter ones that came into play after the mention of car thieves.

Examples like this can easily be multiplied. They form the basis of standard arguments for contextualism and sensitive invariantism.

2.2 Embedded occurrences of “know”

Temporal and modal operators shift the circumstances of evaluation. But we seem to use the same epistemic standard in evaluating “know” when it is embedded in the scope of temporal or modal operators as we do when it occurs unembedded. We don’t seem to *mix* different standards at a single context of use, even when we’re considering putative knowers in very different circumstances.

To take up the example from the last section, when I concede that I don’t know that my car is parked in the driveway, I won’t insist that I *did* know this two minutes ago, before

the bothersome question raised the standards. I will say that I did not know it then either. In deciding whether I knew it then, I use the standards in play now, not the standards that were in play then.

Relatedly, we do not say things like “Before the possibility that he might win the lottery became relevant, John knew that he would not be able to afford health insurance, but now he does not know this (though he still believes it),” or “John knows that he won’t be able to afford health insurance, but if he were discussing the possibility that he might win the lottery, he would not know this.” If the judge asks Doris whether she knew on January 13 that her car was in the driveway, it would be positively bizarre for her to answer “I don’t know: I can’t remember whether I was worried about car thieves that day” or “Remind me: what epistemic standards were in play at that time?” All this suggests that at any given context of use, we hold the standards that one must meet in order to count as “knowing” constant over all circumstances of evaluation. Observations such as these form the basis of standard arguments against sensitive invariantism.³

2.3 Truth ascriptions and retraction

When standards have been raised, I will say not only that I don’t know that my car is in my driveway, and that I *didn’t* know this earlier, but that my earlier assertion of “I know that my car is in the driveway” was *false*.⁴ In part, this is because we tend to report knowledge claims homophonically, even when they were made in very different epistemic contexts.⁵ Thus, I will report myself as having asserted that I knew that my car was in the driveway. Since I now take myself not to have known this, I must reckon my earlier assertion false.

I won’t just *say* that it was false; I will *treat* it as false. If challenged, I will *retract* my

³See DeRose 2000 and DeRose, forthcoming.

⁴Noticed by Feldman 2001: 77, Rosenberg 2002: 164, and Hawthorne 2004: 163, among others.

⁵See Hawthorne 2004: §2.7.

earlier claim, rather than reformulating it in a way that shows it to be consistent with my current claim—for example, by saying, “What I asserted was merely that I met the standard for ‘knowing’ that was in place *when I was making the claim.*”⁶ I will have correlative expectations about when others ought to retract their knowledge claims. If yesterday Sally asserted “I know that the bus will be on time,” and today she admits that she didn’t know yesterday that the bus would be on time, I will expect her to retract her earlier assertion. I will find it exceedingly bizarre if she replies by saying that her assertion was true, even if she adds “by the standards that were in place yesterday.”

In these respects “know” functions very differently from ordinary indexicals like “here” and from other expressions generally regarded as context-sensitive, like “flat” and “tall.”⁷ Suppose I’m on a moving train. At 3:30 we pass some big factories and tenement houses, and I say “It’s very urban here.” By 3:31 we have passed into suburbs, and I say “It’s not very urban here.” I *won’t* retract my earlier claim. If it is challenged, I’ll say: “When I said a minute ago ‘It’s very urban here,’ what I said was true, and I stand by that, even though it’s not very urban *here.*” To avoid confusion, I may reformulate my earlier claim: “What I asserted was that it was very urban where we were a minute ago.” Similarly, if I find myself in a scientific context where tiny bumps and ridges are important, I might assert “The table is not flat,” but I would not regard this as any reason to withdraw my assertion, made earlier in an everyday context, of “The table is flat.” If pressed, I would say: “I only committed myself to the table’s being flat by everyday standards.”

If we are correct in ascribing truth and falsity to our earlier knowledge claims in light of present standards, and retracting or standing by them accordingly, then it seems that we do not take the epistemic standards one must meet in order to count as “knowing”

⁶As Stephen Schiffer notes, “no ordinary person who utters ‘I know that *p*,’ however articulate, would dream of telling you that what he meant and was implicitly stating was that he knew that *p* relative to such-and-such standard.” (Schiffer 1996: 326–7). See also Feldman 2001: 74, 78–9, Hawthorne 2004: §2.7.

⁷See Stanley 2004 for a detailed discussion of differences between “know” and various kinds of context-sensitive expressions.

to vary across contexts of use. This fact forms the basis of standard arguments against contextualism.

3 Assessing the standard views

Let's assemble the upshots of these observations. The apparent *variability of standards* suggests that the truth of sentences containing "know" depends somehow on varying epistemic standards. That would rule out strict invariantism. The facts about *embedded occurrences* suggest that the semantics of "know" is circumstance-invariant. That would rule out sensitive invariantism. And the facts about *truth ascriptions and retraction* suggest that the semantics of "know" is use-invariant. That would rule out contextualism. Taken at face value, then, our three facts about use seem to rule out all three standard views about the semantics of "know."

What should we conclude? I think we have three basic options:

1. We can argue that one of our three facts about use is a misleading guide to the semantics of "know," either
 - (a) because it can be explained pragmatically, in terms of our broader communicative purposes, or
 - (b) because it can be attributed to systematic and widespread error on the part of ordinary speakers.
2. We can argue that our practice in using "know" is so confused and incoherent that knowledge-attributing sentences cannot be assigned definite truth conditions. Instead of doing semantics, we can advocate reform, perhaps through the introduction of new, unconfused terms of epistemic assessment.

3. We can try to make conceptual space for a semantics for “know” that is use-invariant and circumstance-invariant, but still somehow sensitive to changing epistemic standards.

My aim in this paper is to explore the last of these options, which I will take up in §4, below. But first I want to say a bit about why I find the other options unpromising.

3.1 Pragmatic explanations of the data

One of the most important lessons of philosophy of language in the 1960s was that the connection between meaning and use is indirect.⁸ Even if we restrict ourselves to sincere, knowledgeable informants, the most we can discern directly from their use of sentences are the conditions in which they find it reasonable to use these sentences to make assertions. And these are not the same as the truth conditions. It is often reasonable to make assertions using sentences one knows to be literally false—not just because it is sometimes reasonable to lie, but because it is often reasonable to engage in hyperbole, harmless simplification, irony, and metaphor. Conversely, it is often reasonable to *refrain* from asserting something that is true, germane to the topic, and potentially informative. For example, one might refrain from asserting that Harvard has one of the fifty largest university libraries in the world—though this is true—because doing so would encourage certain audiences to infer that Harvard is closer to number fifty than to number one.

Thus the facts about use catalogued in the previous section do not *by themselves* rule out any proposal about the semantics of “know.” These facts may tell us something about when people find it reasonable to use certain sentences containing “know” to make assertions, but they do not directly tell us anything about the truth-conditions of these sentences. To get from use to truth-conditions, we must rule out the possibility that it is reasonable to use

⁸See Grice 1989, Searle 1969: ch. 6.

these sentences *despite* their falsity, or to refrain from using them *despite* their truth. I know of no fully general way of doing this: all we can do is examine putative explanations one by one and show how they fail. Because we will consider the possibility of speaker error in §3.2, we will assume in this section that speakers are under no relevant substantive or semantic misapprehensions: when they utter false sentences, they know that they are false, and when they refrain from uttering true sentences, they know that they are true.

3.1.1 Variability of standards

The variability data is primarily a problem for strict invariantists. Strict invariantists come in two varieties. *Sceptical* invariantists hold that the fixed epistemic standards are very stringent, perhaps so stringent that human beings never meet them (at least with respect to empirical facts). *Moderate* invariantists hold that the standards are meetably lax. The two kinds of invariantists face different challenges in giving a pragmatic explanation of the variability data, so I will consider them separately.

(a) Fixed high standards

If standards are fixed and high, we need to explain why speakers should so frequently find it reasonable to claim to know things they are fully aware they don't know. (Remember, we are saving the possibility that speakers are unaware of their own ignorance for later.) One possible explanation is that they are trying not to mislead others who do not realize that the standards for knowledge are very high, and who would conclude from a denial of knowledge that the speaker was in a much poorer epistemic position than is actually the case. But this explanation applies only to the discourse of an enlightened sceptic talking to the unenlightened masses—surely a very special case. To explain the masses' own low-standards attributions of knowledge, an error theory would be needed.

Another possibility is that speakers are prone to hyperbole. Just as I might say “I could

eat a horse!” instead of saying, more accurately, “I could eat ten pancakes and a four-egg omelette,” so I might say that I *know* my car is in my driveway instead of saying merely that I have pretty good reason to believe this. If this kind of hyperbole were systematic and widespread, it might explain why we often claim to know things even when our grounds fall short of being conclusive.⁹ But I find the prospects of such an explanation dim. Hyperbole must be deliberate: if I really believed that I could eat a horse, I would not be exaggerating in saying that I could. However, ordinary speakers don’t seem to regard their ordinary knowledge claims as exaggerations. Nor do they mark any distinction between what they literally know and what they only hyperbolically “know.” When their knowledge claims are challenged, they don’t say “I was speaking hyperbolically,” the way I would if you replied to my horse-eating boast by saying, “Not even a grizzly bear can consume an entire horse in one sitting.”

In defense of the hyperbole view, Jonathan Schaffer notes that hyperbole can be “non-obvious,” particularly when it is highly formulaic (forthcoming, n. 3). We are so accustomed to the trope “I’m dying of thirst” that we no longer pause to consider its literal significance; instead, we jump directly to the intended meaning. Schaffer concludes that “the fact that ‘I know that I have hands’ is not obviously hyperbolic is no objection.” But my point is not about obviousness. Even if speakers do not realize at first that in saying “I’m dying of thirst” they are speaking hyperbolically, they will immediately concede this when it is pointed out to them. “Of course I’m not literally *dying*,” they will say, “and I never meant to suggest that I was.” In contrast, those who say “I know that I have hands” will not, in general, concede that they were speaking hyperbolically, even when confronted with sceptical counterpossibilities. No one reacts to the sceptic by saying, “I never meant to suggest that I literally *knew* that I had hands!”

A third approach would appeal to the *inconvenience* of adding all the pedantic hedges

⁹See Schaffer, forthcoming.

and qualifications that would be needed to make our ordinary knowledge claims strictly true. As long as no one is likely to be misled, it may be more efficient to assert (falsely) that one knows that p than to assert (truly, but clumsily) that one knows that probably p , unless of course q ; or that one has ruled out possibilities X , Y , and Z , but not W . For the same reason, one might say “My tank holds 15 gallons” when it really holds 14.5. As the potential misleadingness of unqualified and strictly false knowledge claims varies with the conversational context, so does our willingness to make them.

Like the hyperbole view, however, this approach fails to explain how we actually react when our ordinary knowledge claims are challenged. If I say “My tank holds 15 gallons” and someone calls me on it—“But the manual says it holds 14.5!”—I will say, “I was speaking loosely: what I meant was that it holds *about* 15 gallons.” But if I say “I know that my car is in my driveway” and someone calls me on it—“How can you rule out the possibility that it has been stolen?”—I will *not* say, “I was speaking loosely: what I meant was that I know that my car is *most likely* in my driveway,” or “What I meant was that I know that my car is in my driveway, provided it has not been stolen or moved in some other abnormal way.” In this respect I believe I am representative of ordinary speakers: otherwise, sceptical arguments would be greeted with shrugs, not surprise.

(b) *Fixed low standards*

If standards are fixed and low, then what needs explaining is why we sometimes *deny* that people know, even when they clearly meet these standards. Patrick Rysiew has suggested that we sometimes deny that we know because we do not want to implicate that we can rule out certain salient but irrelevant counterpossibilities.¹⁰ In asserting that p , one ordinarily represents oneself as knowing that p . If I make this implicit knowledge claim explicit by saying “I *know* that my car is parked in my driveway,” my choice of words will

¹⁰Rysiew 2001: 492, 499.

be noticed. My hearers may well wonder why I did not simply say “My car is parked in my driveway,” and they may assume I meant to imply that I could rule out the conversationally salient possibility that my car had been stolen. Even if I do not need to rule out this possibility in order to count as knowing, I do not want to be taken to be implying that I can rule it out. So, Rysiew argues, I have reason to disavow knowledge.

This is an ingenious explanation, but it fails on two counts. First, although worries about misleading implicatures may be good reasons to refrain from asserting something, they aren’t good reasons to assert its *negation*. Before Cal has played any games, I will refrain from asserting (truly) that Cal has won all of its games so far this season, because my doing so would misleadingly imply that Cal has played at least one game already. But these considerations do not give me any reason to assert that Cal has *not* won all of its games so far this season. Similarly, even if Rysiew’s story can explain why it would be rational for me to refrain from saying that I know, it cannot explain why I should say that I *don’t* know.

Second, even if Rysiew’s explanation worked in the first-person case, it could not be extended to third-person knowledge attributions. It is essential to Rysiew’s explanation that the question arises, “Why did the speaker say that he *knows* that *p* rather than just that *p*?” The question does not arise in the same way in third-person cases. In saying that *p*, one does not ordinarily implicate that someone else, *X*, knows that *p*. So an assertion that *X* knows that *p* does not call attention to itself in the same way as a first-person knowledge ascription. Thus, Rysiew’s explanation does not generalize to third-person knowledge attributions. But the phenomenon it seeks to explain does extend to third-person attributions. So the explanation fails.

3.1.2 Embedded “know”

There is an easy pragmatic explanation for the infelicity of asserting “I knew that p earlier, but now that standards have gone up, I don’t know that p .”¹¹ In asserting that I knew that p earlier, I represent myself as knowing that I knew that p . But in representing myself as knowing that I knew that p , I also represent myself as knowing that p , since it is common knowledge that knowledge is factive. Thus there is a clash between what I commit myself to in asserting “I don’t know that p now” and what I represent myself as knowing in asserting “I knew that p earlier.”

But this explanation only takes us so far. It explains why we do not assert “I don’t know now that p , but I knew then that p .” But it does not explain our tendency to *deny* that we knew then that p .¹² Nor does it explain why it is infelicitous to assert “If p is true, then I knew that p before standards went up, though I don’t know that p now,”¹³ or “Joe doesn’t know now that p , but he knew then that p ,” or “I know now that p , but I didn’t know then that p ,” when all that has changed are the standards. Here, it seems, a defender of circumstance-variable semantics must resort to an error theory.

3.1.3 Truth ascriptions and retraction

It might be suggested that the *inconvenience* of reformulating knowledge claims in a way that reflects their dependence on past standards sometimes makes it reasonable to treat them as if they had been made in light of current standards—even if this means saying that they were false when we know that they were true. The differences in usage between “know” and ordinary indexicals might then be attributed to the comparative ease of reformulating claims made using ordinary indexicals when the relevant contextual factors have changed.

¹¹For a slightly different version of this explanation, directed at third-person knowledge ascriptions rather than past-tensed ones, see Hawthorne 2004: 160.

¹²See DeRose 2002: §3.

¹³Hawthorne 2004: 166.

If I say “I am tired now” at 3:30 p.m. today, others can easily re-express the content of my claim tomorrow by using the sentence “he was tired at 3:30 p.m. yesterday.” But when it comes to “know”—supposing that “know” is context-sensitive—things are messier. How can we re-express a knowledge claim made in one context in another, where standards are different? I might say something like this: “I asserted that I knew, by the relatively low standards for knowing in place at the time, that my car was in my driveway.” Or perhaps: “I said something that is true just in case I met the standards in place at the time for knowing that my car was in my driveway.” But these reformulations are cumbersome and not very informative.¹⁴ Even if they are correct, it may seldom be worth the trouble to use them; in many cases, it may be more efficient simply to withdraw the earlier knowledge claim. In this way, a contextualist might attempt to explain away the data about truth ascriptions and retraction that suggest a use-invariant semantics for “know.”

But if this is the explanation of our retraction behavior, there ought to be *some* cases in which the disadvantages of retracting outweigh the inconvenience of reformulating. Suppose Sam is in the courtroom:

Judge: Did you know on December 10 that your car was in your driveway?

Sam: Yes, your honor. I knew this.

Judge: Were you in a position to rule out the possibility that your car had been stolen?

Sam: No, I wasn't.

Judge: So you didn't know that your car was in the driveway, did you?

¹⁴More informative reformulations would require a way of specifying epistemic standards directly, rather than as the standards in play at such-and-such a context. We do not consider speakers masters of the indexicals “here” and “now” unless they are in command of coordinate systems for specifying places and times independently of utterance events (“in Berkeley, California,” “at 3:30 p.m. GMT on October 14, 2003”), which they can use to reiterate claims made using these indexicals in other contexts. Ordinary speakers possess no comparable coordinate system for specifying epistemic standards.

Sam: No, I suppose I didn't, your honor.

Judge: But you just said you did. Didn't you swear an oath to tell the whole truth, and nothing but the truth?

However inconvenient it would be for Sam to reply,

My claim was that on December 10 I knew, by the standards for knowledge that were in play before you mentioned car thieves, that my car was in my driveway. That was true, your honor, so I did not speak falsely,

it would surely be more inconvenient for him to be charged with perjury. Nonetheless, I think that Sam, if he is like most ordinary speakers, will concede that his previous assertion was false and promise to be more careful in his future answers. This suggests that the calculus of inconvenience alone cannot explain why speakers tend to abandon their earlier knowledge claims when they are shown to be false in light of present standards.

3.2 Error theories

A sincere speaker who wants to speak the literal truth and avoid literal falsity may fail to do so if she has false beliefs, either about the facts or about the literal meanings of the words she uses.¹⁵ If I believe (as I once did) that "gravy" is the name of a vitamin-deficiency disease, I will refrain from asserting "I like gravy," even if I do like meaty sauce. And if I believe that whales are fish, I may assert "Whales are fish," even though this is false. Before we make any inferences from facts about ordinary use to truth conditions, then, we

¹⁵Although I doubt that a clean distinction between semantic and substantive error can be made, a rough and ready distinction will suffice for our purposes here. Note that it is ignorance about literal meaning that is at stake here, not ignorance about speaker's meaning. On many accounts of speaker's meaning, it is implausible to suppose that a speaker could be ignorant of what *she* means. Nonetheless, she can very well be ignorant of what *her words* mean, or of what she has literally *said*. See Rysiew 2001: 483, commenting on §IV of Schiffer 1996.

must rule out the possibility that ordinary speakers are systematically mistaken in certain ways. As before, we'll consider our three facts about use in turn.

3.2.1 Variability of standards

To explain the variability data, moderate strict invariantists must argue that speakers often underestimate their success in meeting the standards for knowledge and as a result disavow knowledge that they actually possess. Sceptical invariantists, by contrast, must argue that speakers systematically *overestimate* their success in meeting the standards for knowledge and as a result claim to know when in fact they do not.

The sceptical version of the error theory is sometimes rejected on the grounds that it rules expressions of paradigm cases of knowledge, like “I know that I have hands,” false. But the paradigm case argument is not a good argument. A supposed paradigm case of *F*-ness can turn out not to be an *F* at all. Whales turned out not to be fish; glass turned out not to be a solid. This might even happen on a large scale. Suppose that in 1750, all the emeralds on earth had been replaced by synthetic duplicates indistinguishable by the technology of the time. Then *none* of the extant “paradigm cases” of emeralds would have been emeralds. The sceptic’s claim that ordinary speakers are mistaken in nearly all of their knowledge claims cannot be rejected out of hand.

Nonetheless, it is fair to ask the sceptical invariantist for an *explanation* of the widespread and uniform error she attributes to speakers. Why do speakers so quickly revert to making everyday knowledge claims even after they have been led through sceptical arguments?¹⁶ Human beings are educable; the fact that the lesson does not stick deserves special explanation. Moreover, the sceptic must explain how “know” comes to have the exacting meaning it has, despite the fact that looser use is the norm. (It would be difficult to argue that “decimate” still means just “to kill one in every ten of,” when it is now routinely used for cases

¹⁶Cf. Hawthorne 2004: 131.

of larger-scale destruction.) Here the sceptic will have to put great weight on certain widely accepted generalizations about knowledge (such as closure principles) that can be exploited in sceptical arguments. But it is not clear why these generalizations should have a better claim to be meaning-constituting than the “paradigm cases” the sceptic rejects. At the very least, the sceptic owes us a fancy story here.

The moderate strict invariantist does not face this problem, since she takes many of our ordinary knowledge claims to be true. But she must explain why speakers find the premises exploited in sceptical arguments so compelling, despite the implausibility of the conclusions to which they lead. If these premises are false, why do speakers not come to see their falsity and stop feeling the pull of sceptical arguments? Presumably a moderate strict invariantist will say that I can sometimes know that my car is in the driveway, even though I have been gone for fifteen minutes and cannot absolutely rule out the possibility of car theft in the interim. Why, then, does the closure-exploiting argument that I cannot know this seem so compelling? These are deep and difficult questions, to be sure. My point here is that until she answers them satisfactorily, the moderate strict invariantist cannot explain away the apparent variability of standards in our knowledge attributions.

There is a further problem with both kinds of error theory, recently emphasized by Keith DeRose and John Hawthorne.¹⁷ Ordinary speakers accept many generalizations linking knowledge with other concepts. For example, one ought not assert something unless one knows it, one ought to decide what to do by reasoning from what one knows, and so on. The sceptical invariantist will have to hold that these generalizations, too, are in error, or else take the hard line that the vast majority of our assertions are improper and our decisions and actions irresponsible. The moderate strict invariantist will have trouble here, too, though less spectacularly, because in some situations (where much is at stake) we seem to require a very high standard of evidence before we will act on or assert a proposition. She must either

¹⁷DeRose 2002, Hawthorne 2004: 132–5.

say that our scruples here are unwarranted or reject the generalizations linking knowledge with assertion and action.

3.2.2 Embedded “know”

According to sensitive invariantism, the fact that speakers use the same epistemic standards in evaluating embedded and non-embedded instances of “know” reflects some kind of systematic error. But what kind? There are two possibilities. First, speakers might take the standards required to count as “knowing” to be fixed, or to be determined entirely by the context of use. Alternatively, instead of being mistaken about the semantics of “know,” speakers might systematically misjudge the standards in play at different circumstances of evaluation.

There is something a bit perverse about the first explanatory strategy. One of the best arguments *in favor of* a circumstance-variable, use-invariant semantics for “know” is that it promises to explain both the variability data and the data about truth ascriptions and retraction. But it cannot explain this data unless it plays some role in guiding speakers’ linguistic behavior. Thus, if we explain away the data about embedded occurrences by arguing that speakers implicitly take “know” to be circumstance-invariant and use it accordingly, we undercut one of the best arguments in favor of sensitive invariantism.

Better, then, to argue that speakers systematically misjudge the standards relevant at alternative circumstances of evaluation. Along these lines, John Hawthorne argues that we tend to “project” the standards currently in play to other putative knowers, times, and circumstances:

... we do have some tendency to suppose that, as more and more possibilities of error become salient to us, we are reaching an ever more enlightened perspective. Thus when we consider someone who is not alive to these possibilities,

we have a tendency to let our (putatively) more enlightened perspective trump his. This tendency, when left unchecked, leads to scepticism. (Hawthorne 2004: 164–5)

This kind of projection is not unprecedented: it is well known that those for whom a recent disaster is salient will overestimate risks in past, future, and counterfactual situations. In much the same way, Hawthorne urges,

Once we have gotten ourselves into the frame of mind of thinking ‘I do not in fact know whether or not I’ll be able to afford the Safari,’ as we frequently do when we use parity reasoning, we are not only unwilling to say ‘However I used to know that;’ we are positively willing to say ‘I never did know that.’ (162–3)

This strategy is worth pursuing, but we should remind ourselves how heavy an explanatory burden it must bear. It *always* seems wrong to say that Joe knew before, but doesn’t know now, when the only thing that has changed are the relevant standards. Projection might explain occasional or even frequent mistakes, but I doubt it can account for our universal unwillingness to shift standards across circumstances of evaluation.

Even if the projection strategy works, it is a double-edged sword. If it succeeds in explaining why we evaluate *embedded* occurrences of “know” in light of present standards, it should also explain why we evaluate occurrences of “know” at other *contexts of use* in light of present standards. That is, it should explain the data about truth ascriptions and retraction. Indeed, Hawthorne suggests as much himself, when he adds, immediately after the second passage quoted above: “And, if pressed, we are willing, moreover, to say that ‘I was mistaken in thinking that I did know that’” (163). The problem is that one of the best arguments for an invariantist semantics for “know” is that it explains the data about truth ascriptions and retraction. If that data is explained instead by the story about projection,

then the argument for preferring sensitive invariantism to contextualism is significantly weakened.

3.2.3 Truth ascriptions and retraction

The data about truth ascriptions and retraction is most straightforwardly explained by a use-invariant semantics for “know.” A contextualist must explain this data in some other way. We have ruled out a pragmatic explanation (§3.1.3), so it seems that a contextualist must appeal to an error theory here. Many contextualists are explicit about this: for example, Stewart Cohen (2001) says that “We *mistakenly* think that knowledge ascriptions we make in everyday contexts conflict with the skeptical judgements we make in stricter contexts” (89, emphasis added).¹⁸

As before, there are two options: the contextualist can suppose either that ordinary speakers are wrong about the semantics of “know”—treating it as use-invariant when it is not—or that they make systematic errors about what standards are in play in contexts other than their own. The problem is that both forms of error theory threaten to undermine the positive case for contextualism. This is especially clear if the error is semantic in character. If ordinary speakers have a faulty grasp of the *meaning* of “know,” then we cannot confidently appeal to variability in the standards they require someone to meet in order to count as “knowing” as support for a theory about the meaning of “know.” Yet this data is the primary evidence in favor of contextualism.

What about the second option? It is undeniable that speakers often misjudge features of

¹⁸Cohen argues that this error theory is innocuous, on the grounds that speakers make similar mistakes with gradable adjectives like “flat” (90–91). Richard 2004 concedes Cohen’s analogy and rejects his error theories, plumping for a relativist treatment of both “flat” and “know.” For my part, I am not convinced of the analogy: I think that a pragmatic explanation of our retraction and reporting behavior is much more plausible for “flat” than for “know.” When standards change so that the surface imperfections on pancakes count as “bumps” and “holes,” a speaker might retract an earlier assertion of “pancakes are flat,” but only to avoid pedantry, not because she thinks she’s really contradicted herself. If enough were at stake, she would no doubt find an appropriate way to reiterate her earlier claim. (Contrast what is alleged about “know” in §3.1.3, above.)

other contexts of use than their own, but if we are to explain the data, the error we posit must be *systematic*. We must explain why speakers *never* allow their previous day's assertion of "I know that *p*" to stand as true while asserting "I did not know that *p* yesterday." I doubt that our tendencies to project features of our present situations onto other situations are nearly strong or uniform enough to explain away the uniform data about truth ascriptions and retraction.

The "double-edged sword" point applies here, too. If the projection story works with contexts of use, it ought to work with circumstances of evaluation, too. So if it explains the data about truth ascriptions and retraction, it ought to explain the data about embedded occurrences of "know" as well. This would significantly weaken the contextualist's case against sensitive invariantism.

As should now be clear, a *general* problem with positing speaker error to explain away facts about use is that such explanations tend to undermine the evidential basis for the semantic theories they are intended to support. All of these semantic theories are justified indirectly on the basis of facts about speakers' use of sentences, and the more error we attribute to speakers, the less we can conclude from these facts. We have seen that the cost of defending sensitive invariantism in this way is that the case against contextualism is severely weakened, and conversely that the cost of defending contextualism in this way is that the case against sensitive invariantism is compromised. It is possible that an error theory can be made to work—perhaps in conjunction with pragmatic explanations—but the prospects do not look good.

3.3 Eliminativism

So far we have looked at ways of showing that one of the standard views is in fact consistent with all of the facts about use we considered in §2. An alternative response would be to

concede that no single account of the semantics of “know” accounts for all of these facts.¹⁹ Perhaps our talk of “knowledge” confuses several distinct notions, in much the same way that prescientific talk of “warmer than” confused *having a higher temperature than*, *having more heat energy than*, and *exchanging heat at a higher rate than*.²⁰ In that case there may be no fully coherent way to assign truth-conditions to our knowledge-attributing sentences. The rational course of action would be to reform our thought and talk by introducing new, unconfused terms of epistemic assessment.

At the risk of use-mention confusion, we might call this approach “eliminativism about knowledge.” Like other eliminativisms, it is radical and should not be accepted unless there is no other good alternative.

3.4 Expanding the field of options

Let us sum up our conclusions so far. Together, our three facts about use suggest that an adequate semantics for “know” must be sensitive to changing epistemic standards, but that it cannot be either use-variable or circumstance-variable. That rules out all three standard views: strict invariantism because it is not sensitive to changing epistemic standards at all, sensitive invariantism because it is circumstance-variable, and contextualism because it is use-variable. We might make room for one of these views by arguing that one of our three facts about use is a poor guide to truth-conditions, but attempts to do this either pragmatically or by positing systematic error on the part of ordinary speakers have so far been unpersuasive. If there is no other option, then, it seems we are left with eliminativism.

But how *could* there be another option? How could there be a semantics for “know” that was use-invariant and circumstance-invariant, but still in some way sensitive to changing epistemic standards? What we would need is another dimension of variability. In the next

¹⁹See Schiffer 1996.

²⁰For a discussion of this example, see Churchland 1979: ch. 2.

section, I am going to open up room for just such a thing. This will make possible a semantics for “know” that neatly explains *all three* facts about use.

4 A relativist semantics for “know”

Here is my proposal. The epistemic standards relevant to determining the extension of “know” are not those in play at the context of use or those in play at the circumstance of evaluation, but those in play at the *context of assessment*.

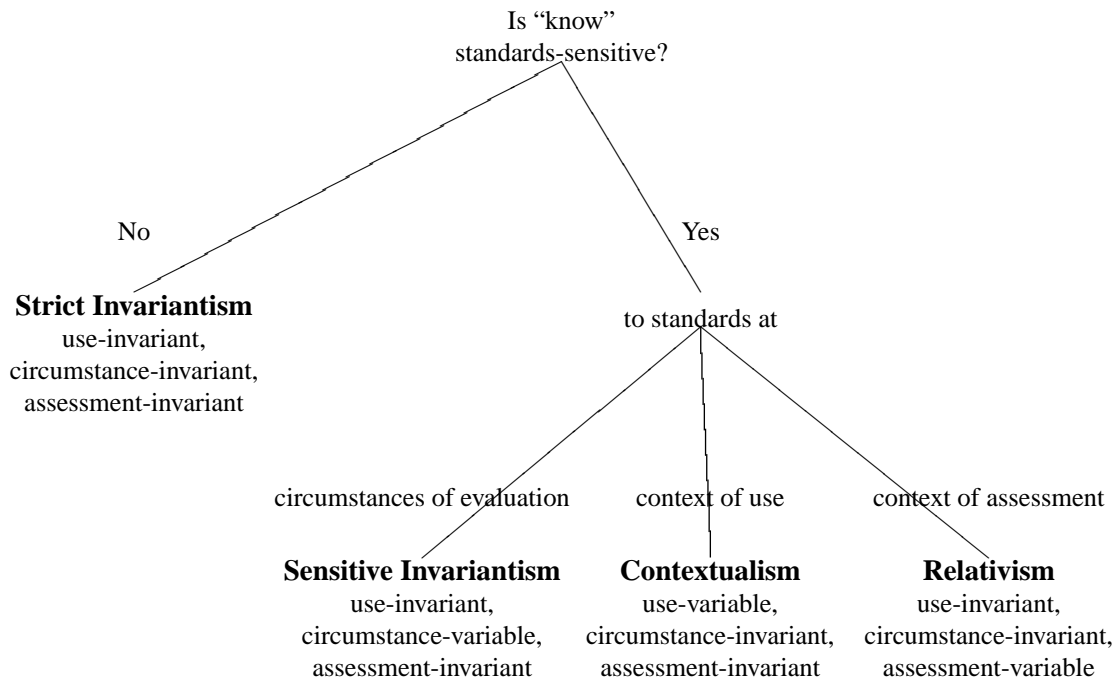
4.1 Assessment sensitivity

The notion of a context of assessment may be unfamiliar, but it is readily intelligible. Just as a context of use is a situation in which a sentence might be *used*, so a context of assessment is a situation in which a (past, present, or future, actual or merely possible) use of a sentence might be *assessed* for truth or falsity. I do not think that there should be any worries about the very *idea* of a context of assessment; even an arch anti-relativist ought to be able to accept it.

What is controversial is the suggestion that we relativize sentence truth not just to a context of use, but to a context of assessment as well. This is certainly a departure from semantic orthodoxy, and I will defend it shortly.²¹ Here I want to focus on what we can do with it. By making sentence truth doubly context-relative, we open up a new way in which sentences can be context-sensitive. A sentence is context-sensitive in the usual way, or *use-sensitive*, if its truth value varies with the context of use (keeping the context of assessment fixed). A sentence is context-sensitive in the new way, or *assessment-sensitive*,

²¹The relativization of truth to a context of assessment should not be confused with the relativization of truth to a “point of evaluation” (e.g., a tuple of time, world, and variable assignment) that is standard in model-theoretic semantics. A point of evaluation is not a context, but a sequence of parameters that can be “shifted” by operators. For more on the difference, see Lewis 1980 and MacFarlane 2003: §V.

Figure 2: Expanded taxonomy of positions on the semantics of “know”



if its truth value varies with the context of assessment (keeping the context of use fixed). Similarly, a subsentential expression is use-sensitive if it is partially responsible for the use sensitivity of (at least some) sentences containing it, and assessment-sensitive if it is partially responsible for the assessment sensitivity of (at least some) sentences containing it.

My proposal is that “know” is sensitive to the epistemic standards in play at the context of assessment. It is a kind of contextualism, then, but not at all the usual kind. To avoid confusion, I will call it “relativism,” reserving the term “contextualism” for the view that “know” is sensitive to the epistemic standards in play at the context of *use* (see Figure 2). Call a semantics for “know” *assessment-variable* just in case it allows the epistemic standard relevant for determining the extension of “know” to vary with the context of assessment, and *assessment-invariant* otherwise. If “know” is assessment-sensitive, then its semantics can be assessment-variable while being use- and circumstance-invariant, and in

this way we can neatly explain all three facts about use:

1. *Variability of standards.* Why is it that I'll happily assert "Joe knows that his car is parked in his driveway" when standards are low, and "Joe doesn't know that his car is parked in his driveway" when standards are high? The relativist semantics affords a simple explanation: the former sentence is true as used and assessed in a context where standards are low, and the latter is true as used and assessed in a context where standards are high. Because I can properly assess each sentence as true at the context in which I utter it, there is no need to appeal to pragmatic explanations or error theories to explain the variability data. (Note that in the special case where the context of use and the context of assessment coincide, the relativist semantics yields exactly the same truth-value assignments as the standard contextualist semantics. So the relativist is in just as good a position to explain the variability data as the contextualist. This is significant, because one of the primary selling points of the contextualist account is its ability to explain this data.)
2. *Embedded "know."* Why is it that we don't say things like "Before the standards went up, Harry knew that his car was in the driveway, but now he doesn't know this"? Or: "Harry doesn't know that his car is in the driveway, but he would know this if the possibility of car theft weren't relevant"? The relativist semantics has a straightforward explanation: the semantics of "know" is not circumstance-variable. This is the same explanation that contextualism and strict invariantism offer.
3. *Truth ascriptions and retraction.* Why is it that when standards go up, leading us to assert "Joe doesn't (and didn't) know that his car is in his driveway," we expect Joe to retract his earlier assertion of "I know that my car is in my driveway," and to concede that what he asserted was false? That is, why do we tend to use the standards appropriate to the present context in assessing past utterances? Where contextualism

is forced to appeal to an error theory, the relativist semantics offers an easy, semantic explanation. Namely: the present standards *are* the appropriate standards to use in assessing past assertions, even ones that were made when very different epistemic standards were in play. According to the relativist, knowledge claims are always properly assessed in light of the standards in play at the assessor's current context.²²

4.2 Expressive relativism and propositional relativism

Contextualists typically hold that “know” expresses different relations at different contexts of use, and that this indexicality is the source of its use sensitivity. Must the relativist, then, hold that “know” expresses different relations relative to different contexts of assessment? This would be an odd view. It would require us to give up the idea that “knowledge” attributers are making determinate claims. Assessors at different contexts could disagree about what was said, and they could all be right! According to this *expressive relativism*, there would be no non-relative fact of the matter about what proposition was expressed by the sentence used, at its context of use.²³

I go back and forth about the coherence of expressive relativism. It sometimes seems to me that with a bit of imagination, we can make sense of it. Even if we can make sense of it, however, it does not seem to be a very attractive view. It would require significant changes in orthodox theories of meaning. For example, we could no longer say, with Stalnaker 1978, that the effect of assertion is to add the proposition asserted to a “common ground” of presupposed propositions, for there may be no common fact of the matter about

²²A full discussion of the truth ascription data would require giving a semantics (at least a naive semantics) for the monadic object-language predicate “true.” I will not pause to do that here. It turns out, not surprisingly, that in a language containing assessment-sensitive expressions, object-language “true” must also be assessment sensitive.

²³Sometimes a single use of a sentence may express multiple propositions, as when a teacher says to a class of thirty: “Of the three people sitting nearest to you, only two are likely to finish this class.” I take it that in this case the teacher has asserted thirty singular propositions, not a single general one. This does not amount to expressive relativism, since all parties can agree about which propositions were expressed.

which proposition *was* asserted.²⁴ Moreover, although expressive relativism might help us understand *speech acts* made using “know,” it would leave it rather mysterious what it is to *believe* that Joe knows that his car is in his driveway.

Fortunately, assessment sensitivity can be had without expressive relativism. Call a sentence *use-indexical* if it expresses different propositions at different contexts of use (keeping the context of assessment fixed), and *assessment-indexical* if it expresses different propositions relative to different contexts of assessment (keeping the context of use fixed). Call a subsentential expression *use-* or *assessment-indexical* if it is at least partially responsible for the use or assessment indexicality of sentences containing it. According to expressive relativism, “know” is assessment-indexical, and that is why it is assessment-sensitive. But in fact a sentence or subsentential expression can be assessment-sensitive without being assessment-indexical.

Indeed, although this is often overlooked, a sentence can be *use-sensitive* without being use-indexical. Here is an example: “The number of AIDS babies born in the United States in 2003 is greater than 1000.” This sentence expresses the same proposition at every context of use, so it is not use-indexical. But it *is* use-sensitive, because its truth value varies with the world of the context of use. Uttered in a world in which there were no AIDS babies in 2003, it would express a falsehood; uttered in the actual world, it expresses a truth.²⁵

To see how a sentence can be use-sensitive without being use-indexical, we need to be explicit about the relation between sentence truth at a context and proposition truth at a circumstance of evaluation. (For simplicity, let us forget for a moment about contexts of assessment.)

SENTENCE TRUTH AND PROPOSITION TRUTH I: A sentence *S* is true at a

²⁴This is pointed out by Egan, Hawthorne, and Weatherson, forthcoming, who also give other arguments against what they call “content relativism.”

²⁵David Lewis put this point by saying that “[c]ontingency is a kind of indexicality” (Lewis 1980/1998: 25)—using “indexicality” for what I call “use sensitivity.” Other writers use “context sensitivity” for what I call “use indexicality.” I think it is useful to have distinct terms for both notions.

context of use C just in case for some proposition p ,

1. S expresses p at C , and
2. p is true when evaluated at the circumstances determined by C .²⁶

Notice that the context of use plays two distinct roles here: (1) it determines what proposition is expressed by the sentence, and (2) it determines how that proposition is to be evaluated to yield a truth value for the sentence in context. Indexicality produces use sensitivity via role (1), while contingency produces use sensitivity via role (2).

The situation is much the same when we relativize sentence truth to contexts of assessment as well as contexts of use:

SENTENCE TRUTH AND PROPOSITION TRUTH II: A sentence S is true at a context of use C_U and context of assessment C_A just in case for some proposition p ,

1. S expresses p at C_U and C_A , and
2. p is true when evaluated at the circumstances determined by C_U and C_A .

As before, there are two distinct roles for contexts to play: (1) determining which proposition is expressed and (2) determining how that proposition is to be evaluated to yield a truth value for the sentence in context. Accordingly, there are two ways in which a sentence can be assessment-sensitive: it can be assessment-indexical, or the context of assessment can play a substantive role in determining the circumstances relative to which the proposition

²⁶This definition is a close paraphrase of Kaplan 1989: 522 (cf. 547). As it stands, it is not sufficiently general, for in some frameworks the context of use will not always determine a unique circumstance of evaluation. For example, in indeterministic frameworks allowing overlapping worlds or “histories,” the context of use will not pick out a single history (see Belnap and Green 1994, MacFarlane 2003). For a more generally applicable definition, we could replace (2) with: “ p is true when evaluated at all circumstances of evaluation compatible with C ” (e.g., at all moment/history pairs in which the moment is the moment of C and the history contains the moment of C). We can ignore this complication for present purposes.

it expresses is to be evaluated. We can state the point more simply if we define proposition truth relative to contexts in the natural way:

CONTEXTS-RELATIVE PROPOSITION TRUTH: A proposition p is true at a context of use C_U and context of assessment C_A just in case p is true when evaluated at the circumstances determined by C_U and C_A .

Call a *proposition* assessment-sensitive just in case its truth varies with the context of assessment (keeping the context of use fixed). Then a sentence can be assessment-sensitive either by being assessment-indexical or by expressing an assessment-sensitive proposition. In the former case, we have expressive relativism; in the latter case, *propositional relativism*.

It might be thought that propositional relativism would require even more radical departures from orthodox semantics than expressive relativism. But that is not so. Granted, the form of propositional relativism I am advocating does require that the circumstances of evaluation to which proposition truth is relativized include an epistemic standards parameter in addition to a world parameter. But quite a few non-relativists have countenanced parameters of circumstances of evaluation besides the world parameter, so this hardly counts as a radical departure from standard semantic assumptions. For example, Kaplan's (1989) circumstances of evaluation include a time parameter. On Kaplan's account, a tensed sentence like "Socrates is sitting" expresses the same proposition at every context of use; it nonetheless has different truth values at different contexts of use, because different contexts determine different circumstances (times and worlds) with respect to which this proposition is to be evaluated. King, forthcoming, contemplates relativizing propositional truth to both worlds and standards of precision. On the view he considers (without endorsing or rejecting it), "France is hexagonal" expresses the same proposition at every context of use; it is true at a context of use just in case this proposition is true when evaluated with

respect to the world of the context of use and the standards of precision in play at the context of use. Despite their appeal to parameters of circumstances of evaluation besides worlds, neither Kaplan nor King is a propositional relativist, because neither countenances assessment-sensitive propositions.

Nor is there anything particularly novel about having an *epistemic standards* parameter in the circumstances of evaluation. The non-relativist form of contextualism about “knows” defended in Kompa 2002 requires one, too. On Kompa’s view, “know” expresses the same relation at every context of use, but this relation is “unspecific,” in the sense that “what counts as having the property” can vary with “the context at hand” (88)—that is, the context of use. Although Kompa does not develop her view in formal detail, it is hard to see how she could do so without adding an epistemic standards parameter to the circumstances of evaluation. The intension of the relation expressed by “know” would then be a function from worlds, times, and epistemic standards to extensions. This relation would be “unspecific” in the sense that there would be no answer to the question whether a particular person and fact fall into its extension at a time and world: only when an epistemic standard was specified would it have a definite extension. Kompa could then define context-relative sentence truth as in *Sentence Truth and Proposition Truth II*, above, taking the circumstances determined by a context of use C_U (and context of assessment C_A) to be $\langle w, t, e \rangle$, where w = the world of C_U , t = the time of C_U , and e = the epistemic standards in play at C_U . “Know” would thus turn out to be use-sensitive but not use-indexical, just like tense on Kaplan’s view, vague expressions on the view explored by King, and contingent eternal sentences on just about everyone’s view.

The only difference between the relativist view I am advocating and Kompa’s non-relativist, non-indexicalist form of contextualism is that I take the circumstances determined by a context of use C_U and context of assessment C_A to be the ordered pair $\langle w, t, e \rangle$, where w = the world of C_U , t = the time of C_U , and e = the epistemic standards in play at C_A

(not C_U).²⁷ In every other respect I can agree with Kompa. We can agree that proposition truth must be relativized to an epistemic standards parameter (in addition to a world and perhaps a time parameter). We can agree about which propositions are expressed by which sentences at which contexts of use. We can both accept the schematic principle *Sentence Truth and Proposition Truth II* (provided Kompa does not mind the relativization to a context of assessment, which plays no substantive role in her account but also does no harm). We will of course disagree about the extension of the relation “ S is true at context of use C_U and context of assessment C_A ,” because we disagree about how this schematic principle is to be filled in. (That is, we disagree about which circumstances of evaluation are “determined” by which contexts of use and assessment.) But this disagreement does not concern the theory of propositions. I conclude that if propositional relativism is objectionable, it is not because it requires radical revision to our existing theories of propositions.

5 Making sense of relative truth

I anticipate two objections to my proposal. First, that it is ad hoc. Good scientific practice dictates that we make central modifications to our theories only when they have great and wide-ranging explanatory value. Surely it is not a good idea to make structural changes to our semantic framework just to accommodate knowledge attributions. Second, that it is incoherent. It is one thing to talk of propositions or sentences being true with respect to one context of assessment, and not with respect to another. It is quite another thing to make sense of that talk, and there are reasons for doubting that any sense *can* be made of it.

²⁷I include the time parameter for illustrative purposes only; nothing hangs on its presence. King, forthcoming, may be right that the best treatment of tense does not call for a time parameter in circumstances of evaluation, in which case I am happy to remove it.

5.1 *Ad hoc?*

To the first objection I have two replies. First, I believe that assessment sensitivity is not limited to knowledge-attributing sentences. I believe it is also the key to adequate semantic treatments of future contingents, epistemic modals, accommodation (in the sense of Lewis 1979 and 1980), terms like “delicious,” and perhaps much else.²⁸ So the modification I propose is not tailor-made for a single use, but has much wider application.

Second, the structural changes that are required are less radical than one might think. As I have argued, propositional relativism requires minimal changes to existing semantic frameworks. These changes are conservative. They allow us to *describe* assessment sensitivity, but they leave open the possibility that there is no assessment sensitivity in any natural language. Existing accounts of the semantics of expressions that are not assessment-sensitive can be carried over essentially unchanged. (In these cases, the relativization to contexts of assessment will be an idle wheel, but a harmless one, because truth will not vary with the context of assessment.) Thus although one might object to the claim that “know” is assessment-sensitive, it is hard to see on what grounds one might object to the framework that makes it possible—unless one thinks that assessment-relative truth is simply incoherent.²⁹

5.2 *Incoherent?*

What on earth can it *mean* to say that an assertion is true as assessed by me, now, but false as assessed by me, later; or true as assessed by me, but false as assessed by you? This

²⁸I discuss future contingents in MacFarlane 2003 and the other issues in a book manuscript, in progress. For independent arguments for a relativist treatment of epistemic modals, see Egan, Hawthorne, and Weatherston, forthcoming. For a relativist treatment of accommodation (somewhat different from mine), see Richard 2004. For a relativism motivated by “delicious” and the like, see Kölbel 2002.

²⁹One might also worry that the extra degree of freedom I offer in constructing semantic theories does not come with sufficient counterbalancing constraints. The connections between contexts-relative truth and norms for assertion which I propose in the next section are meant to address this concern.

is not the kind of question that can be answered by *defining* “true at a context of use C_U and context of assessment C_A .” Indeed, we have already done that, for both sentences (p. 30) and propositions (p. 31).³⁰ But our definitions leave us in more or less the position of Dummett’s (1959) Martian anthropologists, who know what counts as a winning position in chess and other games, but fail to grasp the *significance* of winning. We don’t know what to *do* with a claim that a sentence (or proposition) is true relative to a context of use C_U and context of assessment C_A . And until we know that, we do not really understand relative-truth talk.

The charge of incoherence arises because a standard story about the significance of “true at a context of use C_U ” cannot be extended to “true at a context of use C_U and context of assessment C_A .” According to this story, truth is the internal aim of assertion. Of course, people may have all kinds of goals in making assertions—influencing others, showing off, giving directions, offering reassurance—and these goals may sometimes be better served by speaking falsely than by speaking the truth. But there is a sense in which a false assertion is always incorrect *qua* assertion, even if it succeeds in promoting these other goals. It may be useful to lie, but once your assertion has been shown to be false, you must withdraw it as mistaken. Dummett argues that the concept of truth gets its significance from this normative connection to the practice of assertion: just as it is part of the concept of winning a game that a player aims to win, so “it is part of the concept of truth that we aim at making true statements” (1959/1978: 2). But if our primary grip on the notion of truth comes from our understanding of it as the internal aim of assertion, then the idea that truth might be relativized to a context of assessment just looks incoherent.³¹ It does not make *sense* to aim to assert a proposition that is true at the context of use and the context of assessment, because there is no such thing as *the* context of assessment: each assertion

³⁰These definitions are of course schematic, but when the semantic details are filled in, they will determine an extension for “true at context of use C_U and context of assessment C_A .”

³¹For a development of this argument, drawing on Evans 1985, see Percival 1994: 196–8.

can be assessed from indefinitely many distinct contexts.

At this point relativists typically say that the aim of assertion is to assert something that is true relative to the context of use and the asserter's own current context of assessment (which will of course be identical with the context of use).³² But this only gives a significance to "true at C_U, C_A " for the special case where $C_U = C_A$. The relativist has not told us what to do with "true at C_U, C_A " where C_U and C_A are distinct. As a result, the anti-relativist might justly charge that the relativist's "true at C, C " is just a notational variant of her own "true at C ," and that "true at C_U, C_A " has not yet been given a sense when $C_U \neq C_A$.

In my view, the relativist should instead reject the whole idea of understanding truth as "the aim of assertion." This idea is pretty obscure anyway. Even if truth is *an* internal normative aim of assertion, it is certainly not the only such aim: it is also part of the practice of assertion that we strive to say what is relevant to the conversation at hand, and to say things that are appropriately justified (or on some accounts, known). Indeed, in his 1972 Postscript to "Truth," Dummett emphasizes that his talk of truth as the aim of assertion was intended as a placeholder for a more complex story about the role truth plays in our practice of assertion: "What has to be added to a truth-definition for the sentences of a language, if the notion of truth is to be explained, is a description of the linguistic activity of making assertions; and this is a task of enormous complexity" (Dummett 1978: 20).

Having given up the "aim of assertion" idea, what else might we say about the role truth plays in our practice of assertion? One plausible and widely accepted idea is that an assertion is a *commitment* to the truth of what is asserted.³³ To make an assertion—even an insincere or otherwise defective one—is, *inter alia*, to commit oneself to the truth of the proposition asserted (relative to its context of use).³⁴ But what is it to commit oneself

³²See Kölbel 2002: 125, Egan, Hawthorne, and Weatherson, forthcoming: 29.

³³See e.g. Searle 1979: 12.

³⁴*Inter alia*, because presumably asserting a proposition involves more than simply committing oneself to

to the truth of a proposition? How does one honor or violate such a commitment? Some philosophers seem to find “commitment to truth” intelligible without further analysis, but in my view, “commitment to [noun phrase]” is intelligible only when it can be glossed in terms of commitment to *do* something. For example, we can make sense of “being committed to Al Gore,” but only as meaning something like “being committed to *working for* (or perhaps *supporting*) Al Gore.” When no obvious agentive complement presents itself, we can’t make any sense of the deontic construction at all. What would it mean, for example, to be committed to the color of the sky, or to the texture of a damp rose petal?

So, in committing myself to the truth of a proposition at a context of use, what exactly am I committing myself to *doing* (or refraining from doing)? Well, suppose I assert “Jake is in Boston.” If you ask “How do you know?” or challenge my claim more directly, by giving reasons for thinking it false, then it seems to me that I have an *obligation* to respond, by giving adequate reasons for thinking that my claim was true, or perhaps by deferring to the person who told me. If I can’t discharge this obligation in a way that meets the challenge, I must “uncommit myself” by retracting my assertion. If I neither withdraw the assertion nor reply to the challenge, I am shirking an obligation I incur not *qua* moral agent or friend or member of polite society, but simply *qua* asserter.

These observations suggest an answer to our question. What I have committed myself to doing, in asserting that Jake is in Boston, is vindicating my claim when it is challenged.³⁵ There may be no specific sanction for failing to follow through on this commitment. But if I fail too blatantly or too frequently, others may stop treating me as a being that is capable of undertaking this kind of commitment. They may still take my utterances as expressions of my beliefs, as we take a dog’s excited tail wagging as an expression of its psychological

its truth. Plausibly, the commitment must be undertaken publicly, by means of an overt utterance; perhaps there are other conditions as well.

³⁵For this way of looking at assertoric commitment as a “conditional task responsibility” to vindicate a claim when it is challenged, see Brandom 1983 and chapter 3 of Brandom 1994. I do not develop the idea in quite the same way as Brandom, but I am much indebted to his work.

state. They may even regard my utterances, if found to be reliable, as useful bits of information. But they will be treating me as a measuring instrument, not as an asserter. They will not take me to be *committing myself* to the truth of anything.

If this is right, then we should understand the “commitment to truth” incurred by an assertion as follows:

ASSERTORIC COMMITMENT: In asserting that p at a context C_U , one commits oneself to providing adequate grounds for the truth of p (relative to C_U), in response to any appropriate challenge, or (when appropriate) to deferring this responsibility to another asserter on whose testimony one is relying. One can be released from this commitment only by withdrawing the assertion.³⁶

The principle is schematic along many dimensions: to make it less schematic, one would have to say something about what kinds of challenges count as “appropriate,” what grounds count as “adequate” responses to what kinds of challenges, and when it is appropriate to defer responsibility. I won’t attempt to do any of this here. What is important for our purposes is that this account can be extended in a very natural way to allow for assessment-relative truth. For whenever an assertion is challenged, there are always two relevant contexts: the context in which the assertion was originally made and the context in which the challenge must be met. A natural way to give significance to doubly context-relative truth, then, would be to say that what must be established when an assertion is challenged is truth relative to the original context of use and the asserter’s *current* context of assessment (at the time of the challenge):

³⁶Several philosophers have suggested to me that this account overgeneralizes, taking the norms of the seminar room to apply to assertions in general. It may be that ordinary asserters recognize no general obligation to justify their claims in the face of reasoned challenges (though it would not follow that they are not *bound by* such a norm). But even these sceptics ought to be able to accept a weaker norm requiring withdrawal of assertions that have been shown to be untrue (relative to the context of use). This would be enough for my purposes here.

ASSERTORIC COMMITMENT (DUAL CONTEXTS): In asserting that p at a context C_U , one commits oneself to providing adequate grounds for the truth of p (relative to C_U and one's current context of assessment), in response to any appropriate challenge, or (when appropriate) to deferring this responsibility to another asserter on whose testimony one is relying.³⁷ One can be released from this commitment only by withdrawing the assertion.³⁸

Note that although this account assumes that it makes sense to talk about contexts of assessment, it does not assume that propositional truth actually *varies* with the context of assessment. So non-relativists should be able to accept it, though for them the mention of "one's current context of assessment" will be an idle wheel. What we have, then, is a plausible story about the role of truth in our practice of assertion that gives a significance to talk of truth relative to a context of assessment, without prejudging the question whether we can actually assert anything whose truth is relative in this way. Indeed, this account gives us a way to test particular semantic hypotheses that make use of relative truth, by settling the *normative* consequences of these hypotheses.

We can now get a better feel for the assessment-sensitive semantics for "know" by examining its normative consequences. Suppose that Linda asserts, in a "low standards" context C_1 , that Joe knew on March 10 that his car was in his driveway. If the assertion is challenged at this point, Linda must defend it by showing that

(a) Joe's car was in his driveway on March 10, and

³⁷In speaking of "grounds for truth," I do not mean to imply that the justification must be explicitly semantic. One can give grounds for the truth of a proposition p relative to context of use C_{U1} and context of assessment C_A simply by asserting another proposition at C_{U2} whose truth relative to C_{U2} and C_A entails, or is evidence for, the truth of p relative to C_{U1} and C_A . (The grounds one offers can themselves be challenged, of course, as can their status as grounds.)

³⁸Those who accept only the weaker account of assertoric commitment in note 36, above, may modify it as follows to accommodate assessment-relative truth: one who asserts that p at C_U is obliged to withdraw this assertion in context of assessment C_A if p is shown to be false relative to C_U and C_A .

- (b) Joe’s epistemic position with respect to this fact was good enough on March 10 to meet the (low) standards in play at C_1 .³⁹

Suppose that a little while later, standards are raised. If Linda’s assertion is challenged in this new context, C_2 , she must defend it by showing that

- (a) Joe’s car was in his driveway on March 10, and
- (b) Joe’s epistemic position with respect to this fact was good enough on March 10 to meet the (higher) standards in play at C_2 .

In asserting that Joe knew on March 10 that his car was in his driveway, Linda takes on an open-ended commitment to show, whenever her assertion is (appropriately) challenged at a context C , that what she asserted is true by the standards in play at C —even if these standards are different from those that were in play when she made the assertion. If she lacks the resources to reply to a challenge, or if the challenge is unanswerable, then she is obliged to withdraw her assertion.

It seems to me that there is nothing incoherent about taking on such a commitment. Indeed, the facts about our use of “know” surveyed in §2, above, suggest that we implicitly take ourselves to be bound by just such a commitment whenever we attribute knowledge.

6 Conclusion

According to the “relativist” semantics I have proposed, the epistemic standards relevant for determining whether someone can be truly said to “know” something are determined by the context of assessment, not the context of use. Consequently, in assessing knowledge claims made at different contexts for truth or falsity, one need not keep track of the standards that

³⁹To simplify the exposition, I am ignoring the possibility of deferring to the word of another.

were in place at these contexts. The only relevant standards are the ones *now* in place.⁴⁰ This is why knowledge attributions can be reiterated and reported homophonically.

On this view, knowledge attributions are not as robustly objective as ordinary claims about the world. We must be prepared to withdraw a knowledge attribution if standards change, even if the subject's epistemic position is just as we thought it was. Relatedly, when we challenge others for having made false knowledge claims, we may be assessing them in light of standards higher than the ones they recognized when they made them. Isn't this unfair? Not unless retracting an assertion is always tantamount to admitting that the assertion was made irresponsibly: and of course it is not, even without assessment sensitivity in the picture. When standards rise, speakers withdraw their knowledge attributions and take them to have been false, but they needn't (and typically don't) take themselves to have acted irresponsibly in making them. One indication of this is that when standards fall again, they go right back to their old ways, rather than becoming more cautious in attributing knowledge. This is not so strange if we think of knowledge attributions as temporary record-keeping devices—tools for keeping track of a normative status keyed to ever-changing present circumstances—rather than straightforward statements of fact.

If I am right, then knowledge attributions made blamelessly and with full access to the relevant facts must sometimes be withdrawn as false. In my view, philosophers have been too quick to find this incoherent. Sceptics argue that we are right to withdraw our knowledge claims in the face of sceptical challenges; they conclude that these claims were not responsibly made in the first place. Dogmatists and contextualists argue that we are

⁴⁰Keith DeRose might call this view “single scoreboard semantics run amok” (see DeRose 2004). On this view, there is a “single scoreboard” not just for all parties to a single conversation, but for *all* uses of “know” as assessed from any one perspective. It seems to me that the very same arguments DeRose uses to support his view that the various parties to a conversation share a single scoreboard can be applied transtemporally to show that the semantics for “know” must be use-invariant. If we shouldn't ignore the fact that speakers in a single conversation take themselves to be contradicting and agreeing with each other in making knowledge claims, then we shouldn't ignore the fact that speakers take present knowledge claims to contradict or agree with past ones, even ones made when different standards were in play.

wrong to withdraw our knowledge claims, precisely because they were responsibly made. I say that both sides have part of the truth: we are right to withdraw our knowledge claims in the face of certain sceptical challenges, even though they were responsibly made and we haven't learned anything new. A relativist semantics for "know" allows us to understand how this can be.

References

- Belnap, N. and M. Green (1994). Indeterminism and the thin red line. *Philosophical Perspectives* 8, 365–388.
- Brandom, R. (1983). Asserting. *Noûs* 17, 637–650.
- Brandom, R. (1994). *Making It Explicit*. Cambridge: Harvard University Press.
- Churchland, P. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.
- Cohen, S. (2001). Contextualism defended: Comments on Richard Feldman's 'Skeptical problems, contextualist solutions'. *Philosophical Studies* 103, 87–98.
- DeRose, K. (2000). Now you know it, now you don't. In *Proceedings of the Twentieth World Congress of Philosophy, Volume V: Epistemology*, Bowling Green, Ohio, pp. 91–106. Philosophy Documentation Center.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review* 111, 167–203.
- DeRose, K. (2004). Single scoreboard semantics. *Philosophical Studies* 119, 1–21.
- DeRose, K. (forthcoming). The problem with subject-sensitive invariantism. *Philosophy and Phenomenological Research*.

- Dummett, M. (1959). Truth. *Proceedings of the Aristotelian Society n.s.* 59, 141–62.
- Dummett, M. (1978). *Truth and Other Enigmas*. Cambridge: Harvard University Press.
- Egan, A., J. Hawthorne, and B. Weatherson (forthcoming). Epistemic modals in context. In G. Preyer and P. Peter (Eds.), *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Evans, G. (1985). Does tense logic rest upon a mistake? In *Collected Papers*, pp. 343–363. Oxford: Oxford University Press.
- Feldman, R. (2001). Skeptical problems, contextualist solutions. *Philosophical Studies* 103, 61–85.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan*, pp. 481–566. Oxford: Oxford University Press.
- King, J. (forthcoming). Tense, modality, and semantic value. *Philosophical Perspectives*.
- Kölbel, M. (2002). *Truth Without Objectivity*. London: Routledge.
- Kompa, N. (2002). The context sensitivity of knowledge ascriptions. *Grazer Philosophische Studien* 64, 79–96.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic* 8, 339–359.
- Lewis, D. (1980). Index, context, and content. In S. Kanger and S. Öhman (Eds.), *Philosophy and Grammar*. Dordrecht: Reidel.

- Lewis, D. (1998). *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.
- MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly* 53(212), 321–336.
- Percival, P. (1994). Absolute truth. *Proceedings of the Aristotelian Society* 94, 189–213.
- Richard, M. (2004). Contextualism and relativism. *Philosophical Studies* 119, 215–242.
- Rosenberg, J. F. (2002). *Thinking About Knowing*. Oxford University Press.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs* 35(4), 477–514.
- Schaffer, J. (forthcoming). Skepticism, contextualism, and discrimination. *Philosophy and Phenomenological Research*.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society* 96, 317–333.
- Searle, J. (1969). *Speech Acts*. Cambridge: Cambridge University Press.
- Searle, J. R. (1979). *Expression and Meaning*. Cambridge: Cambridge University Press.
- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics, Vol. 9: Pragmatics*. New York: Academic Press.
- Stanley, J. (2004). On the linguistic basis for contextualism. *Philosophical Studies* 119, 119–146.